OpenDaylight Documentation Documentation

Release Boron

OpenDaylight Project

Jun 04, 2018

Contents

1 Content for OpenDaylight Users 3

2 Content for OpenDaylight Developers 781

3 Content for OpenDaylight Contributors 1077

This handbook provides details on various aspects of OpenDaylight from the user guides to the developer guides and tries to act as a single point of contact for all documentation related articles in OpenDaylight. If you would like to contribute to the Handbook please refer to the *Documentation Guide*.

CHAPTER 1

Content for OpenDaylight Users

The following content is intended for people who would like to deploy, use, or just learn more about OpenDaylight.

1.1 Getting Started Guide

1.1.1 Introduction

The OpenDaylight project is an open source platform for Software Defined Networking (SDN) that uses open protocols to provide centralized, programmatic control and network device monitoring. Like many other SDN controllers, OpenDaylight supports OpenFlow, as well as offering ready-to-install network solutions as part of its platform.

Much as your operating system provides an interface for the devices that comprise your computer, OpenDaylight provides an interface that allows you to connect network devices quickly and intelligently for optimal network performance.

It's extremely helpful to understand that setting up your networking environment with OpenDaylight is not a single software installation. While your first chronological step is to install OpenDaylight, you install additional functionality packaged as Karaf features to suit your specific needs.

Before walking you through the initial OpenDaylight installation, this guide presents a fuller picture of OpenDaylight's framework and functionality so you understand how to set up your networking environment. The guide then takes you through the installation process.

What's different about OpenDaylight

Major distinctions of OpenDaylight's SDN compared to traditional SDN options are the following:

- A microservices architecture, in which a "microservice" is a particular protocol or service that a user wants to enable within their installation of the OpenDaylight controller, for example:
  - A plugin that provides connectivity to devices via the OpenFlow or BGP protocols
  - An L2-Switch or a service such as Authentication, Authorization, and Accounting (AAA).

- Support for a wide and growing range of network protocols beyond OpenFlow, including SNMP, NETCONF, OVSDB, BGP, PCEP, LISP, and more.
- Support for developing new functionality comprised of additional networking protocols and services.

Note: A thorough understanding of the microservices architecture is important for experienced network developers who want to create new solutions in OpenDaylight. If you are new to networking and OpenDaylight, you most likely won't design solutions, but you should comprehend the microservices concept to understand how OpenDaylight works and how it differs from other SDN programs.

What you'll find in this guide

To set up your environment, you first install OpenDaylight followed by the Apache Karaf features that offer the functionality you require. The OpenDaylight Getting Started Guide covers feature descriptions, OpenDaylight installation procedures, and feature installation.

The Getting Started Guide also includes other helpful information, with the following organization:

1. An overview of OpenDaylight and common use models
2. Who should use this guide?
3. OpenDaylight concepts and tools
4. Explanations of OpenDaylight Apache Karaf features and other features that extend network functionality
5. OpenDaylight system requirements and Release Notes
6. OpenDaylight installation instructions
7. Feature tables with installation names and compatibility notes

1.1.2 Overview

OpenDaylight performs the following functions:

- Logically centralizes programmatic control of the physical and virtual devices in your network.
- Controls devices with standard, open protocols.
- Provides higher-level abstractions of its capabilities so experienced network engineers and developers can create new applications to customize network setup and administration.

Common use cases for SDN are as follows:

1. Centralized network monitoring, management, and orchestration
2. Proactive network management and traffic engineering
3. Chaining packets through the different VMs, which is known as service function chaining (SFC). SFC enables Network Functions Virtualization (NFV), which is a network architecture concept that virtualizes entire classes of network node functions into building blocks that may connect, or chain together, to create communication services.
4. Cloud - managing both the virtual overlay and the physical underlay beneath it.

1.1.3 Who should use this guide?

OpenDaylight is for users considering open options in network programming. This guide provides information for the following types of users:

1. Those new to OpenDaylight who want to install it and select the features they need to run their network environment using only the command line and GUI. Such users include:
   (a) Students
   (b) Network administrators and engineers.
2. Network engineers and network application developers who want to use OpenDaylight's REST APIs to manage their network programmatically.
3. Network engineers and network application developers who want to write their own OpenDaylight services and plugins for greater functionality. This group of users needs a significant level of expertise in the following areas, which is beyond the scope of this document:
   (a) The YANG modeling language
   (b) The Model-Driven Service Abstraction Layer (MD-SAL)
   (c) Maven build tool
   (d) Management of the shared data store
   (e) How to handle notifications and/or Remote Procedure Calls (RPCs)
4. Developers who would like to join the OpenDaylight community and contribute code upstream. People in this group design offerings such as applications/services, protocol implementations, and so on, to increase OpenDaylight functionality for the benefit of all end-users.

Note: If you develop code to build new functionality for OpenDaylight and push it upstream (not required), it can become part of the OpenDaylight release. Users can then install the features to implement the solution you've created.

1.1.4 OpenDaylight concepts and tools

In this section we discuss some of the concepts and tools you encounter with basic use of OpenDaylight. The guide walks you through the installation process in a subsequent section, but for now familiarize yourself with the information below.

- To date, OpenDaylight developers have formed more than 50 projects to address ways to extend network functionality. The projects are a formal structure for developers from the community to meet, document release plans, code, and release the functionality they create in an OpenDaylight release.

  The typical OpenDaylight user will not join a project team, but you should know what projects are as we refer to their activities and the functionality they create. The Karaf features to install that functionality often share the project team's name.

- **Apache Karaf** provides a lightweight runtime to install the Karaf features you want to implement and is included in the OpenDaylight platform software. By default, OpenDaylight has no pre-installed features.

- After installing OpenDaylight, you install your selected features using the Karaf console to expand networking capabilities. In the Karaf feature list below are the ones you're most likely to use when creating your network environment.

  As a short example of installing a Karaf feature, OpenDaylight offers Application Layer Traffic Optimization (ALTO). The Karaf feature to install ALTO is odl-alto-all. On the Karaf console, the command to install it is:

feature:install odl-alto-all

- **DLUX** is a web-based interface that OpenDaylight provides for you to manage your network. Its Karaf feature installation name is "odl-dlux-core".

  1. DLUX draws information from OpenDaylight's topology and host databases to display the following information:

     (a) The network

     (b) Flow statistics

     (c) Host locations

  2. To enable the DLUX UI after installing OpenDaylight, run:

     feature:install odl-dlux-core

     on the Karaf console.

- **Network embedded Experience (NeXt)** is a developer toolkit that provides tools to draw network-centric topology UI elements that offer visualizations of the following:

  1. Large complex network topologies
  2. Aggregated network nodes
  3. Traffic/path/tunnel/group visualizations
  4. Different layout algorithms
  5. Map overlays
  6. Preset user-friendly interactions

  NeXt can work with DLUX to build OpenDaylight applications. Check out the NeXt_demo for more information on the interface.

- Model-Driven Service Abstraction Layer (MD-SAL) is the OpenDaylight framework that allows developers to create new Karaf features in the form of services and protocol drivers and connects them to one another. You can think of the MD-SAL as having the following two components:

  1. A shared datastore that maintains the following tree-based structures:

  1. The Config Datastore, which maintains a representation of the desired network state.
  2. The Operational Datastore, which is a representation of the actual network state based on data from the managed network elements.

  2. A message bus that provides a way for the various services and protocol drivers to notify and communicate with one another.

- If you're interacting with OpenDaylight through DLUX or the REST APIs while using the the OpenDaylight interfaces, the microservices architecture allows you to select available services, protocols, and REST APIs.

1.1.5 OpenDaylight Karaf Features

This section provides brief descriptions of the most commonly used Karaf features developed by OpenDaylight project teams. They are presented in alphabetical order. OpenDaylight installation instructions and a feature table that lists installation commands and compatibility follow.

- *AAA*
- *ALTO*
- *Border Gateway Protocol (including Link-state Distribution (BGP)*
- *Border Gateway Monitoring Protocol (BMP)*
- *Control and Provisioning of Wireless Access Points (CAPWAP)*
- *Controller Shield*
- *Device Identification and Driver Management (DIDM)*
- *DLUX*
- *Fabric as a Service (FaaS)*
- *Group Based Policy (GBP)*
- *Internet of Things Data Management (IoTDM)*
- *Link Aggregation Control Protocol (LACP)*
- *Location Identifier Separation Protocol (LISP) Flow Mapping Service (LISP)*
- *NEMO*
- *NETCONF*
- *NetIDE*
- *OVSDB-based Network Virtualization Services*
- *OpenFlow Configuration Protocol (OF-CONFIG)*
- *OpenFlow plugin*
- *Path Computation Element Protocol (PCEP)*
- *Secure Network Bootstrapping Interface (SNBi)*
- *Service Function Chaining (SFC)*
- *SNMP Plugin*
- *SNMP4SDN*
- *Source-Group Tag Exchange Protocol (SXP)*
- *Topology Processing Framework*
- *Time Series Data Repository (TSDR)*
- *Unified Secure Channel (USC)*
- *Virtual Tenant Network (VTN)*

AAA

Standards-compliant Authentication, Authorization and Accounting Services. RESTCONF is the most common consumer of AAA, which installs the AAA features automatically. AAA provides:

- Support for persistent data stores
- Federation and SSO with OpenStack Keystone

This release of AAA includes experimental support for having the database of users and credentials stored in the cluster-aware MD-SAL datastore.

ALTO

Implements the Application-Layer Traffic Optimization (ALTO) base IETF protocol to provide network information to applications. It defines abstractions and services to enable simplified network views and network services to guide application usage of network resources and includes five services:

1. Network Map Service - Provides batch information to ALTO clients in the forms of ALTO network maps.
2. Cost Map Service - Provides costs between defined groupings.
3. Filtered Map Service - Allows ALTO clients to query an ALTO server on ALTO network maps and/or cost maps based on additional parameters.
4. Endpoint Property Service - Allows ALTO clients to look up properties for individual endpoints.
5. Endpoint Cost Service - Allows an ALTO server to return costs directly amongst endpoints.

Border Gateway Protocol (including Link-state Distribution (BGP)

Is a southbound plugin that provides support for Border Gateway Protocol (including Link-state Distribution) as a source of L3 topology information.

Border Gateway Monitoring Protocol (BMP)

Is a southbound plugin that provides support for BGP Monitoring Protocol as a monitoring station.

Control and Provisioning of Wireless Access Points (CAPWAP)

Enables OpenDaylight to manage CAPWAP-compliant wireless termination point (WTP) network devices. Intelligent applications, e.g., radio planning, can be developed by tapping into the operational states made available via REST APIs of WTP network devices.

Controller Shield

Creates a repository called the Unified-Security Plugin (USecPlugin) to provide controller security information to northbound applications, such as the following:

- Collating the source of different attacks reported in southbound plugins
- Gathering information on suspected controller intrusions and trusted controllers in the network

Information collected at the plugin may also be used to configure firewalls and create IP blacklists for the network.

Device Identification and Driver Management (DIDM)

Provides device-specific functionality, which means that code enabling a feature understands the capability and limitations of the device it runs on. For example, configuring VLANs and adjusting FlowMods are features, and there may be different implementations for different device types. Device-specific functionality is implemented as Device Drivers.

DLUX

Web based OpenDaylight user interface that includes:

- An MD-SAL flow viewer
- Network topology visualizer
- A tool box and YANG model that execute queries and visualize the YANG tree

Fabric as a Service (FaaS)

Creates a common abstraction layer on top of a physical network so northbound APIs or services can be more easily mapped onto the physical network as a concrete device configuration.

Group Based Policy (GBP)

Defines an application-centric policy model for OpenDaylight that separates information about application connectivity requirements from information about the underlying details of the network infrastructure. Provides support for:

- Integration with OpenStack Neutron
- Service Function Chaining
- OFOverlay support for NAT, table offsets

Internet of Things Data Management (IoTDM)

Developing a data-centric middleware to act as a oneM2M-compliant IoT Data Broker (IoTDB) and enable authorized applications to retrieve IoT data uploaded by any device.

Link Aggregation Control Protocol (LACP)

LACP can auto-discover and aggregate multiple links between an OpenDaylight-controlled network and LACP-enabled endpoints or switches.

Location Identifier Separation Protocol (LISP) Flow Mapping Service (LISP)

LISP (RFC6830) enables separation of Endpoint Identity (EID) from Routing Location (RLOC) by defining an overlay in the EID space, which is mapped to the underlying network in the RLOC space.

LISP Mapping Service provides the EID-to-RLOC mapping information, including forwarding policy (load balancing, traffic engineering, and so on) to LISP routers for tunneling and forwarding purposes. The LISP Mapping Service can serve the mapping data to data plane nodes as well as to OpenDaylight applications.

To leverage this service, a northbound API allows OpenDaylight applications and services to define the mappings and policies in the LISP Mapping Service. A southbound LISP plugin enables LISP data plane devices to interact with OpenDaylight via the LISP protocol.

NEMO

Is a Domain Specific Language (DSL) for the abstraction of network models and identification of operation patterns. NEMO enables network users/applications to describe their demands for network resources, services, and logical operations in an intuitive way that can be explained and executed by a language engine.

NETCONF

Offers four features:

- odl-netconf-mdsal: NETCONF Northbound for MD-SAL and applications
- odl-netconf-connector: NETCONF Southbound plugin - configured through the configuration subsystem
- odl-netconf-topology: NETCONF Southbound plugin - configured through the MD-SAL configuration datastore
- odl-restconf: RESTCONF Northbound for MD-SAL and applications

NetIDE

Enables portability and cooperation inside a single network by using a client/server multi-controller architecture. It provides an interoperability layer allowing SDN Applications written for other SDN Controllers to run on OpenDaylight. NetIDE details:

- Architecture follows a client/server model: other SDN controllers represent clients with OpenDaylight acting as the server.
- OpenFlow v1.0/v1.3 is the only southbound protocol supported in this initial release. We are planning for other southbound protocols in later releases.
- The developer documentation contains the protocol specifications required for developing plugins for other client SDN controllers.
- The NetIDE Configuration file contains the configurable elements for the engine.

OVSDB-based Network Virtualization Services

Several services and plugins in OpenDaylight work together to provide simplified integration with the OpenStack Neutron framework. These services enable OpenStack to offload network processing to OpenDaylight while enabling OpenDaylight to provide enhanced network services to OpenStack.

OVSDB Services are at parity with the Neutron Reference Implementation in OpenStack, including support for:

- L2/L3
  - The OpenDaylight Layer-3 Distributed Virtual Router is fully on par with what OpenStack offers and now provides completely decentralized Layer 3 routing for OpenStack. ICMP rules for responding on behalf of the L3 router are fully distributed as well.
  - Full support for distributed Layer-2 switching and distributed IPv4 routing is now available.
- Clustering - Full support for clustering and High Availability (HA) is available in the this OpenDaylight release. In particular, the OVSDB southbound plugin supports clustering that any application can use, and the Openstack network integration with OpenDaylight (through OVSDB Net-Virt) has full clustering support. While there is no specific limit on cluster size, a 3-node cluster has been tested extensively as part of the release.

- Security Groups - Security Group support is available and implemented using OpenFlow rules that provide superior functionality and performance over OpenStack Security Groups, which use IPTables. Security Groups also provide support for ConnTrack with stateful tracking of existing connections. Contract-based Security Groups require OVS v2.5 with contract support.
- Hardware Virtual Tunnel End Point (HW-VTEP) - Full HW-VTEP schema support has been implemented in the OVSDB protocol driver. Support for HW-VTEP via OpenStack through the OVSDB-NetVirt implementation has not yet been provided as we wait for full support of Layer-2 Gateway (L2GW) to be implemented within OpenStack.
- Service Function Chaining
- Open vSwitch southbound support for quality of service and Queue configuration Load Balancer as service (LBaaS) with Distributed Virtual Router
- Network Virtualization User interface for DLUX

OpenFlow Configuration Protocol (OF-CONFIG)

Provides a process for an Operation Context containing an OpenFlow Switch that uses OF-CONFIG to communicate with an OpenFlow Configuration Point, enabling remote configuration of OpenFlow datapaths.

OpenFlow plugin

Supports connecting to OpenFlow-enabled network devices via the OpenFlow specification. It currently supports OpenFlow versions 1.0 and 1.3.2.

In addition to support for the core OpenFlow specification, OpenDaylight also includes preliminary support for the Table Type Patterns and OF-CONFIG specifications.

Path Computation Element Protocol (PCEP)

Is a southbound plugin that provides support for performing Create, Read, Update, and Delete (CRUD) operations on Multiprotocol Label Switching (MPLS) tunnels in the underlying network.

Secure Network Bootstrapping Interface (SNBi)

Leverages manufacturer-installed IEEE 802.1AR certificates to secure initial communications for a zero-touch approach to bootstrapping using Docker. SNBi devices and controllers automatically do the following:

1. Discover each other, which includes:
   (a) Revealing the physical topology of the network
   (b) Exposing each type of a device
   (c) Assigning the domain for each device
2. Get assigned an IP-address
3. Establish secure IP connectivity

SNBi creates a basic infrastructure to host, run, and lifecycle-manage multiple network functions within a network device, including individual network element services, such as:

- Performance measurement
- Traffic-sniffing functionality

- Traffic transformation functionality

SNBi also provides a Linux side abstraction layer to forward elements as well as enhancements to feature the abstraction and bootstrapping infrastructure. You can also use the device type and domain information to initiate controller federation processes.

Service Function Chaining (SFC)

Provides the ability to define an ordered list of network services (e.g. firewalls, load balancers) that are then "stitched" together in the network to create a service chain. SFC provides the chaining logic and APIs necessary for OpenDaylight to provision a service chain in the network and an end-user application for defining such chains. It includes:

- YANG models to express service function chains
- SFC receiver for Intent expressions from REST & RPC
- UI for service chain construction
- LISP support
- Function grouping for load balancing
- OpenFlow renderer for Network Service Headers, MPLS, and VLAN
- Southbound REST interface
- IP Tables-based classifier for grouping packets into selected service chains
- Integration with OpenDaylight GBP project
- Integration with OpenDaylight OVSDB NetVirt project

SNMP Plugin

The SNMP southbound plugin allows applications acting as an SNMP Manager to interact with devices that support an SNMP agent. The SNMP plugin implements a general SNMP implementation, which differs from the SNMP4SDN as that project leverages only select SNMP features to implement the specific use case of making an SNMP-enabled device emulate some features of an OpenFlow-enabled device.

SNMP4SDN

Provides a southbound SNMP plugin to optimize delivery of SDN controller benefits to traditional/legacy ethernet switches through the SNMP interface. It offers support for flow configuration on ACLs and enables flow configuration via REST API and multi-vendor support.

Source-Group Tag Exchange Protocol (SXP)

Enables creation of a tag that allows you to filter traffic instead of using protocol-specific information like addresses and ports. Via SXP an external entity creates the tags, assigns them to traffic appropriately, and publishes information about the tags to network devices so they can enforce the tags appropriately.

More specifically, SXP Is an IETF-published control protocol designed to propagate the binding between an IP address and a source group, which has a unique source group tag (SGT). Within the SXP protocol, source groups with common network policies are endpoints connecting to the network. SXP updates the firewall with SGTs, enabling the firewalls to create topology-independent Access Control Lists (ACLs) and provide ACL automation.

SXP source groups have the same meaning as endpoint groups in OpenDaylight's Group Based Policy (GBP), which is used to manipulate policy groups, so you can use OpenDaylight GPB with SXP SGTs. The SXP topology-independent policy definition and automation can be extended through OpenDaylight for other services and networking devices.

Topology Processing Framework

Provides a framework for simplified aggregation and topology data query to enable a unified topology view, including multi-protocol, Underlay, and Overlay resources.

Time Series Data Repository (TSDR)

Creates a framework for collecting, storing, querying, and maintaining time series data in OpenDaylight. You can leverage various data-driven applications built on top of TSDR when you install a datastore and at least one collector.

Functionality of TDSR includes:

- Data Query Service - For external data-driven applications to query data from TSDR through REST APIs
- NBI integration with Grafana - Allows visualization of data collected in TSDR using Grafana
- Data Aggregation Service - Periodically aggregates raw data into larger time granularities
- Data Purging Service - Periodically purges data from TSDR
- Data Collection Framework - Data Collection framework to allow plugging in of various types of collectors
- HSQL data store - Replacement of H2 data store to remove third party component dependency from TSDR
- Cassandra data store - Cassandra implementation of TSDR SPIs
- NetFlow data collector - Collect NetFlow data from network elements
- NetFlowV9 - version 9 Netflow collector
- SNMP Data Collector - Integrates with SNMP plugin to bring SNMP data into TSDR
- sFlowCollector - Collects sFlow data from network elements
- Syslog data collector - Collects syslog data from network elements

TSDR has multiple features to enable the functionality above. To begin, select one of these data stores:

- odl-tsdr-hsqldb-all
- odl-tsdr-hbase
- odl-tsdr-cassandra

Then select any "collectors" you want to use:

- odl-tsdr-openflow-statistics-collector
- odl-tsdr-netflow-statistics-collector
- odl-tsdr-controller-metrics-collector
- odl-tsdr-sflow-statistics-collector
- odl-tsdr-snmp-data-collector
- odl-tsdr-syslog-collector

See these TSDR_Directions for more information.

Unified Secure Channel (USC)

Provides a central server to coordinate encrypted communications between endpoints. Its client-side agent informs the controller about its encryption capabilities and can be instructed to encrypt select flows based on business policies.

A possible use case is encrypting controller-to-controller communications; however, the framework is very flexible, and client side software is available for multiple platforms and device types, enabling USC and OpenDaylight to centralize the coordination of encryption across a wide array of endpoint and device types.

Virtual Tenant Network (VTN)

Provides multi-tenant virtual network on an SDN controller, allowing you to define the network with a look and feel of a conventional L2/L3 network. Once the network is designed on VTN, it automatically maps into the underlying physical network and is then configured on the individual switch, leveraging the SDN control protocol.

By defining a logical plane with VTN, you can conceal the complexity of the underlying network and better manage network resources to reduce network configuration time and errors.

1.1.6 OpenDaylight Experimental Features

- *Network Intent Composition (NIC)*
- *UNI Manager Plug-in (Unimgr)*
- *YANG-PUBSUB*

Network Intent Composition (NIC)

Offers an interface with an abstraction layer for you to communicate "intentions," i.e., what you expect from the network. The Intent model, which is part of NIC's core architecture, describes your networking services requirements and transforms the details of the desired state to OpenDaylight. NIC has four features:

- odl-nic-core-hazelcast: Provides the following:
  - A distributed intent mapping service implemented using hazelcast, which stores metadata needed to process Intent correctly
  - An intent REST API to external applications for Create, Read, Update, and Delete (CRUD) operations on intents, conflict resolution, and event handling
- odl-nic-core-mdsal: Provides the following:
  - A distributed Intent mapping service implemented using MD-SAL, which stores metadata needed to process Intent correctly
  - An Intent rest API to external applications for CRUD operations on Intents, conflict resolution, and event handling
- odl-nic-console: Provides a Karaf CLI extension for Intent CRUD operations and mapping service operations
- Four renderers to provide specific implementations to render the Intent:
  - Virtual Tenant Network Renderer
  - Group Based Policy Renderer
  - OpenFlow Renderer

- Network MOdeling Renderer

UNI Manager Plug-in (Unimgr)

Formed to initiate the development of data models and APIs that facilitate OpenDaylight software applications' and/or service orchestrators' ability to configure and provision connectivity services.

YANG-PUBSUB

An experimental feature Plugin that allows subscriptions to be placed on targeted subtrees of YANG datastores residing on remote devices. Changes in YANG objects within the remote subtree can be pushed to OpenDaylight as specified and don't require OpenDaylight to make continuous fetch requests. YANG-PUBSUB is developed as a Java project. Development requires Maven version 3.1.1 or later.

1.1.7 Other features

OpFlex

Provides the OpenDaylight OpFlex Agent , which is a policy agent that works with Open vSwitch (OVS), to enforce network policy, e.g., from Group-Based Policy, for locally-attached virtual machines or containers.

Network embedded Experience (NeXt)

Provides a network-centric topology UI that offers visualizations of the following:

1. Large complex network topologies
2. Aggregated network nodes
3. Traffic/path/tunnel/group visualizations
4. Different layout algorithms
5. Map overlays
6. Preset user-friendly interactions

NeXt can work with DLUX to build OpenDaylight applications. NeXt does not support Internet Explorer. Check out the NeXt_demo for more information on the interface.

1.1.8 API

We are in the process of creating automatically generated API documentation for all of OpenDaylight. The following are links to the preliminary documentation that you can reference. We will continue to add more API documentation as it becomes available.

- mdsal
- odlparent
- yangtools

1.1.9 Installing OpenDaylight

You complete the following steps to install your networking environment, with specific instructions provided in the subsections below.

Before detailing the instructions for these, we address the following: Java Runtime Environment (JRE) and operating system information Target environment Known issues and limitations

Install OpenDaylight

Downloading and installing OpenDaylight

The default distribution can be found on the OpenDaylight software download page: http://www.opendaylight.org/software/downloads

The Karaf distribution has no features enabled by default. However, all of the features are available to be installed.

Note: For compatibility reasons, you cannot enable all the features simultaneously. We try to document known incompatibilities in the *Install the Karaf features* section below.

Running the karaf distribution

To run the Karaf distribution:

1. Unzip the zip file.
2. Navigate to the directory.
3. run `./bin/karaf`.

For Example:

```
$ ls distribution-karaf-0.5.x-Boron.zip
distribution-karaf-0.5.x-Boron.zip
$ unzip distribution-karaf-0.5.x-Boron.zip
Archive:  distribution-karaf-0.5.x-Boron.zip
   creating: distribution-karaf-0.5.x-Boron/
   creating: distribution-karaf-0.5.x-Boron/configuration/
   creating: distribution-karaf-0.5.x-Boron/data/
   creating: distribution-karaf-0.5.x-Boron/data/tmp/
   creating: distribution-karaf-0.5.x-Boron/deploy/
   creating: distribution-karaf-0.5.x-Boron/etc/
   creating: distribution-karaf-0.5.x-Boron/externalapps/
...
  inflating: distribution-karaf-0.5.x-Boron/bin/start.bat
  inflating: distribution-karaf-0.5.x-Boron/bin/status.bat
  inflating: distribution-karaf-0.5.x-Boron/bin/stop.bat
$ cd distribution-karaf-0.5.x-Boron
$ ./bin/karaf

    _________                       _________                  .__   .__          .__      __
    \_____   \ ______   ____   ____ \_______ \ ______  ___.__.\|   \| \|__\| ____ \|   \|__\|__
↪_/  \|_
     /   \|   \\____ \_/ __ \ /    \ \|    \|  \\__  \<   \|   \|\|   \| \|   \|/ ___\\|␣
↪ \|  \   __\
```

(continues on next page)

(continued from previous page)

```
    /    \|    \  \|_> >  ___/\|   \|  \\|    `   \/ __ \\___  \|\|  \|_\|  / /_/  > ␣
↪ Y  \  \|
    \________  /   __/ \___  >___\|  /________  (____  / ____\|\|____/__\___  /\|___\| ␣
↪/__\|
             \/\|__\|        \/     \/        \/     \/\/            /_____/      \/
```

- Press `tab` for a list of available commands
- Typing `[cmd] --help` will show help for a specific command.
- Press `ctrl-d` or type `system:shutdown` or `logout` to shutdown OpenDaylight.

Note: Please take a look at the *Deployment Recommendations* and following sections under *Security Considerations* if you're planning on running OpenDaylight outside of an isolated test lab environment.

Install the Karaf features

To install a feature, use the following command, where feature1 is the feature name listed in the table below:

```
feature:install <feature1>
```

You can install multiple features using the following command:

```
feature:install <feature1> <feature2> ... <featureN-name>
```

Note: For compatibility reasons, you cannot enable all Karaf features simultaneously. The table below documents feature installation names and known incompatibilities.Compatibility values indicate the following:

- **all** - the feature can be run with other features.
- **self+all** - the feature can be installed with other features with a value of **all**, but may interact badly with other features that have a value of **self+all**. Not every combination has been tested.

Uninstalling features

To uninstall a feature, you must shut down OpenDaylight, delete the data directory, and start OpenDaylight up again.

Important: Uninstalling a feature using the Karaf feature:uninstall command is not supported and can cause unexpected and undesirable behavior.

Listing available features

To find the complete list of Karaf features, run the following command:

```
feature:list
```

To list the installed Karaf features, run the following command:

```
feature:list -i
```

Features to implement networking functionality provide release notes, which you can find in the *Project Specific Release Notes* section.

Karaf running on Windows 10

Windows 10 cannot be identify by Karaf (equinox). Issue occurs during installation of karaf features e.g.:

```
opendaylight-user@root>feature:install odl-restconf
Error executing command: Can't install feature odl-restconf/0.0.0:
Could not start bundle mvn:org.fusesource.leveldbjni/leveldbjni-all/1.8-odl in␣
↪feature(s) odl-akka-leveldb-0.7: The bundle "org.fusesource.leveldbjni.leveldbjni-
↪all_1.8.0 [300]" could not be resolved. Reason: No match found for native code:␣
↪META-INF/native/windows32/leveldbjni.dll; processor=x86; osname=Win32, META-INF/
↪native/windows64/leveldbjni.dll; processor=x86-64; osname=Win32, META-INF/native/
↪osx/libleveldbjni.jnilib; processor=x86; osname=macosx, META-INF/native/osx/
↪libleveldbjni.jnilib; processor=x86-64; osname=macosx, META-INF/native/linux32/
↪libleveldbjni.so; processor=x86; osname=Linux, META-INF/native/linux64/
↪libleveldbjni.so; processor=x86-64; osname=Linux, META-INF/native/sunos64/amd64/
↪libleveldbjni.so; processor=x86-64; osname=SunOS, META-INF/native/sunos64/sparcv9/
↪libleveldbjni.so; processor=sparcv9; osname=SunOS
```

Workaround is to add

    org.osgi.framework.os.name = Win32

to the karaf file

    etc/system.properties

The workaround and further info are in this thread: http://stackoverflow.com/questions/35679852/karaf-exception-is-thrown-while-installing-org-fusesource-leveldbjni

Karaf OpenDaylight Features

Table 1: Karaf OpenDaylight features

Feature Name	Feature Description	Karaf feature name	Compatibility
Authentication	Enables authentication with support for federation using Apache Shiro	odl-aaa-shiro	all
BGP	Provides support for Border Gateway Protocol (including Link-State Distribution) as a source of L3 topology information	odl-bgpcep-bgp	all
BMP	Provides support for BGP Monitoring Protocol as a monitoring station	odl-bgpcep-bmp	all
DIDM	Device Identification and Driver Management	odl-didm-all	all
Centinel	Provides interfaces for streaming analytics	odl-centinel-all	all
DLUX	Provides an intuitive graphical user interface for OpenDaylight	odl-dlux-all	all

Continued on next page

Table 1 – continued from previous page

Feature Name	Feature Description	Karaf feature name	Compatibility
Fabric as a Service (Faas)	Creates a common abstraction layer on top of a physical network so northbound APIs or services can be more easiliy mapped onto the physical network as a concrete device configuration	odl-faas-all	all
Group Based Policy	Enables Endpoint Registry and Policy Repository REST APIs and associated functionality for Group Based Policy with the default renderer for OpenFlow renderers	odl-groupbasedpolicy-ofoverlay	all
GBP User Interface	Enables a web-based user interface for Group Based Policy	odl-groupbasedpolicyi-ui	all
GBP FaaS renderer	Enables the Fabric as a Service renderer for Group Based Policy	odl-groupbasedpolicy-faas	self+all
GBP Neutron Support	Provides OpenStack Neutron support using Group Based Policy	odl-groupbasedpolicy-neutronmapper	all
L2 Switch	Provides L2 (Ethernet) forwarding across connected OpenFlow switches and support for host tracking	odl-l2switch-switch-ui	self+all
LACP	Enables support for the Link Aggregation Control Protocol	odl-lacp-ui	self+all
LISP Flow Mapping	Enables LISP control plane services including the mapping system services REST API and LISP protocol SB plugin	odl-lispflowmapping-msmr	all
NEMO CLI	Provides intent mappings and implementation with CLI for legacy devices	odl-nemo-cli-renderer	all
NEMO OpenFlow	Provides intent mapping and implementation for OpenFlow devices	odl-nemo-openflow-renderer	self+all
NetIDE	Enables portabilty and cooperation inside a single network by using a client/server multi-controller architecture	odl-netide-rest	all
NETCONF over SSH	Provides support to manage NETCONF-enabled devices over SSH	odl-netconf-connector-ssh	all
OF-CONFIG	Enables remote configuration of OpenFlow datapaths	odl-of-config-rest	all
OVSDB OpenStack Neutron	OpenStack Network Virtualization using OpenDaylight's OVSDB support	odl-ovsdb-openstack	all
OVSDB Southbound	OVSDB MDSAL southbound plugin for Open_vSwitch schema	odl-ovsdb-southbound-impl-ui	all
OVSDB HWVTEP Southbound	OVSDB MDSAL hwvtep southbound plugin for the hw_vtep schema	odl-ovsdb-hwvtepsouthbound-ui	all
OVSDB NetVirt SFC	OVSDB NetVirt support for SFC	odl-ovsdb-sfc-ui	all
OpenFlow Flow Programming	Enables discovery and control of OpenFlow switches and the topoology between them	odl-openflowplugin-flow-services-ui	all
OpenFlow Table Type Patterns	Allows OpenFlow Table Type Patterns to be manually associated with network elements	odl-ttp-all	all

Continued on next page

Table 1 – continued from previous page

Feature Name	Feature Description	Karaf feature name	Compatibility
Packetcable PCMM	Enables flow-based dynamic QoS management of CMTS use in the DOCSIS infrastructure and a policy server	odl-packetcable-policy-server	self+all
PCEP	Enables support for PCEP	odl-bgpcep-pcep	all
RESTCONF API Support	Enables REST API access to the MD-SAL including the data store	odl-restconf	all
SDNinterface	Provides support for interaction and sharing of state between (non-clustered) OpenDaylight instances	odl-sdninterfaceapp-all	all
SFC over L2	Supports implementing Service Function Chaining using Layer 2 forwarding	odl-sfcofl2	self+all
SFC over LISP	Supports implementing Service Function Chaining using LISP	odl-sfclisp	all
SFC over REST	Supports implementing Service Function Chaining using REST CRUD operations on network elements	odl-sfc-sb-rest	all
SFC over VXLAN	Supports implementing Service Function Chaining using VXLAN tunnels	odl-sfc-ovs	self+all
SNMP Plugin	Enables monitoring and control of network elements via SNMP	odl-snmp-plugin	all
SNMP4SDN	Enables OpenFlow-like control of network elements via SNMP	odl-snmp4sdn-all	all
SSSD Federated Authentication	Enables support for federated authentication using SSSD	odl-aaa-sssd-plugin	all
Secure tag eXchange Protocol (SXP)	Enables distribution of shared tags to network devices	odl-sxp-controller	all
Time Series Data Repository (TSDR)	Enables support for storing and querying time series data with the default data collector for OpenFlow statistics the default data store for HSQLDB	odl-tsdr-hsqldb-all	all
TSDR Data Collectors	Enables support for various TSDR data sources (collectors) including OpenFlow statistics, NetFlow statistics, NetFlow statistics, SNMP data, Syslog, and OpenDaylight (controller) metrics	odl-tsdr-openflow-statistics-collector, odl-tsdr-netflow-statistics-collector, odl-tsdr-snmp-data-collector, odl-tsdr-syslog-collector, odl-tsdr-controller-metrics-collector	all
TSDR Data Stores	Enables support for TSDR data stores including HSQLDB, HBase, and Cassandra	odl-tsdr-hsqldb, odl-tsdr-hbase, or odl-tsdr-cassandra	all
Topology Processing Framework	Enables merged and filtered views of network topologies	odl-topoprocessing-framework	all
Unified Secure Channel (USC)	Enables support for secure, remote connections to network devices	odl-usc-channel-ui	all
VTN Manager	Enables Virtual Tenant Network support	odl-vtn-manager-rest	self+all
VTN Manager Neutron	Enables OpenStack Neutron support of VTN Manager	odl-vtn-manager-neutron	self+all

Other OpenDaylight features

Table 2: Other OpenDaylight features

Feature Name	Feature Description	Karaf feature name	Compatibility
OpFlex	Provides OpFlex agent for Open vSwitch to enforce network policy, such as GBP, for locally-attached virtual machines or containers	n/a	all
NeXt	Provides a developer toolkit for designing network-centric topology user interfaces	n/a	all

Experimental OpenDaylight Features

The following functionality is labeled as experimental in this OpenDaylight release and should be used accordingly. In general, it is not supposed to be used in production unless its limitations are well understood by those deploying it.

Table 3: Other features

Feature Name	Feature Description	Karaf feature name	Compatibility
Authorization	Enables configurable role-based authorization	odl-aaa-authz	all
ALTO	Enables support for Application-Layer Traffic Optimization	odl-alto-core	self+all
CAPWAP	Enables control of supported wireless APs	odl-capwap-ac-rest	all
Clustered Authentication	Enables the use of the MD-SAL clustered data store for the authentication database	odl-aaa-authn-mdsal-cluster	all
Controller Shield	Provides controller security information to northbound applications	odl-usecplugin	all
GBP IO Visor Renderer	Provides support for rendering Group Based Policy to IO Visor	odl-groupbasedpolicy-iovisor	all
Internet of Things Data Management	Enables support for the oneM2M specification	odl-iotdm-onem2m	all
LISP Flow Mapping OpenStack Network Virtualization	Experimental support for OpenStack Neutron virtualization	odl-lispflowmapping-neutron	self+all
Network Intent Composition (NIC)	Provides abstraction layer for communcating network intents (including a distributed intent mapping service REST API) using either Hazelcast or the MD-SAL as the backing data store for intents	odl-nic-core-hazelcast or odl-nic-core-mdsal	all
NIC Console	Provides a Karaf CLI extension for intent CRUD operations and mapping service operations	odl-nic-console	all
NIC VTN renderer	Virtual Tenant Network renderer for Network Intent Composition	odl-nic-renderer-vtn	self+all
NIC GBP renderer	Group Based Policy renderer for Network Intent Composition	odl-nic-renderer-gbp	self+all
NIC OpenFlow renderer	OpenFlow renderer for Network Intent Composition	odl-nic-renderer-of	self+all
NIC NEMO renderer	NEtwork MOdeling renderer for Network Intent Composition	odl-nic-renderer-nemo	self+all
OVSDB NetVirt UI	OVSDB DLUX UI	odl-ovsdb-ui	all
Secure Networking Bootstrap	Defines a SNBi domain and associated white lists of devices to be accommodated to the domain	odl-snbi-all	self+all
UNI Manager	Initiates the development of data models and APIs to facilitate configuration and provisioning connectivity services for OpenDaylight applications and services	odl-unimgr	all
YANG PUBSUB	Allows subscriptions to be placed on targeted subtrees of YANG datastores residing on remote devices to obviate the need for OpenDaylight to make continuous fetch requests	odl-yangpush-rest	all

Install support for REST APIs

Most components that offer REST APIs will automatically load the RESTCONF API Support component, but if for whatever reason they seem to be missing, install the "odl-restconf" feature to activate this support.

1.1.10 Release Notes

Target Environment

For Execution

The OpenDaylight Karaf container, OSGi bundles, and Java class files are portable and should run on any Java 7- or Java 8-compliant JVM to run. Certain projects and certain features of some projects may have additional requirements. Those are noted in the project-specific release notes.

Projects and features which have known additional requirements are:

- TCP-MD5 requires 64-bit Linux
- TSDR has extended requirements for external databases
- Persistence has extended requirements for external databases
- SFC requires addition features for certain configurations
- SXP depends on TCP-MD5 on thus requires 64-bit Linux
- SNBI has requirements for Linux and Docker
- OpFlex requires Linux
- DLUX requires a modern web browser to view the UI
- AAA when using federation has additional requirements for external tools
- VTN has components which require Linux

For Development

OpenDaylight is written primarily in Java project and primarily uses Maven as a build tool Consequently the two main requirements to develop projects within OpenDaylight are:

- A Java 8-compliant JDK
- Maven 3.1.1 or later

Applications and tools built on top of OpenDaylight using it's REST APIs should have no special requirements beyond whatever is needed to run the application or tool and make the REST calls.

In some places, OpenDaylight makes use of the Xtend language. While Maven will download the appropriate tools to build this, additional plugins may be required for IDE support.

The projects with additional requirements for execution typically have similar or more extensive additional requirements for development. See the project-specific release notes for details.

Known Issues and Limitations

Other than as noted in project-specific release notes, we know of the following limitations:

- Migration from Helium, Lithium and Beryllium to Boron has not been extensively tested. The per-project release notes include migration and compatibility information when it is known.
- There are scales beyond which the controller has been unreliable when collecting flow statistics from OpenFlow switches. In tests, these issues became apparent when managing thousands of OpenFlow switches, however this may vary depending on deployment and use cases.

Security Advisories

All OpenDaylight Security Advisories can be found on the Security Advisories wiki page. Of particular note to OpenDaylight Boron users are:

- CVE-2017-1000357
- CVE-2017-1000358
- CVE-2017-1000361

There are known and documented mitigations described on the Security Advisory page linked above. Because of the efficacy of the mitigations, we do not intend to release another version of Beryllium to address them. Instead, we encourage all of those who are using Beryllium to carefully understand the mitigations in the context of their deployments.

The following two CVEs were fixed in Boron-SR3, but affect Boron-SR2 and before:

- CVE-2017-1000359
- CVE-2017-1000360

Major Changes

- Bug 2594 Restconf PUT now returns 201 status code instead of 200 when a resource has been created. Before, when creating new resource with PUT method, status code 200 OK is returned. But RESTCONF Protocol draft-bierman-netconf-restconf-04 says: Consistent with [RFC2616], if the PUT method creates a new resource, a "201 Created" Status-Line is returned. If an existing resource is modified, either "200 OK" or "204 No Content" are returned.

Project Specific Release Notes

For the release notes of individual projects, please see the following pages on the OpenDaylight Wiki.

- ALTO
- Atrium Router
- Authentication, Authorization and Accounting (AAA)
- BGP PCEP
- Cardinal
- Centinel
- Control And Provisioning of Wireless Access Points (CAPWAP)
- Controller
- Controller Shield
- DLUX
- Device Identification and Driver Management (DIDM)
- Documentation
- Energy Management Plugin
- Fabric As A Service (FaaS)
- Genius

- Group Based Policy (GBP)
- Honeycomb Virtual Bridge Domain
- Infrastructure Utilities
- Integration/Distribution
- Internet of Things Data Management (IoTDM)
- L2 Switch
- LISP Flow Mapping
- Link Aggregation Control Protocol (LACP)
- MD-SAL
- NAT Application Plugin
- NETCONF
- NEtwork MOdeling (NEMO)
- NeXt UI Toolkit
- NetIDE
- Network Intent Composition (NIC)
- Network Virtualization
- Neutron Northbound
- ODL Root Parent
- ORI C&M Protocol (OCP)
- OVSDB Integration
- OpenFlow Configuration Protocol (OF-CONFIG)
- OpenFlow Plugin
- OpenFlow Protocol Library
- Packet Cable/PCMM
- SDN Interface Application (SDNi)
- SNMP Plugin
- SNMP4SDN
- Secure Network Bootstrapping Infrastructure (SNBI)
- Secure tag eXchange Protocol (SXP)
- Service Function Chaining
- Table Type Patterns (TTP)
- Time Series Data Repository (TSDR)
- Topology Processing Framework
- Unified Secure Channel (USC)
- User Network Interface Manager (UNIMGR)
- Virtual Tenant Network (VTN)

- YANG PUBSUB
- YANG Tools

Boron-SR1 Release Notes

This page details changes and bug fixes between the Boron Release and the Boron Stability Release 1 (Boron-SR1) of OpenDaylight.

Projects with No Noteworthy Changes

The following projects had no noteworthy changes in the Boron-SR1 Release:

- ALTO
- Atrium Router
- Cardinal
- Control And Provisioning of Wireless Access Points (CAPWAP)
- Controller Shield
- Device Identification and Driver Management (DIDM)
- Energy Management Plugin
- Fabric As A Service (FaaS)
- Integration/Distribution
- Internet of Things Data Management (IoTDM)
- Link Aggregation Control Protocol (LACP)
- NAT Application Plugin
- NEtwork MOdeling (NEMO)
- NeXt UI Toolkit
- Network Intent Composition (NIC)
- ORI C&M Protocol (OCP)
- OpenFlow Configuration Protocol (OF-CONFIG)
- Packet Cable/PCMM
- SNMP Plugin
- SNMP4SDN
- Secure Network Bootstrapping Infrastructure (SNBI)
- Table Type Patterns (TTP)
- Time Series Data Repository (TSDR)
- Topology Processing Framework
- Unified Secure Channel (USC)
- YANG PUBSUB

Authentication, Authorization and Accounting (AAA)

- 304660 BUG-6956 - Do not wrap Guava as a bundle in the feature definition
- b4aacb Auto-detect secure HTTP in the idmtool script

BGP PCEP

- 40a2e9 BUG-6737: bgp:show-stats Karaf CLI causes NPE
- 81050d BUG-6781: Inbound and outbound connection attempts from controller are not synchronized - created new peer session listener registry in BGPPeerRegistry for the outbound connection establishment logic to get notified when new peer session is created or destroyed - updated outbound connection establishment logic to attempt a connection only when no existing session is present - updated unit-tests
- 7309aa BUG-7004: NPE when configuring BGP peer using OpenConfig API twice - handle scenario where peer not having AFI-SAFI info is reconfigured using OpenConfig API - updated unit-test
- e789e8 BUG-6622 - ClusterSingletonService registration race condition
- e07ac3 Do not wrap Guava as a bundle in features' definition
- 617ca0 BUG-6889: BGPCEP Boron Autorelease Breaking - if server is not ready when client connects, wait for client reconnection before checking for test pass/fail criteria
- 53e8e4 BUG-6955: Fix BGP TestTool
- 827a46 BUG-6954: Create Application Peer with Route Counter
- 67dcc4 BUG-6809: PMSI attribute's mandatory leaves are always enforced
- 3093fa BUG-6257: Implement PMSI tunnel attribute handler
- 7b0516 BUG6257 Add BGP attribute PMSI tunnel to the EVPN Yang
- bf9d2b BUG-6889: BGPCEP Boron Autorelease Breaking
- 873f97 BUG-6788: peer singleton service closed just after initialization
- 4fbc6b BUG-6811: wrong namespace for binding-codec-tree-factory
- 15baa0 BUG-6835: Missing "simple-routing-policy" knob in OpenConfig BGP Neighbor configuration
- 363448 BUG-6675: add missing cluster-id configuration knob
- efe39b BUG-6616: BGP synchronization can happen after the session was closed
- 9f31c0 BUG-6747: Race condition on peer connection
- a55a84 BUG-6647 Increase code coverage and clean up IV
- 078654 BUG-6647 Increase code coverage and clean up III
- adbc08 BUG-6734: Generate correct L3VPN route key
- 5b10d8 BUG-6799: IllegalAccessException on install bgp
- 9c40c9 BUG-6647 Increase code coverage and clean up II
- c807b0 BUG-6647 Increase code coverage and clean up I
- 98fc76 BUG-6784 - Failed to fully assemble schema context for ..
- a1b3b8 BUG-6662: On connection reset by peer, sometimes re-connection attempt stops after HoldTimer expired error

- 63cd93 BUG-4827 - BGP add-path unit tests
- ef40e4 OpenConfig BGP more defensive
- 1a0e80 BUG-6651: Route Advertisement improvement

Centinel

- a1d4a7 BUG-7040 - Deploy centinel UI to Nexus

Controller

- c8356b Do not wrap Guava as a bundle in features' definition
- 13c9db Configurable update-strategy for clusteredAppConfig
- 5ef954 BUG-5700 - Backwards compatibility of sharding api's with old api's
- e12c3c BUG-6910: Fix anyxml node streaming
- c5b1b3 BUG-6540: EOS - handle edge case with pruning pending owner change commits
- 7c89dc DataBrokerTestModule: use AbstractDataBrokerTest without inheritance
- 87cdec BUG-5280: expose backing client actor reference
- 35dbc0 BUG-5280: make EmptyQueue public
- f6d208 BUG-5280: fix a few warnings
- d82a58 BUG-5280: add ExplicitAsk utility class
- c2964f BUG-5280: Create AbstractProxyHistory class
- 44d363 BUG-5280: move proxy instantiation to AbstractClientHistory
- ee7e56 BUG-5280: separate request sequence and transmit sequence
- 6e2480 BUG-5280: split out cds akka client substrate
- e70b8f Move MessageTrackerTest
- 081550 BUG-5280: add maxMessages field to ConnectClientSuccess
- d833b7 BUG-5280: add FrontendMetadata
- cc9e36 BUG-6540: EOS - Prune pending owner change commits on leader change
- 692916 BUG-6540: Fix journal issues on leader changes
- 82f5b4 BUG-6540: Move LeaderInstallSnapshotState to FollowerLogInformation
- 12c069 BUG-6540: Refactor FollowerToSnapshot to its own class
- 54d6e3 Fix relativePaths for mdsal-it-parent under controller

DLUX

- 771965 BUG-6956 - Do not wrap Guava as a bundle in the feature definition

Documentation

- ce7361 Update requirements for Tox
- 5f1abe BGP user guide reworked
- 2449ff Add warning about RtD not cleaning up between runs
- ce5b0b Replace supported admonitions with rst directives
- d39f1b Note that nested formatting isn't supported
- 1364a2 Fix two typos
- 0c45de Update PacketCable User-Guide
- 8ce9c9 Update Unimgr Documentation for Boron Release
- 2347d5 Remove non-participating project's features from Boron docs
- ca9eb6 Change image to figure
- d036d7 Fix sphinx warnings (and some formatting)
- 8fbdde Update tutorial to use OOR instead of LISP
- 668436 Add documentation for SalFlatBatchService in OFP
- 648a06 Update tutorial docs to replace add mapping RPCs with RESTCONF calls

Genius

- cecdfc BUG-6765: Overriding in_port in table0 with Zero value
- 2f201d Fixes for IT base
- eb07cb Add pom for commons
- d10198 BUG-6278: Switch to use odlparent's karaf-parent
- ec321a IdManager Performance Improvements
- f5be80 Enhancements to improve DJC transaction retry mechanisms
- e49433 Upstreaming ITM cache impl and monitoring bug fix
- bb9f02 ODL BUG-6095, bundle:diag failing for ITM bundle. UT:- RemoveExternalEndpoint is pointing to a vpnservice package which is causing the issue, Started the Karaf and checked the bundle status and diag. coming up jjst fine.
- b16704 Make local variables creation and assignment in a single statement. Some other minor formatting (removing commented code, etc.)
- 8be9b2 Checkstyle and formatting.
- d76bde BUG-6786: L3VPN is not honoring VTEP add or delete in operational cloud
- 1826f3 BUG-6726 : Loss of traffic during ODL Cluster reboot
- 08b545 Arp cache feature changes
- 9e74d4 BUG-6776 - Bad instructions returned by genius RPC
- c977fb Intro. new TestIMdsalApiManager implements IMdsalApiManager
- 46b8e6 Adding the Add/Remove ExternalEndpoint commands.

- 42e57e BUG-6838: Retry Mechanism for Batched Transaction
- 3aac36 BUG-6642 - Improvising Batching code
- 8bdc93 Implement an action type nx_load_in_port
- 14e9d6 Fixing overflow in long-to-IPv4 address conversion
- c547a9 Replace some collection.size() > 0 for !collection.isEmpty() to improve readability. Some other minor changes.
- 338db8 Add SFC relevant service binding constants
- 3c1775 Add JavaDoc to AsyncDataTreeChangeListenerBase init() re. @PostConstruct
- 5c8895 Add support to the ITM to create Transport Zones with different UDP: VxLAN: default port VxLAN-GPE: 4880
- 9c5d78 Improved error message for jobs
- a36863 Add fcapsapplication-impl XML config to features/pom.xml
- f18f59 AsyncDataTreeChangeListenerBase @PreDestroy close() for easier DI
- 2e8028 NPE in InterfaceTopologyStateListener
- 631a2e Reverting Overriding in_port in table0 with Zero value
- eddde4 Implement action types required for ping responder
- 749c4b Performs a residual cleanup of ElanPseudoPort flows
- 20d32c BUG-6765 : Overriding in_port in table0 with Zero value
- b7834a BUG-6748: Added support for match on nxm_reg5
- d9fbcb VM Migration: Flows not programmed in new DPN
- 6bd6b9 Arp cache feature changes
- 3e0a4e BUG-6689 - long delays between vm boot and flow installation
- a45578 Add VxLAN-GPE to the interface types list handled by the IFM
- 0b877b BUG-6493 - Interface-Manager performance optimizations
- de231c BUG-6557 : NPE thrown during Interface-mgr RPCs call
- 01704e BUG-6610 Moving ACL service as highest among all the services.

Group Based Policy (GBP)

- 6e665c BUG-6953: fix renderer-node overwriting
- 375d4c Increasing coverage on faas-renderer
- 334cef introducing vpp-demo
- ca2892 Improved GBP-VBD communication process
- dfe97d BUG-6858: adapt to ise api, fix sgt-generator
- ebba85 BUG-6858: adapt to ise api, fix NPE in listener when missing masterDB
- 27a291 BUG-6858: adapt to ise api, wire harvestAll to template-provider
- 4adb23 BUG-6858: adapt to ise api, change lookup from ise

- 91afd6 BUG-6858: adapt to ise api, simultaneous queries
- e19aad BUG-6650: ep-ip/sgt, bump sxp dep.version to SR1
- ce403c BUG-6650: ep-ip/sgt, remove sxp-mapper
- 36336d BUG-6650: ep-ip/sgt, update/rename models and yangs for sxp-ise-adapter
- ed1db1 BUG-6650: ep-ip/sgt, propose sxp-generator
- 7f9c8a BUG-6650: ep-ip/sgt, implement and wire template provider
- 0a5a00 BUG-6650: ep-ip/sgt, prepare removal of deprecated sxp-mapper
- 82d300 BUG-6650: ep-ip/sgt, propose initial template provider api
- 152d62 BUG-6650: ep-ip/sgt, rename groupbasedpolicy-ise-adapter to sxp-ise-adapter
- 68133e BUG-6650: ep-ip/sgt, move groupbasedpolicy-ise-adapter into sxp-integration
- e97dcf BUG-6650: ep-sgt/ip, propose initial sxp-ep-provider
- 508454 [ios-xe-renderer] Increases coverage for PolicyWriterUtil
- 8d1960 BUG-6743 switch to clustered data listener
- eeeec5 BUG-6743: Service group identifier set for neutron-vpp-mapper and vpp-renderer

Honeycomb Virtual Bridge Domain

- f92940 Reference to DataBroker added into VBD blueprint/instance
- 2b7628 Added current status about bridge domain processing
- c323d9 added support for blueprint and ClusterSingletonService

Infrastructure Utilities

- 58ca84 Remove SingletonWithLifecycle, because @Singleton is not inherited
- 6c5388 Fix broken build
- b14251 @Inject convenience helper (org.opendaylight.infrautils.inject)

L2 Switch

- 29f52d BUG-6655 - arphandler unable to flood arp packet
- 0aea1f Using incremental numbers for initial flow can easily conflict with the flows installed through config data store. To make a simple fix, this patch adds L2switch prefix with the incremental flow-id
- 430922 BUG-6278: Switch to use odlparent's karaf-parent

LISP Flow Mapping

- 9d63b2 BUG-7035: Fix race condition in HashMapDbTest
- e56d74 BUG-5047: Use Netty Epoll in SB when available
- 6bd3b1 JUnit Test - RadixTrie test with random IPs.

- 95a77d BUG-6925: Fix NPE in SimpleMapCache
- 387553 BUG-6782: Fix parent insertions for empty children
- 31ef90 Add RadixTrie parent insertion unit test
- 2a6b14 BUG-6782: Fix RadixTrie parent insertions
- d95b1b BUG-6759: Fix NPE when request for expired mapping
- 0e94ae BUG-6754: Add serializer for IPv6 prefix SimpleAddress
- 8d605d Add OOR conf files in tutorial
- 238406 Fix RESTCONF collection for delete IPv4 Key call.

MD-SAL

- 67197e BUG-7009: fix invalid model
- 999641 Remove augmentableToAugmentations maps
- 6f071d Clean up apparently dead (and not thread safe) code
- efc5ff BUG-5561: retain SchemaContext order for bits
- d07e90 Convert to using BatchedListenerInvoker
- f17c5a Move transaction-invariants into producer
- 7723a3 Add cursor lookup fast-path
- e47199 Fix a raw type warning
- 13ed3b Fix raw types
- d49ac5 Make sure we optimize DOMDataTreeIdentifier
- fb75a6 Do not allow transaction creation with an empty shard map.
- 9d2575 Remove public keyword
- c182e1 Encapsulate ShardedDOMDataTreeProducer layout
- 7452aa Fix warnings in AbstractDOMShardTreeChangePublisher
- 3653b3 Do not instantiate iterator for debugging
- 1b1273 Perform delegate cursor enter/exit first
- 23e32b Move lookup check
- a2aa3d Eliminate ShardedDOMDataTreeWriteTransaction.doSubmit()'s return
- d64f50 Do not use entrySet() where values() or keySet() suffices
- 0b4eee Do not use ExecutorService unnecessarily
- b143da Use ImmutableMap instead of Collections.emptyMap()
- 41c7b4 Speed up InmemoryDOMDataTreeShardWriteTransaction's operations
- 2ea7c1 Switch to using StampedLock
- 11da30 Remove mdsal-binding-util from features because it's only a pom file
- 5f693a Improve ShardedDOMDataTreeProducer locking
- 4c7bb2 Improve ShardedDOMDataTreeProducer locking

- 6ffa81 Improve ShardedDOMDataTreeWriteTransaction performance
- 74425f Optimize InMemoryDOMDataTreeShardProducer
- dca009 Fix InMemory shard transaction chaining.
- 395348 Add batching of non-isolated transaction in ShardedDOMDataTreeProducer
- c37d38 checkStyleViolationSeverity=error implemented for mdsal-dom-broker Resolved the merge conflicts. Implemented code review comments. Implemented another set of code review comments.
- 093b38 Use a bounded blocking queue in InmemoryDOMDataTreeShards.
- 41c34c checkStyleViolationSeverity=error implemented for mdsal-dom-inmemory-datastore Changed the local variable indVal to index. An unwanted folder was added accidentally, removed. Code review comments are implemented.

NETCONF

- c16afa Remove unused imports
- b7c112 Update netconf-topology-singleton.xml file formatting
- 38935a Add serialVersionUID to all java.io.Serializable messages
- a7f406 Add the RemoteDeviceId at the begining of the log
- d4e0ec BUG-6714 - Use singleton service in clustered netconf topology
- 07000c BUG-6256 - OpenDaylight RESTCONF XML selects wrong YANG model for southbound NETCONF
- 1ebd12 Fix tests after merging Change 47121 to Yangtools
- 7999d7 BUG-6272 - support RESTCONF PATCH for mounted NETCONF nodes
- 1ad4d5 Add xml config dependency to features pom
- 08a3d1 BUG-6023 - Adress for config subsystem netconf endpoint is not configurable
- 91be81 BUG-6936 - Fix post request
- 362ab0 Unit test for PostDataTransactionUtil class
- c389ac Unit test for RestconfInvokeOperationsUtil class
- a90a3e BUG-5615 - Netconf connector update overwriting existing topology data
- 6d5c49 BUG-6848 - update url pattern of restconf from 16 to 17
- 054442 BUG-6848 - repackage providers for jersey+create xml and json reader for restconf draft17
- 14efd6 BUG-6848 - upgrade XML media type
- 2e946b BUG-6848 - upgrade namespace of notification container
- 3608c0 BUG-6848 - Renaming to draft17
- d575fc Do a proper disconnect when deleting a connector.
- efe5c7 BUG-6099 - ControllerContext#addKeyValue ignores key type when key is derived type from instance-identifier
- 5db0cc BUG-6797 - Fix deadlock on cached schema-changed notifications
- 11655d BUG-6664 - upgrade draft15 to draft16 - change media types
- b996bc BUG-6664 - upgrade draft15 to draft16 - renaming

- 1f5873 Fix broken ApiDocGeneratorTest
- 0607c0 BUG-6343 - Incorrect handling of configuration failures in SAL netconf connector

NetIDE

- 464082 BUG-6812: Add fix to handing of NetIP protocol version

Network Virtualization

- 815885 Fix for BUG-7059
- 1ca70c BUG-7024: When router is associated to L3VPN , VRF entry creations takes long time
- 2ea687 BUG-6089: Fix the wrong implementation for ICMPV6
- e7917c BUG-7031: Implement ping responder for router interfaces
- 56fe0c BUG-6476 : After configuring NAPT, table 26 and table 46 are not programmed
- 4793ff Changed the AsyncDataChangeListenerBase to AsyncDataTreeChangeListenerBase in the NAT reated files
- f55516 Fix missing init for VpnPseudoPortListener
- 594ad8 BUG-6717 - Output to external network group entry is not installed on NAPT FIB table for new DPN
- 456698 BUG-6831: support for l3 directly connected subnet After the fix only unique mac values will be stored in the vpn interface adjacency. This values will be used for the group programming. No duplicate groups will be created.
- 59afa8 BUG-6778 - VPN interface for external port is deleted when clearing router gw interface
- 3ec9cd BUG-6395: Fixed the Problems in using ODL and neutron-l3-agent in Openstack
- 4297eb BUG-6089:Fix for TCP/UDP and ICMP communication between VM's using learn Action according to SG
- eb448b InterVpnLink cache
- 5366c3 BUG-6934: VpnPseudoPort flows not moved to a new DPN
- 8d24e4 BUG-6863 - Router interfaces incorrectly include network interfaces
- 57a4b6 AclServiceTest with http://immutables.org "depluralize" option
- acc05f Cleanup: remove unnecessary boxing/unboxing
- 8919f8 Cleanup: use Java 8 lambdas
- 29e541 BUG-6482: ERROR Log Observations - CSIT (Boron-Legacy)
- e5fdbf Fixes BUG-6909 ACLs TCP/UDP port ranges for the case of all ports 1-65535) should not use port masking at all
- 3b63e9 fix learn security groups
- 4ee773 Arp cache feature changes
- 69affd BUG-6643 fixed broken l2gw functionality
- 001624 BUG-6816: NAT breakage fix for GRE provider type
- 8607b7 BUG-6831: Retain subnetroute with l3 directly-connected subnet

- 8b8b63 BUG-6843 : NPE in router-add leading to failure of router related cases
- 7c8e2e BUG-6779 -After a Cluster Reboot, 10 VPNintfs seen
- 01f9ab BUG-6824 - floating IP rules deleted upon unrelated neutron port delete
- 44f658 Increase AclServiceTest coverage significantly (from 66% to 84%)
- 104259 BUG-6923 - sfc-translation-layer : OVS data path locator options (nsp,nsi,nshc*) are not required.
- 58a846 BUG-6922 - sfc-translation-layer : Do not explictly set RSP name
- 3e7fcd BUG-6921 - SFC-Translation-Layer : Skip acl classifier write before chain creation
- 48fc20 BUG-6395: Fixed the Problems in using ODL and neutron-l3-agent in Openstack
- 54d0ee BUG-6920 : Fix for ACL portSecurityUpdate to work with DjC + listed fixes
- b75028 De-static-ify aclservice utility classes methods and fields
- 3818f1 aclservice end-to-end test, with a bunch of cool new patterns
- 4c4488 Remove unneeded alivenessmonitor-xml css dependency
- 785cad BUG-6474 : Fixed the issue when using ODL with VXLAN Gateway
- 83f1d4 Add clear ping status
- e02b38 Fix BROKEN aclservice listeners
- d4e1ca Fixes logging exceptions, plus few formatting changes
- 0a2af8 Drop Maven prerequisite
- 759bea ipv6: Use versions from odlparent
- 2199c7 Remove duplicate lockmanager bean
- 82979d Modification cloud-servicechain-state.yang key
- 28597f BUG-6861 : Fix for proper tableId in punt action
- e9160f Clean up logging tests
- f39b5a BUG-6841: Few Remote flows not deleted on DPNs
- 9d0dda BUG-6840: New karaf CLI commands
- 19b1d3 Fixes bgpmanager-api folder structure
- 98d6cb BUG-6589 adding support for hwvtep devices ha
- f8921c BUG-6842 : Incorrect error msg upon associating router to VPN with non-existing VPN-ID
- b71f78 BUG-6823 : Performance improvement in DHCP
- edba2b BUG-6770 - Fixes DjC for NPortCL + snmaps serialized + listed changes
- f683c3 BUG-6825:- “BgpManager not started” error when trying to configure Bgp peer For commands class, bgpmanager not supplied as parameter
- 4eb05b BGP-configuration read is failing as shard leader is not available implemeted retry mechanism in bgp-get-config (100Seconds) for MDSAL read
- 76abf7 fix whitespace
- 93d86e modified stale route cleanup timer to 600Sec, in case nothing configured. enabled route removal on stale-path timer expiry

- a5f5de minor fixes related to BGP - command output: F-bit always set to true and fetch Stale-path time from config, show GR-stalepathTime as default in case not-configured.
- cc5d42 set FBIT for bgp to true (always), as we expect to keep the forwarding state (of CSS) eventhough the controller goes down.
- b715b9 BGP networks update callback is triggered even if the content remain same Fix: On Update callback, verify old and new values and act on it
- 81ae16 BUG-6839: Fixes for import/export RT and router dissociation in L3Vpn
- dbd173 BUG-6673: DCN to DTCN changes
- 58abb3 BUG-6725: fix contains below issues
- e66046 BUG-6446: Concurrency changes related to NeutronPortChangeListener
- 35e7c6 BUG-6668 - Security Groups (all implementations) - port_security extension and default DHCP/ICMP drop rules
- 5a158d BUG-6831: support for l3 directly connected subnet
- e2f944 Flow Entries to match ARP packets in GwMacTable(19)
- aa8246 BUG-6721: first few ping requests to a floating IP are receiving multiple responses
- 94efae BUG-6773: Floating IP response answered from all
- f58d8b Performs a residual cleanup of ElanPseudoPort flows
- fd8fd4 BUG-6758: Remove inter-VPN link state even if error
- 57d40e BUG-6673 : DCN to DTCN Changes for various modules
- adc66f BUG-6691: Fix exceptions in natservice for a dual-stack network
- e13013 BUG-6089: Fix for communication between VM's according to SG.
- e98862 Thrift interface changes to support EVPN operations over Quagga BGP stack
- e2e329 BUG-6716:Fix NPE in NeutronvpnNatManager
- 033052 Mask IPv6Prefix in ACL flows
- 28d2f3 BUG-6589 adding support for hwvtep devices ha
- d2e1ad bgp logging fixes
- 9c217e aclservice-impl Listener without dumb @PreDestroy super.close()
- 863fd2 Fix WARNING when port is updated with allowed_address_pairs
- 749762 Fix 6693 -DHCP Server responds to DHCP requests punted from its table(60) only -DHCP server should not run at all when the controller-dhcp-enabled flag is false
- a3d16b BUG-6708 Neighbor NAPT switches group table buckets remain empty Fix race by triggering NAPT neighbor group table update upon tunnel interface state addition
- e9c655 BUG-6727 ExternalRouterListener ignore multiple routers implementation
- f41a80 BUG-6628 - DMAC for L3 entities flows installation only after reversal
- a303b5 Fix wiring issue in openstack.sfc-translator-impl
- eea48c BUG-6741: eth1 flows on table 0 are missing from d2 ovs
- 307a1e BUG-6707 - FIB table rules are not created when DPNTEPInfo is not available
- d88d71 BUG-6732: ARP Replies Intermittent for Floating IP Addresses

- 047979 BUG-6690 - when mixing dpdk & non-dpdk OVS with the same ODL no way to configure different datapath types
- 3a57d5 BUG-6742 FloatingIPHanlder should use the external interface-name
- ded594 BUG-6756: Fix related to missing ACL flows
- 3e2f52 BUG-6748: ACL mechanism uses reg5 instead of reg6.
- dbedee Fixes default SG remote groups rules. 1))Remote default SG rules are not added with Ip addres asn same is fixed 2)Flow id is fixed for ipv4 and ipv6 rules.
- 5a1ae8 BUG-6752: DHCP service is not bound
- 837ac6 aclservice-impl use infrautils AbstractLifecycle
- 68766e BUG-6452: Error logs when deleting neutron network
- fd9b21 aclservice AclInterfaceStateListener update() TODO replace with comment
- 0c7c0a BUG-6677: Create ext-routers when a router is created with ext-gw
- 9feab4 BUG-6687: Fix NPE when updating ExternalNetwork
- 1876ba BUG-6688 - Patch port is not correctly associated to ELAN
- 3f99ad Code for myMAC in the L3VPN pipeline
- 9ff1e5 BUG-6666: Making sure no 0 datapathID is used when adding interfaces to the model, and when the node updated with the datapathID, create the relevant interfaces
- 80ccfd BUG-6628 - Handling missing router entities DMAC table flows
- b3d85d Support multiple routers per external GW
- 4a0531 Fix bad design of AclClusterUtil to make it pluggable for e2e tests
- ee219d BUG-6609: when 2 vm belonging to the same NETWORK/SUBNET get created in different COMOUTE NODE - ping between those 2 did not work

Neutron Northbound

- 21ae38 BUG-6865 - Transcriber skips the uppercase protocol field
- 7ac89e BUG-6736: Fix incorrect destination ip prefix value

ODL Root Parent

- 597e62 Disable stack trace trimming
- 7d99fd Copy in supporting bouncycastle PKIX/CMS/EAC/PKCS/OCSP/TSP/OPENSSL packages
- 817ade BUG-6790: use non-blocking /dev/urandom
- 6125cd BUG-6712: fix bin/shell's classpath

OVSDB Integration

- 18c226 BUG-6851 fix: handling OtherConfig column clear
- 2c79ca BUG-6692: don't mock Optional

- 9ee9ec BUG-6692: remove InstanceIdentifierCodec instance from SBU
- ff745e Change cipher suites for SSL connection
- 31f29c BUG-6854: remove unnecessary generics from OVSDB schema classes
- 1ffd3c BUG-5995: remove org.json
- bca1f7 BUG-6692: use non-deprecated firstKeyOf() variant
- 194a08 BUG-6692: improve TyperUtils tests (first pass)
- 41c08b Fixed inappropriate WARN message.
- 46366f BUG-6692: clean up MonitorRequestBuilder
- 43b759 Clean up SouthboundConstants
- aa1893 Fix clear bug related to "num" in JSON Node

OpenFlow Plugin

- d2ad11 Fix direct statistics
- ff2b50 Fix flow matching function
- a3e97a Remove RoleManager and RoleContext
- e8c17a BUG-6890: Edit to cfg file reflecting that statistics collection is turned on by default
- 0641dc BUG-6890:
- 7f05ed BUG-6930 Notiifcation-suppliers was broken
- ee3af8 BUG-6890: Updated the cfg file with detailed description of usage
- c7e373 BUG-6890: Enabling statistics collection through a config parameter in openflowplugin.cfg
- 983cb0 BUG-6890: Enabling echo timout configurable through config file
- e1d998 BUG-6890: Enabling barrier configurability through cfg file
- 37cdc5 Optimize LLDP packet check
- a6fece BUG-6745 SimplifiedOperationalListener optimation
- 4e32df BUG-6745 Do not skip first data for reconciliation
- 4eca5f Create SemaphoreKeeper inside decorators
- 82b167 BUG-6745 Improve compression queue locking and handle InterruptedException
- 1dd929 Add finals and move thread name constant to provider
- be6a07 BUG-6745 Set compression semaphore to fair
- 4bc17c BUG-6745 Do not ignore syncup return value
- a2299a BUG-6745 Remove thread renaming and unnecessary logging
- 2b99a1 BUG-6745 Fix replacing in compression queue
- 0d0777 Write SwitchFeatures to operational datastore
- bb1ea2 Remove excessive (trace) logging in FRS
- c44de6 Fix translation to packet.received.MatchBuilder
- 8da0e5 Create DeviceMasterShipManager before forwarders

- c63deb BUG-6633 : NXM_OF_IN_PORT support in openflowplugin
- 0ab612 Update comments and imports after DataChangeListener changes
- 01f583 BUG-6749: Set the openflow connection config at xml file
- 5fffad Fix connection closing before initialization
- 56def6 BUG-6665 Clean code
- ab966a ClusterSingletonService cleaning FRM/FRS
- 74500b SONAR TD - StatisticsContextImpl, StatisticsManagerImpl
- db2f2f Update comments around flat-batch service
- f12541 Convert openflowplugin-it to use DTCL instead of DCL
- 731845 Convert OF samples to use DTCL instead of DCL
- 9a08ed Update old links in code to deprecated DataChangeListener
- 4dec3f BUG-6665 - Fix switches scalability
- 018ab3 BUG-6118: making the OFentityListener aware of the InJeopardy() flag
- aba015 BUG-6542 FRS - prevent concurrent reconciliation node config add

OpenFlow Protocol Library

- 0d1629 BUG-6744 - the parameters of the function of registerMeterBandSerializer need to be more refined
- d068a6 BUG-6674 - the key of the serialization function registered by the vendor is not refinement enough

SDN Interface Application (SDNi)

- ddc016 Use of DataBroker from CSS instead of OFP for boron branch OF SDNINTERFACEAPP

Secure tag eXchange Protocol (SXP)

- a36aee BUG-6849 - PurgeAll message is not propagated to other domains
- 02a8a0 BUG-6448 - Add blueprint and clustering support to sxp-controller
- c5a4ab BUG-6448 - Add blueprint and clustering support to sxp-controller

Service Function Chaining

- 07b43c BUG-7039: Fixing when RSPs are deleted by SF/SFF updates

User Network Interface Manager (UNIMGR)

- f49cbd BUG-6767 - Null pointer exception when adding an EVC with no UNIs

Virtual Tenant Network (VTN)

- 5e248d BUG-6278: Use odlparent's karaf-parent to build local distro only for IT.
- b33c56 BUG-6846: Fixed bug in VTN Coordinator shutdown sequence.
- 69cf85 BUG-6278: Switch to use odlparent's karaf-parent.
- e9b37c BUG-6715:Fixed issue in maven-site generation.
- beb01e BUG-6632: Fixed VTN coordinator build error on Fedora 24.

YANG Tools

- 3a6559 Hide BuildGlobalContext methods
- 3ae284 BUG-4456: rework leaker integration
- 9774e1 BUG-6522: do not re-read models
- 4705f8 BUG-6522: grow namespaces lazily
- fc760c BUG-6522: Optimize toString() methods
- ab7e06 Cleanup SupportedExtensionsMapping
- e3757f BUG-6522: streamline QNameToStatementDefinitionMap API
- 05cf5b BUG-6964: reuse StatementDefinitionNamespace
- fb8b10 BUG-6522: intern cached QNames
- fd79ee BUG-6522: Remove an unneeded branch
- 4e6f54 BUG-7025: make the parser tree immutable
- 89a15d BUG-7025: Use ParseTreeWalker.DEFAULT
- 57fee4 BUG-6522: lower map sizing
- 3f498d BUG-6522: do not accidentally grow lists
- d4b39b BUG-6972: optimize SourceSpecificContext
- 936f49 BUG-6522: Improve parser reactor logging
- 032498 BUG-6522: grow StatementContext collections lazily
- f40575 BUG-6972: refactor copy checking
- 613764 BUG-6972: Remove GroupingUtils.needToCreateNewQName()
- 65e018 BUG-6972: eliminate AugmentUtils
- bedd5f BUG-6669: Mandatory nodes cannot be added to node from another module via augment
- 5f14ad BUG-6972: Consolidate copy operations
- bbde9e BUG-6979 - yang.model.util.EffectiveAugmentationSchema is mutable
- a7427e BUG-6329: Parser fails when target node of uses-augment is an unknown node
- e6ceb0 BUG-6522: create a specialized CopyHistory object
- 75ec31 BUG-6522: share instances of Config(Effective)Statement
- 992f98 BUG-6522: allocate copyHistory lazily

- 954cfd BUG-6522: Adjust collection sizes
- 056f73 Improve ContextBuilder error reporting
- 3182c7 BUG-6965: introduce BuiltinDeclaredStatements
- c157a9 BUG-6964: define a namespace to hold model-defined statements
- 9d06db BUG-6491: Some imports are not exposed in Module.getImports
- 755a50 BUG-6961: SchemaContext.getAllModuleIdentifiers() doesnt work for submodules
- 124100 Reuse cardinality constants
- 29c174 Fix inefficient ConstraintEffectiveStatements
- 97896a Fix GroupingEffectiveStatementImpl's unknown nodes
- 80c95e BUG-4456: add RecursiveExtensionResolver
- 91f43b BUG-6757: revert fix for BUG-4456
- 0de136 Fix failing unit test in QueuedNotificationManagerTest
- ad1f74 Use lambdas instead of anonymous classes
- 29596d Use lambdas instead of StmtContextUtils.build{Declared,Effecive}
- 0e2419 BUG-5561: use canonical Bits order
- e4b8cc Fix a bunch of warnings
- 903b93 Intro. new yangtools.testutils artifacts (incl. Mikito)
- 086f01 Add YangInstanceIdentifier fast paths
- a92b31 Use YangInstanceIdentifier.EMPTY
- c87622 BUG-6771: Problem with typedefs nested in augment
- 0ab997 Allow QueuedNotificationManager to batch notifications
- 21782e Add batching to QueuedNotificationManager
- 8fb1f4 Cleanup QueuedNotificationManager
- 920c87 BUG-6551: Support for third-party Yang extensions implementation
- c55522 Fix ConstraintDefinition inconsistency

Boron-SR2 Release Notes

This page details changes and bug fixes between the Boron Stability Release 1 (Boron-SR1) and the Boron Stability Release 2 (Boron-SR2) of OpenDaylight.

Projects with No Noteworthy Changes

The following projects had no noteworthy changes in the Boron-SR2 Release:

- ALTO
- Atrium Router
- Authentication, Authorization and Accounting (AAA)
- Cardinal

- Centinel
- Control And Provisioning of Wireless Access Points (CAPWAP)
- Controller Shield
- Device Identification and Driver Management (DIDM)
- Energy Management Plugin
- Fabric As A Service (FaaS)
- Infrastructure Utilities
- Integration/Distribution
- Internet of Things Data Management (IoTDM)
- L2 Switch
- Link Aggregation Control Protocol (LACP)
- NAT Application Plugin
- NEtwork MOdeling (NEMO)
- NeXt UI Toolkit
- NetIDE
- ORI C&M Protocol (OCP)
- OpenFlow Configuration Protocol (OF-CONFIG)
- OpenFlow Protocol Library
- Packet Cable/PCMM
- SNMP Plugin
- SNMP4SDN
- Secure Network Bootstrapping Infrastructure (SNBI)
- Table Type Patterns (TTP)
- Time Series Data Repository (TSDR)
- Topology Processing Framework
- Unified Secure Channel (USC)
- Virtual Tenant Network (VTN)
- YANG PUBSUB

BGP PCEP

- 20f0ea BUG-7222: Fix clustering BGPPeer NPE
- 2bbe79 Pcepy: Address pep8 violations
- 537319 BUG-7003: Remove Thread.sleep() calls in tests
- 0ab7e5 BUG-7006: Unit-tests sometimes hangs during execution
- e931d3 BUG-7003: Introduce CheckUtil
- 0cbce4 BUG-7098: Fix Sonar complains

Controller

- d6e8fd BUG-6003: Create config-filtering-parent

DLUX

- 834781 YangUI - quickfix operational list form Yangman - fix mountpoint disconnect Yangman - fix loading Yin schemas for mountpoints
- 335429 Yangman - requests settings
- 0b8364 Logout button added
- be6883 Yangman - make elements accessible via ids - part3
- d49783 Yangman - make elements accessible via ids - part2
- 59e637 Yangman - make elements accessible via ids - part1
- a06cbb Yangman - view switched to json when request is run from history
- 4666bf Yangman - hide show previous item icon if there are no data
- 7fceb3 Remove version tags from modules pom.xml
- 6c1e09 Yangman - delayed progress bar is displayed
- 268a48 Yangman - zero out Status and Time-update
- 70fd6d Yangman - Show all items box showed for a moment-update
- 9b42a6 Yangman - Rpc output list is appending elements instead of replacing
- 0b06e2 Nodes app doesn't display nodes at all

Genius

- 5d42fc Enhancing DataStoreJobCoordinaor logs
- ad7d96 Add utility apis
- 882092 Introduce DataStoreJobCoordinator counters
- 91e5fa Cleanup unwanted exceptions in interfacemanager
- ba5a09 BUG-7220 :port updates are not getting reflected in Table 220
- 80c847 New match Reg4 type and temporary SMAC table definitions
- adde04 BUG-7220 - OVS egress table (220) contains stale rules that send the packet to the wrong port
- ac8710 Fcaps: changing alarm text parameter to be same while raising and clearing
- c46818 Gateway mac table should have unique MAC address for vhu hosts other than 00:00:00:00:00:00
- f6bd5b BUG-6952: DPN can't be added in multiple TZ
- f3fccb BUG-6791 adding async clustered listeners for hwvtep
- 4560ed BUG-7230: tun_id from vxlan tunnel is incorrectly stored into gre key
- 4d10ae Upstreaming BFD monitoring fixes
- b05d26 Addition of constants for ARP Responder

- 5871b9 BUG-7205 l2gw itm mesh is not getting built
- ec7d05 Bug Fix: 7203 Wrong handling of binding service to a tunnel
- 21791d BUG-6589 adding support for hwvtep devices ha
- 90983b BUG-7178: DataStoreCoordinator code and related classes missing
- 4a5654 move interface utilities from ElanUtils, undeprecate Genius IIM
- 990095 Add egress split horizon drop flows for external interfaces
- fa6b37 Cleanup: use plain String concatenation
- fcf814 Cleanup: various performance issues
- d6c81e Cleanup: remove unnecessary type casts
- 9870ee Cleanup: remove redundant type declarations
- 0b7dce Cleanup: remove redundant modifiers
- 947ce3 BUG-6726 : Loss of traffic during ODL Cluster reboot
- dc01a8 BUG-6836 - No access to external network
- 73b09a BUG-6836 - No access to external network
- 792862 Add info to log message with ARP response details on transmit
- c4b0a1 Fix for merge build breakage
- bbb993 BUG-6626 Packet IN handler thread in deadlock after high ARP rate
- db93f4 Added postman collections for id-manager
- b3697e Adding resourcemanager postman collection
- ad917b target-ide/ on .gitignore
- 716657 BUG-7048 - Update to OF port does not change 220 flow
- f011dc Fix for fcaps application module config push error

Group Based Policy (GBP)

- af0d40 BUG-6898 - fixed too slow build in GBPSFC demo
- 47c806 Wire ip-sgt-distribution-service - renderer part
- ddb351 Wire ip-sgt-distribution-service - service part
- e04b67 ip-sgt-distribution-service
- 3bb3a7 Fix collision between VBD UI and GBP UI
- 253e7b BUG-7241: Fix logging for VPP node
- 5ebf40 BUG-7174: stop propagating mandatory/min-elements in configuration nodes
- f652ad Stop manadatory flag propagation in range-value/*

Honeycomb Virtual Bridge Domain

- 010d7c GUI - fixed various REST calls

LISP Flow Mapping

- 0b94ac Fix radix trie lookup exact
- 8fa29b BUG-7272: Fix prefix removal in MultiTableMapCache
- a6d25c BUG-7293: Fix radix trie removals that update the root (2)
- 0429b7 BUG-7293: Fix radix trie node removals that update the root
- ad596c BUG-7018: Fix NPE when SMR and no locators

MD-SAL

- 1b56fa BUG-6163: fixed number of argument when resolving rpc input
- ecbf2b BUG-6135: Java binding v1: IAE from provideTypeForLeafref
- 62a591 BUG-6710 - Close ClusterSingletonServiceRegistration fix
- 8216db BUG-7013: do not rely on default character encoding
- 3abb16 BUG-7064: yang-to-source error on description

NETCONF

- f5d851 BUG-6911 - RPC support in singleton
- 10de38 Add mdsal-singleton-common-api to singleton pom
- 8f4b6f Remove old clustered netconf topology implementation
- 84ff02 BUG-7172 - Correct error-info for missing-attribute errors
- 90c9b6 BUG-7240 - Restconf returns Status.Ok if delete fails
- 655930 BUG-6324 - Notifications stream output is not same as restconf data
- d2591f Set mdsal version to Boron-SR2 version
- ba253b Add logging in tx facade along with the RemoteDeviceId
- 9651a9 Move SubmitFailedReply in the appropriate package
- 43302f Use SerializationUtils to (de)serialize NormalizedNode and YangInstanceIdentifier

Network Intent Composition (NIC)

- b02d45 Fix FlowBuilder issue
- 2206ff Fix issue related to 'Flows aren't pushed to switches'
- ead10c Removed 'min-element' restriction in Intent

Network Virtualization

- 1d12f3 BUG-7368: VPN Engine unable to process external interfaces
- 54d7e3 BUG-7343 - NETVIRT Boron Autorelease Breaking

- 6a2d10 Cleanup ArpNotificationHandler code
- 5e4b78 BUG-7077 - NAPT inbound rules never Expire
- c6f457 BUG-7333: Fix for Arp flows were not deleted for DHCP port in Control node.
- c7638e BUG-7319: thread.sleep in group installation
- 72d9a5 BUG-7305: DHCP fails for Dual stack ports
- 6b8760 BUG-7312: modify param reference of AclNodeListener
- 5dfd60 Netvirt IT: Assert return value of ping in Netvirt IT tests
- 48a428 Fixed BGP AS number field size
- d6a948 BUG-7331: CLI command to create VPNs allows creation of two VPNs with the same RD
- ac846f BUG-6589 Adding retry mechanism to listener
- b6cca3 BUG-7264 Fix missing flows for Remote SG rule Problem: Missing SG remote flows to the first VM associated with the SG. The sample reproduction scenario as below: 1. create empty SG1 2. create VM1 with the SG1 3. add custom TCP rule to SG1 with remote group id as SG1 4. create VM2 with the SG1 There is flow for TCP rule added for VM1 to VM2 but the flow from VM2 to VM1 is missing
- 4ed9a5 BUG-7324: Stale FIB Entries are not getting Removed
- 1768e8 7280 - ARP Responder fix for Floating Ips (extension of BUG-6726)
- 2f454b BUG-7236: handle high rate of src mac learning packet-ins
- b50fe0 BUG-7081 - NAPT is not functional
- 2d0fdb BUG-7081 - NAPT is not functional
- 53d48f BUG-7253: Added learn support for other protocols rule (protocol Number)
- ea6d27 BUG-7128: Added learn support for other protocols rule (ANY)
- 0bd8b8 BUG-7250: Add IPv6 integration tests
- 55cf96 BUG-6998: Fix for VM Instance . Ip Address Not Assigned
- 846d06 BUG-7302: Enable Ping Responder for router interface IPs , on a BGPVPN that has this router associated.
- 5f8434 IT for provider network
- db8803 cleanup unused dependencies for ipv6
- c8031e cleanup unused dependencies for it
- b3a97e cleanup from previous cherry-picks
- ba807d BUG-7236: Add temporary SMAC learning table
- fc38df BUG-7239 : Multiple FIB entries for extra route get created when neutron route-update is done
- 1bdc3e BUG-7298: NPE in vpn manager
- 209217 BUG-7294: Use delete_learned flag on learn flows
- b25b7a BUG-6589 Logging exception
- 44e845 Fcaps: changing alarm text parameter to be same while raising and clearing
- c3ee9c BUG-7239 : Multiple FIB entries for extra route get created when neutron route-update is done
- c54c8a Fix Bug #7289 Set delete_learned flag to Acl Learn flow Entries

- 9d5885 BUG-7263: Spread InterVpnLinks among available DPNs
- c75d04 BUG-6589 adding support for hwvtep devices ha
- fa0d4d BUG-6589 adding support for hwvtep devices ha
- 773a1c BUG-6668 - Security Groups (all implementations) - port_security extension and default DHCP/ICMP drop rules
- 9b78ad BUG-6833: InterVpnLink FIB routes not populated when no VM on VPN
- c80380 BUG-7283: Fix exception in VpnInterfaceManager for IPv6 subnets
- 9d0413 BUG-7233: Multiple VLAN external network communication failed while using compute nodes.
- 9c1ffe BUG-7247-The BGP configuration is getting configured as "router bgp 0"
- 86e1ae BUG-7233: Multiple VLAN external network communication failed while using compute nodes.
- f6b06c BUG-7278 - SC to Elan handover flows priority is wrong
- 8d8505 BUG-7170: ARP thread is sleeping 2s
- a078b6 BUG-7282 - The egress table flows(table 220) are not deleted on port delete.
- debe02 BUG-6589 adding support for hwvtep devices ha
- b0f5aa Add configurable timeouts for acl security groups in Legacy NetVirt. Also sync up default timeout values with those in current new NetVirt (see aclservice-config.yang)
- a95429 BUG-6589 adding support for hwvtep devices ha
- 07dea3 legacy netvirt: forcibly disable port security for network port
- 47558e Listen on Topology Node instead of Inventory's
- f7a9d7 BUG-7234 : Placeholder for BGP minor fixes
- d2f8c9 BUG-6786: L3VPN is not honoring VTEP add or delete in operational cloud
- df867d BUG-6726 : Arp Responder for Internal Subnet Gateway IPAddress
- a212e4 BUG-6589 adding support for hwvtep devices ha
- 947917 BUG-6589 adding support for hwvtep devices ha
- 54a2d9 BUG-7096: After disassociation/association of Router to VPN , it takes 2minutes to update in FIB table
- ca50b0 BUG-6786: L3VPN is not honoring VTEP add or delete in operational cloud
- 1c8b4e Refactor the code that updates the vpn-to-dpn list
- 0ee5c8 BUG-7192 - Inter-VPN link routes BGP leaking not working
- 320bb6 Added default Security Group to test modes other than transparent.
- 193b5d Apply checkstyle fixes on cloud-servicechain
- 05de78 BUG-7208: Import-Export RT feature is not working on stable-boron
- dcf199 BUG-7119: gw arp didn't resolve
- 0d1472 BUG-7188: VpnInterface creation is delayed for 90s
- e29938 legacy netvirt: forcibly disable port security for network port
- c0db2c Checkstyle for dhcpservice-impl
- fa7b58 BUG-7168 - MAC Learning from ARP to be allowed on Ext-Interfaces
- 7f250a Switch the NeutronFloatingIP listener to DTCL

- 229228 Checkstyle for dhcpservice-api
- d93bc9 Fix version warning
- 14894a BUG-6089:Add support for All ICMP code and type in SG using learn
- ec6f3d Add aggregator pom for commons
- 7b0640 Make aggregator poms consistent
- 9a48ee Implement InterVpnLink update operation
- 527429 BUG-7162 - legacy netvirt: null pointer exception
- 1fc33f BUG-7105: Fix learned matches for all TCP/UDP SG
- bf38b8 BUG-7075: AlivenessMonitor skip non-neutron ports
- 8d5cb9 BUG-7093
- a0ec4b BUG-7127 - legacy netvirt: null pointer exception SecurityServicesImpl
- e79020 Support for IPv6 East-West Routing
- 421acd BUG-7126 - legacy netvirt: null pointer exception NeutronSubnetInterface.fromMd
- cb4087 BUG-7016 flows fix the bug that flows is not corrrect after the reconnection of ovs
- a24a12 BUG-7157: Modify inter-VPN link model to enable route leaking
- 1c3ef2 BUG-7125 - legacy: ConcurrentModificationException
- 0bc43c BUG-7147: VpnInterfaces not removed when InterVpnLink is removed
- f15d95 BUG-7124 - legacy netvirt: null pointer exception in SouthboundHandler
- 6270ac BUG-6786: L3VPN is not honoring VTEP add or delete in operational cloud
- f4663f BUG-6822: IVpnLink Static routes not removed on cascade
- 903ea2 BUG-6853 : Directly removing floatingip doesn't rmv mac entries from T19
- 5ecc07 BUG-6777 - FIB entries for RNH routed to VxLAN tunnel for flat/VLAN provider networks
- ebb9ed BUG-7116: Change in TunnelInterfaceStateListener.java
- 5eb3b8 BUG-6934: VpnPseudoPort flows not moved to a new DPN (II)
- 304089 BUG-6904 : ELANInstance read causing NPE on cluster reboot while sending notifs.
- 13cdf6 Fixed UT failure in AclServiceTest
- c3c46c BUG-7020: Deletion issue when VM has multiple SGs with same rules
- ea39a9 BUG-7106: Handles static routes at InterVpnLink creation
- a6ad9e BUG-7086: ELAN broadcast group for VLAN network
- 3f04b6 BUG-7107: RouteOrigin updated to support Local routes
- cede4c BUG-7101: Restrict ARP to only learn non-neutron IPs
- c25143 BUG-7120 : NAT Support For GRE TEP add/del is missing
- 14a21b Fixes BUG-7076 SSH between vm in different network on same compute is blocked even with an allow rule.
- 5a39d5 Fix one of the Netvirt IT test cases failure.
- ea3930 Listens to Network-topology nodes instead of inventory ones

- 13938b BUG-7034 : Replace all write_actions by apply_actions.
- deb9d2 small fixes related to BUG-7031
- 38b5e7 BUG-7091 : When Primary NAPT switch goes down, NAPT switch re-election is not happening.
- c2648d BUG-7074 Fix wrong validation check in dynamic tunnel creation logic
- 11b2f8 BUG-7055: Interface removals donot keep Op DS and Cfg DS consistent.
- ccdff1 BUG-7084: Fix for IP is not assigned in single OpenStack node while creating more than one VM at a time.
- 8f0fab BUG-6940 - Avoid TZ subnet per neutron subnet
- 731423 BUG-7045: ACL: Default flows are not programmed in Cluster environment
- ee8c12 Fix for ACL UT failure
- df93c2 BUG-6992 - legacy: ignore IPv6 router interface
- c23e4f Disable SG IT test until learn test is included
- eb6f82 Do not add br-int when manually deleted
- b2b3eb IT - L3 tests
- 0340d3 Do not log frequest NeutronHostConfig updates

Neutron Northbound

- f40027 NeutornLogger: print data when node is deleted

ODL Root Parent

- 1b6ded BUG-6577: package fixed TrieMap
- a82ccf Blacklist the Triple DES cipher suite

OVSDB Integration

- b8c692 Add docs for OVSDB
- 2aa094 Prep for provider network IT
- d5a921 Checkstyle clean-up: Remove useless “final” in interfaces
- cd1bbd BUG-7201 skip monitoring stats tables
- d54079 BUG-7202 upon node reboot hwvtep op ds is missing
- 5a3245 BUG-6643 hwvtep configuration reconcilation
- 121f89 Corrected data type for “src-mac” in hwvtep.yang

OpenFlow Plugin

- 5a9a2b BUG-6820 - Implement SalExperimenterMpMessageService
- e9deca BUG-7209 - Null Pointer Exception in LearnCodecUtil when add learn flow for ipv6
- a47ccf BUG-6890:Flow-Removed Notification configuration
- ac114e lower log level when stats come before flow is written to deviceflowregistry
- 6990bb Implement SalExperimenterMpMessageService
- 4ac927 Improve cleanup after device disconnected event
- fe3ece BUG-7058 - [Helium Plugin]Stats collection issue when controller disconnect the device
- e65e86 BUG-6890: Statistics-polling configuration
- 9d78c3 Optimize port number lookups
- e790c2 BUG-7011 - Race condition in statistics collection related transaction chain handling

SDN Interface Application (SDNi)

- c92ed0 Do not skip deployment of UI artifacts

Secure tag eXchange Protocol (SXP)

- 9d4992 BUG-7121 - SXP filtering model does not contain presence statement
- e0a7da BUG-6760 - Connections in both mode need to be handled separately
- 2a9beb BUG-6999 - Node Listener closes its own datastore access

Service Function Chaining

- a6d185 BUG-7243 : Modify the checkbox text color to be visible

User Network Interface Manager (UNIMGR)

- ce8ddf BUG-7100 - Fix null pointer exception

YANG Tools

- 2bb7fc BUG-5968: Mandatory leaf enforcement does not work in some cases
- cf932e Eliminate use of environment variables
- 42475e Fix IT test instantiation
- f55a31 BUG-5717: eliminate StmtContext.substatements()
- cbe54f BUG-6150: add an import statement special-case
- f10946 BUG-5968: Mandatory leaf enforcement does not work in some cases
- ff8790 Intern SchemaContext.NAME

- 146a53 BUG-6814: Fix (de)serialization of anyxml in JSON
- a125b1 BUG-7057: introduce UntrustedXML class
- 655c13 BUG-5968: Mandatory leaf enforcement does not work in some cases

Boron-SR3 Release Notes

This page details changes and bug fixes between the Boron Stability Release 2 (Boron-SR2) and the Boron Stability Release 3 (Boron-SR3) of OpenDaylight.

Projects with No Noteworthy Changes

The following projects had no noteworthy changes in the Boron-SR3 Release:

- Atrium Router
- Control And Provisioning of Wireless Access Points (CAPWAP)
- DLUX
- Device Identification and Driver Management (DIDM)
- Group Based Policy (GBP)
- Infrastructure Utilities
- Link Aggregation Control Protocol (LACP)
- NeXt UI Toolkit
- Network Intent Composition (NIC)
- OpenFlow Protocol Library
- Packet Cable/PCMM
- SNMP Plugin
- Secure Network Bootstrapping Infrastructure (SNBI)
- Table Type Patterns (TTP)
- Time Series Data Repository (TSDR)
- Topology Processing Framework
- YANG PUBSUB

ALTO

- 2502cc Remove missed import of odl-mdsal-all
- 75de1f Do not pull in odl-mdsal-all
- 2d452b BUG-5222: do not pull in odl-mdsal-xsql

Authentication, Authorization and Accounting (AAA)

- 82b267 BUG-7774: Cherry pick aaa-cert refactoring from master branch
- c6ba3c move aaa-encrypiotn service to blueprint
- 1414a8 Remove RBAC rule implementation

BGP PCEP

- 5a3311 Correct XSD pattern
- 1b62e3 Fix NPE under TopologyProviderTest test
- 7b48d5 BUG-7964: IllegalStateException under testReconectClient test
- 5fd9a9 BUG-7958: Fix IllegalArgumentException
- 2f0a98 BUG-7003: Remove sleeping from Tests
- 412963 BUG-7937: Fix BGPTerminationReason
- 45fc10 BUG-7505: Conflict Modification
- 0148d9 BUG-7386 Remove blocking gracefulshutdown
- 259f28 BUG-7768 Fix missing registration
- 8b71ff BUG-7614: LocRibWriter does not recover well from transaction chain failure
- d2849f BUG-7222: Make BGP DS clean up asynchronous
- 5e1eb9 BUG-4808: remove reference to IdentityHashMap
- 7a866e BUG-7768 Synchronize ServerSessionManager for PCEP
- ec57bc BUG-7732: Improve BGPDeployer synchronization
- 16e0aa BUG-7708: Invalid Route Distinguisher error while processing withdraw update containing VPN route
- 7bb4e2 BUG-7723 Change BGP connection failure log level
- fdce89 BUG-7351: Update IRO reference
- dbd80d BUG-7673: Improve synchonization under BGP/PCEP Session
- 6f6317 BUG-7673: Fix PCC Mock test failures
- dbc2f5 BUG-7215: Fix App Peer ModifiedNodeDoesNotExistException

Cardinal

- a54a34 BUG-5222: do not pull in odl-mdsal-xsql

Centinel

- 492742 Do not pull in odl-mdsal-all
- 58a4d4 BUG-5222: do not pull in odl-mdsal-xsql

Controller

- 8c9dfe BUG-5222: remove xsql from archetype
- faf24d BUG-7814: Fix InvalidActorNameException
- cbdd5d Fix timing issue in testChangeToVotingWithNoLeader
- adcd0c BUG-6856: Rpc definition should implicitly define input/output
- bc9977 BUG-7746: Fix intermittent EOS test failure and synchronization
- be7c84 Fix intermittent failure in testCloseCandidateRegistrationInQuickSuccession
- c71eca Usage of Collections.unmodifiableCollection is unsafe
- 8dfdfb Add OnDemandShardState to report additional Shard state
- 65f9c2 Add DOMDataTreeCommitCohort example for the cars model
- 5b78f9 Add more info logging in sal-akka-raft
- 944822 CDS: updateMinReplicaCount on RemoveServer
- 25f26d BUG-7608: activate action-service element
- c6b367 BUG-7573: add BucketStore source monitoring
- 1b1643 BUG-3128: cache ActorSelections
- cf005e BUG-3128: rework sal-remoterpc-connector
- 9d1222 BUG-7608: Clarify DOMRpc routing/invocation/listener interactions
- e247eb BUG-7697: add defences against nulls
- 03f387 BUG-6937: Add ReachableMember case to Gossiper
- 600bba BUG-3128: do not open-code routed RPC identification
- 285d96 Remove DOMRpcIdentifier.GLOBAL_CONTEXT
- eb7470 BUG-7594: Expand NormalizedNodeData{Input,Output} to handle SchemaPath
- fc008d BUG-6937: correct format string
- fa66b0 Cleanup RemoteDOMRpcFuture
- d78711 BUG-7608: Add ActionServiceMetadata and ActionProviderBean
- 0361c9 BUG-7506: use common DocumentBuilderFactory
- 90fc6d BUG-7608: OpendaylightNamespaceHandler methods can be static
- f2a7e4 BUG-7608: restructure exception throws
- 01941a BUG-7326: Fix ConcurrentModificationException in Blueprint
- 4f323d Fix FindBugs warnings in blueprint and enable enforcement
- 08a954 Checkstyle compliant src/main|test/resources
- 98b630 Fix CS warnings in blueprint and enable enforcement
- 416a6b BUG-3128: Update RPC router concepts
- c3f368 Update dependendency desc properly in RpcServiceMetadata
- 1f0eea BUG-5222: offload XSQLBluePrint creation to first access
- 707da8 BUG-7469: Advertise CDS DOMDataTreeCommitCohortRegistry

- d3293c BUG-7391: Fix out-of-order LeaderStateChange events

Controller Shield

- 3824fe Removed fixed (and ancient) version of maven-bundle-plugin

Energy Management Plugin

- 9e2e52 BUG-5222: do not pull in odl-mdsal-xsql

Fabric As A Service (FaaS)

- 44a7cc fix demo scripts

Genius

- 9da81f BUG-5222: do not pull in odl-mdsal-xsql
- 864b9f BUG-8048: Potential fix for ID Duplication on 1-node
- 8c0ebc BUG-8048 : Ensure unique ids are allocated
- 894e9e BUG-8049 runOnlyInLeaderNode() - out of order event processing
- b7c672 Updated TestIMdsalApiManager.java to support installFlow() with CheckedFuture return value
- 757219 BUG-7864: Specified Id key does not exist in id pool vpnservices
- 88bbb1 Improving ITM performance in a scale setup
- c3141a Handling RACE conditions in bind/unbind service
- 9a49c9 Harden BFD configuration parameters
- cc7f5e Bind/Unbind Service should work irrespective of Port Status
- a5ee0b Enhancing interface-manager logging
- 33ec80 Adding job retries for DJC bind/unbind service jobs
- eb8e18 BUG-7531 : Different ids allocated for same key
- 44a670 Tunnels in DOWN state in scaled scenario
- 0e0c30 Optimizing southbound Tunnel Events
- 8c514f getInterfaceInfoFromOperationalDS Optimisation
- 4664b9 Fix for id duplication for different id keys
- a0939e To fix grep not working for tep:show & tep:show-state on karaf console
- b92562 Inconsistent Maven Bundle Plugin version in ITM
- ab1866 Optimizing tunnel configuration
- 446aca Enhancing service binding logic to support more services
- 9a08b3 Ignoring a Junit test case in Idmanager to unblock autorelease
- 89af25 BUG-7466 - NPE thrown for interface without lport tag

- 648c66 Fix Idmanager JUnit test case
- 7b80d1 BUG-7494 : Idmanager returns the same Id from the same pool for different threads with different id keys
- b42464 Fixes for duplicate tunnels
- 4c281f BUG-7486: ITM perf and scale fixes
- c73bbe Allow Nicira Extension Actions in BoundServices
- ff7b93 BUG-7450 : suppressing unnecessary warning logs
- ca237e Add new ActionInfo implementations for reg load/move
- 85112e Moving interface-manager CLI utils to use cached entries
- ee165f BUG-7419 : Ids from id pool exhausted
- 77a356 flow entries for multiple subports not getting created
- c58a30 Add isIpInSubnet utility API to NwUtil
- 032d4b BUG-7270 Duplicate remote Mcast mac entry in TOR .

Honeycomb Virtual Bridge Domain

- 06db2c BUG-5222: do not pull in odl-mdsal-xsql

Integration/Distribution

- 009615 BUG-4296: Version to inherit from config-filtering-parent

Internet of Things Data Management (IoTDM)

- 4b2567 BUG-5222: do not pull in odl-mdsal-xsql

L2 Switch

- 47082d BUG-5222: do not pull in odl-mdsal-xsql
- 63bffe BUG-5452 - Adding comments that details the configuration.

LISP Flow Mapping

- ad58d1 BUG-6071: Fix fast path Map-Notify auth data
- 5af5ce Add postman collection in FD.io tutorial
- ec27bf WIP: Update Tutorial for FD.io and OOR

MD-SAL

- 0c8723 BUG-7759 - TEST - Getter of BA object fails to construct class instance
- 7b7b26 BindingGenerator v1 "copy-paste" bug in RPCs
- db2d6f BUG-7759 - TEST - Getter of BA object fails to construct class instance
- ea12e8 BUG-6856: Rpc definition should implicitly define input/output
- e54d13 BUG-6856: Rpc definition should implicitly define input/ouput
- abb67f BUG-6028: check value types for encapsulation
- 0819d4 Fix generate of comma before augmentations in toString generator
- 0f6902 BUG-7222: Improve ClusterSingletonService error handling.
- 9c244e BUG-3147 - Binding spec v1: auto generated code by YANGTOOLS could be more efficient
- d92aa2 Fix getValue() of bits in union
- 3c156c BUG-3147 - Binding spec v1: auto generated code by YANGTOOLS could be more efficient
- 96d661 Don't use deprecated SourceIdentifier.create() method anymore
- 7b1ef1 BUG-7425: Recognize instance-identifier in union template
- edcae2 Fix backport damage
- 9be1c8 New test utility AssertDataObjects
- f8094a BUG-6236: Introduce "mdsal.skip.verbose" property, for build speed

NAT Application Plugin

- 822933 BUG-5222: do not pull in odl-mdsal-xsql

NETCONF

- 1dc51d BUG-7906: Fixed json parsing on patch request
- 6779d4 BUG-7987: Json HTTP PATCH: Problem parsing simple leaf value
- 610d85 Fix autorelease fail
- fb4d06 BUG-6856: Rpc definition should implicitly define input/output
- 0d35db BUG-7728 - modules-state conflict kills rest-connector-default-impl
- 44a6e6 BUG-7728 - modules-state conflict kills rest-connector-default-impl
- ae1a78 Fix error reporting for PUT/POST
- b2a304 BUG-7686 - Make notifications defined by yangs automatic loaded for streaming
- f23b06 BUG-7295 - Incorrect handling of device transactions in clustered setting
- 994506 Clear out tx reference immediately after submit.
- 58e0b7 Remove blocking get in read transaction
- 363f5d BUG-5581: Merge empty map instead of exist check
- c30a4c BUG-5581: Minor Exi decoder optimization

- 864eda Clean up Rfc8040
- e20c6e BUG-7735 - Update restconf models by RFC 8040
- fff87a BUG-7429 - Upgrade docgen in restconf to the latest draft of ietf-restconf
- 2c7043 BUG-7206 - Chinese improper code problem.
- edac43 BUG-5679 - implement new service RestconfService
- 56f1b9 BUG-5679 - add new module ietf-restconf
- 6a6144 BUG-5679 - implement ietf-restconf-monitoring - cleanup
- bbf615 BUG-5679 - implement ietf-restconf-monitoring - streams
- 285fb8 BUG-5679 - implement ietf-restconf-monitoring - capabilities
- d0faba BUG-5679 - prepare constants for ietf-yang-library model
- c00740 BUG-5679 - add new module ietf-restconf-monitoring to restconf
- 2be9e7 BUG-5679 - fix generating of schema by ietf-yang-library and path of RestconfSchemaService
- 38c8d0 BUG-5679 - ietf-yang-library module implemetation
- f061af BUG-5679 - prepare constants for ietf-yang-library model
- 31b18b BUG-5679 - new module 'models' for new models in restconf
- f440bd BUG-5679 - renamed package rest.services to base.services
- 132121 Add missing yang - sal-remote
- f08db0 Do not attempt to construct invalid QNames
- 7a5c4a BUG-7231 - Upgrade ietf-restconf draft17 to draft18
- 41429f BUG-7207 - Bad parsing of identityref in typedef for last ietf-restconf draft
- ff6f68 BUG-6951 - Implement Query parameters - with-defaults
- b262be BUG-6947 / BUG-6948 - implement point and insert query params
- 4ada21 BUG-4883 - implement query parameter - filter
- fd7a91 BUG-6949 / BUG-6950 - Implementation of start-time and stop-time query parameters
- 05767a BUG-6935 - RPC in latest draft doesn't work - problem of parsing to instance identifier
- c7579f BUG-6995 - Change event notification subscription usability PART2
- a3ec2f BUG-6995 - Change event notification subscription usability PART1
- d61bb8 BUG-6903 - Implement Query parameters - fields
- cca10c BUG-6895 - Implement Query parameters - depth
- 693e5b BUG-6382 - add apidoc for latest restconf draft
- 0aa729 BUG-6731 - add valid example values
- ef4102 BUG-6931 - Fix unsupported specific type of leaf
- 5019be Use UTF-8 as the standard output
- 62406a InstanceIdentifierCodecImpl JUnit test
- 75b2cc BUG-6679 - api explorer creates false examples
- 74f910 BUG-6272 - support RESTCONF PATCH for mounted NETCONF nodes

- 5e5f9d BUG-6746 - Restconf: Not working GET operation on mount points
- 9b0c35 Fix close() in provider
- 2c244f Fix autorelease - stub channel and eventLoop

NEtwork MOdeling (NEMO)

- 9bf2a1 BUG-5222: do not pull in odl-mdsal-xsql

NetIDE

- 4256d4 BUG-5222: do not pull in odl-mdsal-xsql

Network Virtualization

- 990076 BUG-5222: do not pull in odl-mdsal-xsql
- d6622e BUG-8046 fix for mac movement issue
- a71e06 BUG-7984: IDLE_TIMEOUT check required in onFlowRemoved.
- 978717 BUG-7387 : Netvirt: qos policy applied on the network, not applied on newly created ports of same network
- f345d0 BUG-7966: Fix route origin for some vrfEntries after VM migration
- 5729c6 BUG-7842: ACL: Arp flows missing in ACL tables for overlapping MAC address
- aa12b7 BUG-7826: proper elan djc job retries
- c804af BUG-7896 OptimisticLockFailedException
- e72a98 BUG-7863 - Add Layer 4 Match for flow entries for TCP/UDP security group rule with no min/max
- 1eb928 BUG-7727 : Local and Connected routes do not get imported
- e54df1 Fix potential NPEs in ELAN tunnel handling
- 603a74 BUG-7418 Run local group creation as async task with key equal to subsequent tasks.
- 795d5d BUG-7725: AAP with prefix 0.0.0.0/0 not supported in ACL
- 0d0dd6 Fix for GwMac flow deletion during interface delete
- 8b06f0 BUG-7875: Separated out snmap create and update workflow
- 416e01 Adding some more debug logs to elan module
- dbaaa3 BUG-7817 & BUG-7838: DHCP ARP flow is not added and irrelevant ARP flows are installed in compute node.
- 47b3a7 BUG-7888: handle update of floating ip port
- d46817 BUG-7878: provider interface MACs are installed on remote DPNs
- 1ed61e Rectified incorrect help usage displayed for BGP add-neighbor cli command
- b28049 BUG-7787 - missing flows in T21
- d26650 BUG-7931: SubnetRoue re-election to be triggered on disconnected nodes

- 987db2 BUG-7876 : After router association to L3vpn, one of the VM ip is not removed from router interface to BGPVPN
- a626bd BUG-7885 - CSIT Sporadic failures - tempest.scenario.test_port_security_macspoofing_port
- 24ef0a BUG-7839: ACL: ACL flows are not deleted from source host during VM migration
- 085965 Use the right service name when binding service
- 8eadc9 corrected the population of BGP Total Prefixes counter
- 6cacda BUG-7856: Reverse SNAT flows order to minimize race possibility
- cec17e BUG-7714: VPN Operational Interfaces not getting removed at all.
- 11509a BUG-7831 : BgpRouter receives unnecessary events
- 19da79 BUG-7881 - Traffic drops when not matching UL SC starting in a VPNPseudoPort
- c82cf8 BUG-7861: No ping response from FIP on 1st router when adding 2nd FIP
- 021f7c BUG-7824 ModifiedNodeDoesNotExistException
- 9352c7 Cleanup errors for networks of unsupported type
- 56d147 BUG-7775: Using DJC for NAT Interface-state Listeners
- 87a7d7 BUG-7824 - ModifiedNodeDoesNotExistException
- bd0183 releasing dcn thread once tunnel interface state dcn delivered
- 25cdfd BUG-7780 : NAT RPC's for getting SNAT/DNAT translation information.
- 520a8c BUG-7815: Using DJC for VpnManager Interface-state Listeners
- c60308 BUG-7843 - Missing buckets in ELAN BC group installation during OVS restart
- 36406c BUG-7786 Delete and re add of access port handling
- 8d36d5 BUG-7772 - Service Chaining is not being applied to VMs in the L3VPN
- 1079dc Adding debug statements to track caching of Operational Vpn Instances
- 5f4c62 Fix priority in IntervpnLink flows installed in LFIB
- 5e3b48 adding lport tag for temporary mac learning
- e9b04d Fix several NPEs showing up in CSIT
- ff40ea Fix BFD regression
- c8cab0 BUG7748: Subnet-op-data empty after cluster reboot
- cd4b0f BUG-7790 - Attempting to install RNH on local DPN for FIB with custom instructions
- cd0b95 Use Objects equals instead of == where necessary
- 479c6a Update netvirt guide - correct DB DROP procedure
- 826aa4 BUG-7599 added l2gw validate cli
- 3dfb74 BUG-6589 l2gw cluster reboot fixes
- 7e6eb4 BUG-7606: Fix for missed tunnel flows, after VM live migration
- 9098a7 BUG-7773: Objects should be compared with "equals()".
- 209df4 Improve log messages.
- afdb29 BUG-7392: L2 Forwarding Table=110 Flows Missing

- a3a91c BUG-7729: Remove redundant tunnel drop flow in table 110
- 079b07 BUG7748: Subnet-op-data empty after cluster reboot
- f7137f BUG-7680: Fix Nexthop when advertising to DCGW
- 59eb41 BUG-7680: leaked routes not advertised to DC-GW
- f89bce BUG-7372 - Supress error of NAPT switch selection failure before router-dpn association
- 9002f3 Setup SMAC on routed packets destined to virtual endpoints
- adb379 BUG-7445: Improve the performance on bulk create.
- 62b919 BUG-7733: NeutronVPN: Error out if VxLAN/VLAN network configured without seg-id
- cf78ad BUG-7714 - Vpn Interface not deleted from oper DS
- e5d465 BUG-7720: create/delete VPN CLI handling addition/removal of subnets
- 51f044 BUG-7601: Cleanup Elan instances when a network is deleted
- d42502 BUG-7717 Fix OOM when defining large number of networks
- e2f333 BUG-7488 Add option to disable auto bridge creation
- 3e8f3e BUG-7591: Allow configuration of inactivity_probe and max_backoff for OVS
- 18628c BUG-7667: SNAT table 46&44 are not getting programmed when private BGPVPN
- b83b1c BUG-7489: Add startup config file for elanmanager-config
- 95cd2c BUG-7700: create-l3vpn (REST/CLI) should not allow another VPN to use the same VPNID
- ac59b7 Fix ElanStatusMonitorJMX failing upon bundle reinitialization.
- b89e52 SNAT tests failures
- 36a01f BUG7308: fix leaf to leaf traffic
- d7d621 BUG-7461
- d25478 BUG-7669: Add multi-provider network support to NetVirt for L2 Gateway.
- bfdc0f Lower debug level when truncating provider port name
- cb9359 [BUG-7543] Replace the request function used by restangular
- 0f1728 BUG-7660 Infinite loop while vpn instance removal
- 0b8e3f BUG-7384: CSIT Exception: NPE in deleteVpnInterface
- 8aa1ee BUG-7532 - arp responder rule sometimes missing after vm reboot
- 908b42 BUG-7601 - Cleanup Elan instances when a network is deleted
- 1599b4 BUG-7436: Handle VpnInterfaces of VpnInstance
- fcef7b BUG-7536: Static routes not handled when ivpnlink becomes active
- b3afca BUG-7530 : ElanPacketInHandler mutexes are too coarse
- e04ffc BUG-7533 : Fix for bind/unbind in DHCP service
- d68ba3 BUG-7567: External subnet group is not updated with external gwmac
- eb9509 BUG-7528 : Don't learn the DMAC flows from other DPNs
- eeb431 BUG-7525 - Inter-VPN link static/connected routes leaking not working
- eb0618 BUG-7547 : Ping from DC-GW to invisible ip configured in VM is failing

- d96917 BUG-7497 - NAPT rules missed for second DPN
- 669678 Fix links to openstack images
- 1a03d6 BUG-7478 : SNAT traffic to use router GW MAC
- 4d6439 Spec to setup SMAC on routed packets destined to virtual endpoints
- 5e20e6 Minor updates to openstack doc
- 6dcb23 BUG-7405: IVpnLink routes not removed from BGP on cascade
- f9473e BUG7318: ETREE ODL learning leaf MACS on wrong tag
- edd071 BUG-7520: Avoid creating auto-tunnels for VLAN tenant networks
- c83656 BUG-7488: Autobridge overwrites DpnId if bridge already exists
- 55d659 BUG-7363: Fix for Flows are overlapped when we add custom SG along with ANY rule.
- beac12 Use Developer Guide text in place of Documentation
- 5badc1 Updates to NetVirt docs from some user feedback
- 2d3c0b Add basic initial docs for new layout
- 8d3814 Update specs template
- d19a3b Initial layout for NetVirt docs migration
- 368a0b Added Creative Commons attribution
- 18f52b Update specs-template
- 384e74 Add link to specs
- 927522 Add template for design spec documents
- c71d28 BUG7339:EtreeLeafBG isn't updated with new remotes
- f8be23 BUG-7476: Configure Reachable Time in IPv6 Router Advt
- ab565a Remove unnecessary page headings
- 968884 Remove unneeded code from VpnUtil
- a2ef06 BUG-7403 : getl3vpn RPC behavioural issues
- 962e39 BUG-7355: Remove Vrf Entries in a single transaction
- 186314 BUG-7406: The flows are overridden.
- 35fdd5 BUG-7393: Flows are not getting removed from table:20 and table:90
- 8aaf43 Add Docs for netvirt
- de2eea BUG-7496: Errors and exceptions handling
- d9ff04 Scalability of ServiceChainTag
- 1d7fb1 BUG-7447: Unexpected flows from T21 to T44 for FIP
- 4d1355 BUG-7358 - Inter-VPN traffic is drop when out_port == in_port
- a9471b BUG-7382: NPE while getting the napt primary-switch-id
- 2f759b BUG-7229: Allow certain ICMPv6 NDP packets by default
- 5ea13d BUG-7423: Clean unnecessary leaked flows and fibEntries
- 0ed778 BUG-7448 - External network recreation fails in newton nodl v2

- a5861d BUG-7340: overwritten rule in T28 for multi-tenant
- 46ae4a BUG-7422 Resolve checkstyle errors
- 43a2ac BUG-7426 Adding elantag along with mac-address as key to synchronized block
- 35904c BUG-7444 : External routes are not getting populated
- 83c0d3 BUG-7463: nexthop in leaked routes is wrongly set
- 727b6f BUG-6866 - missed NAPT rules for second router
- 4cd390 BUG-7142 - all VpnPortIpToPort entries are lost from ODL cache after reboot.
- a1655c BUG-7321: ELAN Pseudo-port flows not installed on new DPNs
- d453e7 BUG-7439: Discard internal VPNs for InterVpnLink purposes
- 18a85c BUG-7409 - Traffic Drop from NFip VM to FIP VM
- 623731 subnet-op-data and port-op-data is empty after cluster reboot
- 8bbec1 BUG-7377, BUG-7383: handling unnecessary error log
- d02fbe BUG-7260: no rules in table 26 for default route
- bdeb09 BUG-7359: duplicate local broadcast group backets

Neutron Northbound

- 6ef259 BUG-7848: Allow neutron port create with security disabled.

ODL Root Parent

- c52629 Bump bouncycastle dependencies from 1.54 to 1.56
- 70436d Bump netty to 4.0.44
- 93acc9 git-commit-id-plugin skipped on mvn -Pq, because it slows down a little
- 6cc350 [eclipse] git-commit-id-plugin ignored in M2E by lifecycle-mapping
- 95d620 git-commit-id-plugin cannot fail build for new projects w.o. .git/
- 44163f git-commit-id-plugin to put a META-INF/git.properties in all built JAR
- 6fffad Skip Jacoco in SingleFeatureTest
- 7b8f97 Bump netty to 4.0.43
- c6391b BUG-6236: Add mdsal.skip.verbose to -Pq Quick profile

ORI C&M Protocol (OCP)

- 3d6cd8 This patch remove ^ and $ signs due to https://git.opendaylight.org/gerrit/#/c/53224/.

OVSDB Integration

- d62858 BUG-5222: do not pull in odl-mdsal-xsql
- 66f070 BUG-7160 - operational store still has node after ovs disconnects
- d06892 BUG-5306: Enable the SSL connection for ovs manager
- ec389f BUG-7836 - Make OVSDB southbound plugin listener port configurable
- 9b5598 BUG-2487 - Too large configuration file from OVS
- c880b5 Checkstyle: fix ParenPad violations
- 05afaa BUG-7023 - NPE at org.opendaylight.ovsdb.southbound.OvsdbConnectionManager.getClient
- 9af401 BUG-7563: Add config for inactivity_probe and max_backoff
- da65cb BUG-7559: Add utilities to expose TP by ext-id, get dpnId from TP
- 3c3451 BUG-6579 removed boilerplate code
- 5eb9b5 BUG-6579 checking the depency before transaction
- d7306b BUG-6579 added dependency queue
- b1426c BUG-6579 handling back to back tx updates
- e96963 BUG-7310: Add configurable option to skip columns
- 941b7e BUG-6579 added basic unit tests
- 2b9dca BUG-7373 - ConflictingModificationAppliedException: Node was created by other transaction
- f83cd3 BUG-7414: Tunnel pushed to wrong OVS

OpenFlow Configuration Protocol (OF-CONFIG)

- eda88e fix UT error

OpenFlow Plugin

- 501d4d Fix statistics race condition on big flows
- 86fd39 BUG-7915 - Zero flows populated in all switches when connected to Leader Node
- 49b07d Add arbitrary mask for nxm-reg
- 308285 Fix connection closing on switch IDLE state
- 9b95b1 BUG-7910 - Flow with ethernet mask (ff:ff:ff:ff:ff:ff), get stored under alien-id in operational data store
- 11f1b6 Fix comparison between port numbers in match
- e0030a BUG-7763 - Openflow plugin deletes switch from topology while changing mastership from one controller to another
- 5ea445 BUG-7736 - Forwarding Rules application cluster singleton id should not use the same cluster singleton id as the openflow switch singleton connection handler
- d813c7 BUG-7764 - Do no throw warning on explicit task cancellation
- ed16f7 BUG-7501 - Ensure delete old statistics and create new ones are executed sequentially to ensure stats are updated properly.

- 4683e5 BUG-7500 - TransactionChainManager: fix synchronization issues and error handling when mdsal throws an error.
- 8cdd80 Fix PacketInV10TranslatorTest
- adbc13 BUG-7608: use blueprint action-provider/action-service
- 32a99f BUG-7499 - ensure statistics scheduler does not die and keep trying while the controller keeps the ownership of the device
- 371c65 BUG-7453 - FlowRemoved doesn't have Removed Reason Information
- 5ea638 BUG-6110: Fixed bugs in statistics manager due to race condition.
- de64a9 Fix Direct statistics RPC - actions part
- 747013 RPC opendaylight-direct-statistics:get-flow-statistics not taking nicira extension match
- dd7f35 BUG-5222: do not pull in odl-mdsal-xsql
- 7d65e9 Bug7485 Make statistics poller parameters configurable.
- ffc6b3 BUG-7071: adding support for fin-timeout
- 8917aa BUG-7481 - Flows with nicira actions get corrupted after cluster failure
- 32b316 BUG-6997 supporting OXM_OF_MPLS_LABEL in nicira extensiona
- e059b1 BUG-7415 Reducing the severity of the log message
- 66c19f BUG-7349 - Flow ID not updated in operational after removing and adding a flow with same match
- f9f165 Add LOG.isDebugEnabled to add performance.
- 3456c4 Improve class with lambdas. Change wrong parameters and variables.
- d1ace0 Split long lines (>120)
- 41b15d Remove unused imports, repair checkstyle first sentence.
- 84a143 BUG-7335 - Flow update rejected by switch generates faulty flow entry in operational DS

SDN Interface Application (SDNi)

- 267746 Do not pull in odl-mdsal-all
- b53576 BUG-5222: do not pull in odl-mdsal-xsql
- 800459 Removed fixed (and ancient) version of maven-bundle-plugin

SNMP4SDN

- dbb4de Fix TopologyServices related internal interface binding failure, and TopologyServiceUtil is removed since no effect

Secure tag eXchange Protocol (SXP)

- 103831 BUG-7347 - NodeIdentity listener does not check security fields
- ecc595 BUG-7517 - Both mode does not send PurgeAll on close

Service Function Chaining

- b6646c Support VxLAN-gpe in sfc103 demo setup
- a83a70 BUG-7548 : Delete SFF vxgpe port on SFF delete

Unified Secure Channel (USC)

- 3f9677 BUG-5222: do not pull in odl-mdsal-xsql
- e1a370 Update to work with BouncyCastle 1.55

User Network Interface Manager (UNIMGR)

- a11ecd BUG-5222: do not pull in odl-mdsal-xsql

Virtual Tenant Network (VTN)

- b1971f Fix BUG-7969, to avoid incompatible database queries contructed
- 679f6d BUG-7402: Upgrade Apache Tomcat for VTN coordinator to 7.0.73.
- 8a18c4 BUG-7360: Fixed VTN coordinator build error on Fedora 25.

YANG Tools

- a26d18 BUG-7759: return deserialized type
- 5ff949 BUG-6392: Fix lenient parsing of unkeyed list entries
- 66a191 BUG-5410 - XSD regular expressions are interpreted as Java regexes (2/2)
- 63d9db BUG-5410 - XSD regular expressions are interpreted as Java regexes (1/2)
- 3cbea0 BUG-6856: Rpc definition should implicitly define input/output
- 0e755b BUG-7879 - refine target node was not found
- 62235f BUG-7182: YangToSourcesProcessor deletes output directory
- 2d4178 BUG-7182: add M2E lifecycle mapping
- f5080e BUG-7568: silence YangTextSchemaContextResolver
- c6bbca SourceIdentifier.create() method deprecation explained
- fc8713 Do not confuse statement and argument names
- 7d5cf4 Fix mandatory statement checking
- 5c0d83 BUG-7440: Fix empty parent in deviate “replace”
- 380048 BUG-7267: catch RuntimeExceptions when processing sources
- 1a712d BUG-7038: cleanup parser listener
- 1fd63e BUG-7161: Do not tolerate source-level exceptions
- cf0eb8 BUG-7267: catch null path offenders

Boron-SR4 Release Notes

This page details changes and bug fixes between the Boron Stability Release 3 (Boron-SR3) and the Boron Stability Release 4 (Boron-SR4) of OpenDaylight.

Projects with No Noteworthy Changes

The following projects had no noteworthy changes in the Boron-SR4 Release:

- Atrium Router
- Authentication, Authorization and Accounting (AAA)
- Cardinal
- Centinel
- Control And Provisioning of Wireless Access Points (CAPWAP)
- Controller Shield
- DLUX
- Device Identification and Driver Management (DIDM)
- Energy Management Plugin
- Fabric As A Service (FaaS)
- Group Based Policy (GBP)
- Honeycomb Virtual Bridge Domain
- Infrastructure Utilities
- Internet of Things Data Management (IoTDM)
- L2 Switch
- LISP Flow Mapping
- Link Aggregation Control Protocol (LACP)
- NAT Application Plugin
- NEtwork MOdeling (NEMO)
- NeXt UI Toolkit
- NetIDE
- Network Intent Composition (NIC)
- Neutron Northbound
- ORI C&M Protocol (OCP)
- OpenFlow Configuration Protocol (OF-CONFIG)
- OpenFlow Protocol Library
- Packet Cable/PCMM
- SDN Interface Application (SDNi)
- SNMP Plugin
- SNMP4SDN

- Secure Network Bootstrapping Infrastructure (SNBI)
- Service Function Chaining
- Table Type Patterns (TTP)
- Time Series Data Repository (TSDR)
- Topology Processing Framework
- Unified Secure Channel (USC)
- User Network Interface Manager (UNIMGR)
- YANG PUBSUB

ALTO

- 4f45d0 Fix yang model validity

BGP PCEP

- e5ec0b BUG-8292 Fix BGP flowspec NLRI length read
- 83214e BUG-8548: Pass missing parameter to
- 2801f2 BUG-8156 Fix PCEP topology registration
- 4b1b55 Fix RD pattern in RouteDistinguisherBuilder
- 042b3e BUG-4491: Race condition under PCEP Topology Provider
- 859728 BUG-8306: Fix Pcep failing test
- c99daf BUG-8202: Fix Conflict Modification on odl-programming
- 732599 BUG-8275: Close ReadOnly transaction
- c3c46c BUG-8198: Fix InferenceException
- 66199e BUG-8252 Fix IP prefix length util
- a41992 BUG-6973: Implement test for InstructionDeployer
- 85bc55 Rework CheckUtil Test
- 24ecbf BUG-8007: Code Clean Up
- 036b01 BUG-8007: Code clean Up
- 560438 BUG-8007: Code Clean Up
- 7576c5 BUG-8007: Code clean up
- a58c6b BUG-6975: Integrate Topology Provider with CS Service
- d53c03 BUG-6973: Wire topology-provider with BP
- feee36 BUG-8213: Add missing mandatory attribute on tests
- aa7617 BUG-8007: Programming Clean up
- ac6687 BUG-6975: Integrate Programming with Cluster Singleton Service
- 636a1b BUG-6973: Implement Programming config via DS

- fea98e BUG-6973: Wire Programming with BP
- 5f0f0e BUG-8203: update Netty isRoot method name
- 6c16df Restrict powermock dependency scope to test
- aa71d4 BUG-7976: Race between peer removal and routes update

Controller

- b11aa4 BUG-8038: Ignore testLeadershipTransferOnShutdown
- ca7d83 BUG-8327: GlobalBundleScanningSchemaServiceImpl should be a proxy
- eff76f BUG-7927: stop scanning bundles on framework stop
- 869a74 Turn off visibility of GlobalBundleScanningSchemaServiceImpl#start()
- 09f442 Remove artifacts entries for long-gone RESTCONF
- 5515ab Move sal-remote to sal-rest-connector
- 8edc82 BUG-8219: optimize empty CompositeDataTreeCohort case
- 1374a9 BUG-7783: increase precision of execution times
- 7b9477 BUG-7814: Add counter to make tx actor names unique

Genius

- 5c7f36 Workaround for BUG-7451
- dbfd73 Add getChildInterfaces to IInterfaceManager
- 5d68e4 BUG-8023 CSIT failures with ELAN drop rule

Integration/Distribution

- 18eade BUG-8197: Remove CAPWAP from compatible

MD-SAL

- ba557d BUG-8449 - BindingToNormalizedNodeCodec fails to deserialize union of leafrefs
- 33f90b BUG-8237 - BI to BA conversion not resolving nested nodes
- 14f049 BUG-8327: Introduce DOMYangTextSourceProvider and implement it
- 262d02 BUG-7927: stop scanning bundles on framework stop
- 98bc46 Lazily create schema context in GlobalBundleScanning*
- c2c61d Turn off visibility of OsgiBundleScanningSchemaService#start()
- 45dfd0 Speed up OsgiBundleScanningSchemaService close
- eba650 BUG-8004: handle implicit RPC input

NETCONF

- 8dfd6a BUG-8085: create missing parent augmentation node
- 5fc76d BUG-8566 direct writes to ordered list fail
- 6fdf99 BUG-8455: Yang Patch response is not having the error details
- e76019 BUG-8289 - 409 in cluster restperfclient test
- e98774 BUG-8405: Add close check to NetconfDevice
- 09a3b0 BUG-8311 - Apidoc: Incomprehensible 500 id model is wrong BUG-8266 - Apidoc explorer is broken after installing Boron SR3
- cfa555 BUG-8197: Deregister schema sources on actor stop
- 0569f0 Remove references to controller's sal-remote
- da2126 Refactor netconf clustered topology tests
- 46bc71 BUG-8152: Add way to configure idle timeout
- 3cf390 Do not include duplicate models
- d57368 BUG-8152: Transaction is already opened
- c27b52 Change handling of netconf cluster transactions
- a4199d Close read-only transactions
- 7f898e BUG-7868: perform checks before starting modifications
- 574d8b BUG-8115: Change URI decoding from ISO-8859-1 to UTF-8
- b936ff BUG-8084 - FilterContentValidator.getKeyValues creates invalid YII key values
- a728c0 BUG-5581: Optimize subtree filtering
- 614388 BUG-8072: Fix decoding of URLs with external identityref
- 751b3f Fix LibraryModulesSchemasTest failure
- d44c83 BUG-7812: NPE when NetconfDeviceSalProvider.close
- c52b7e BUG-8037 YANG Patch using "replace" instead of "merge"

Network Virtualization

- 3ddaa2 BUG-8696 fix elan blueprint xml
- 72e3e2 BUG-7988 - Cluster reboot fix
- 995d2d BUG-7809 - NAT snatGroupIdPool is overlapping with Elan Groups
- 34fae9 Bug8484: Non-NAPT Group action is drop for router associated with BGP-VPN
- 32175d BUG-8376: Fix DHCP for external tunnels
- a9945a BUG-7866 fixing remote bc group for vlan provider network
- 022afb BUG-7599 hwvtep ucast mac add performance improv
- 55da67 BUG-7866 adding retries for remote dmac programming during tunnel up event
- 7c94c9 BUG-7758: Use Trunk instead of Transparent port for Flat networks
- f3b171 BUG-8142 : DHCP timeout issue.

- 348193 BUG-8229: fix bad git merge of handleFloatingIpPortUpdated
- 1e3242 BUG-8023 Handling ELAN remote DMAC programming correctly
- a6d99b BUG-7606: Fix for missing table 110 flow in OVS 2.4 after VM live migration
- 8e23ff BUG-7778: VM's FIP are not able communicate to each other in external network provider
- a3a54f BUG-8165 - Learnt IP route does not reappear on DC-GW after OVSRestart
- 31f21b BUG-7922 - Use counter to keep track of duplicate flow entries
- 99a36e BUG-7939 - Remote flows missing in Table 21
- 4ed33a BUG-7939, 7938, 7968, 7997: Potential fix for the four L3VPN bugs
- 770c08 BUG-7939: VpnService Suite and Tempest failures
- 9521be BUG-8019: when the neutron port acting as gateway is deleted, invisible ip is not removed from FIB
- b7d79b BUG-7816: NullPointerException while create a router in external network provider
- d6b892 Fix ACL IPv6 flows to match on ipv6_src/ipv6_dst for remote SG
- 72359b BUG-7952: ACLService to treat Ethertype=IPv6 and Protocol=icmp as a request for ICMPv6
- 96117b BUG-7979: Fix issue where VM is unable to acquire address during IPv6 tests
- 261ad3 BUG-7913: QosInterfaceStateChangeListener IllegalArgumentException
- c74e6c Migrate l3vpn service docs to netvirt
- 4c2608 BUG-8023: Making ELAN to use StateTunnelList listener
- d0c2e1 Correct several equals() bugs

ODL Root Parent

- 195c5b Bump netty to 4.0.45
- 0bb105 Add cli property to un-skip git-commit-id skip flag

OVSDB Integration

- d55457 BUG-8280: TerminationPoint reconciliation fails
- 5145a7 BUG-6692 - Fix checkstyle problems not detected by the current version
- 25f0d0 BUG-7781 update vlan bindings and tunnel ips
- c4b95f Don't ignore .gitignore
- 20613e BUG-7599 rename opData method as operData
- 018c31 BUG-7599 improve perf of ucast mcas
- 3db9fe BUG-7599 avoid mdsal read
- 602d82 BUG-7599 performance improvement for ucast macs
- d02b70 BUG-7781 update vlan bindings and tunnel ips
- c23fdc BUG-7779: Adding try-catch in hwvtep transactions.
- 2d7190 Fix Checkstyle "Redundant Modifier: Redundant 'final' modifier."

- 672607 Fix NPE in OvsdbConnectionService updateConfigParameter
- 8c8e25 BUG-8055: remove unnecessary type parameters
- 2399d5 BUG-8055: remove all unused imports
- 8b259d BUG-8055: various performance issues
- 752044 BUG-8055: remove unnecessary type arguments
- a2dc83 BUG-8055: remove redundant array creation
- cf13e3 BUG-8055: remove unnecessary type casts
- 45e38d BUG-8055: simplify a couple of streams
- c8a4f9 BUG-8055: use method references instead of lambdas
- 7acf2d BUG-8055: use Java 8 predicates instead of Guava
- 3e88dc BUG-8055: use lambdas instead of anonymous classes

OpenFlow Plugin

- c4fe3d Optimize port status and hello message handling
- b25cf4 BUG-8497 - Provide config knob to disable the Forwarding Rule Manager reconciliation
- 62dc27 Fix logging of exception in HandshakeListenerImpl
- c94d17 Improve property-based configuration
- 438465 Fix masked NXM reg length
- 53428e Fix checkstyle api.openflow.md.util
- a1adc8 Fix checkstyle - api.openflow.md.queue
- 53d724 Fix checkstyle warnings.
- c1e1ce Fix checkstyle warnings
- 88445a Fix modifiers order to comply with Java coding guidelines
- 4d9a32 Fix minor issues regarding checkstyle
- f10e19 BUG-8217: Set error information into direct statistics RPC result.
- c7c10d BUG-7901: fix unsynchronized transaction access
- 55623d Fix DeviceFlowRegistry performance regression
- 36aaf6 Fix table miss flow push

Secure tag eXchange Protocol (SXP)

- b7a538 BUG-8368 - UT - ThreadsWorker tests consist of race conditions

Virtual Tenant Network (VTN)

- 826e06 BUG-8211: Fixed bug that failed to associate MD-SAL flow with VTN flow.
- ab723c BUG-8212: Accept data-flow-source that contains no vlan-id.

- 83e8b1 BUG-8184: Suppress unnecessary send-barrier RPC error logs.
- fadf81 BUG-8191: Suppress warning detected by FindBugs.

YANG Tools

- 848b1f Do not create temporary array for module sorting
- bd168f BUG-8566 Introduce a fallback for ChoiceSchemaNode lookup
- b70220 BUG-7844 - Unable to create LeafRefContext for leafref
- 83d0e1 BUG-8123: fix URL naming mixup
- 5fdb34 BUG-7954: Throw an exception when parsing duplicate (sub)modules
- 9f3901 BUG-8123: be better at guessing identifiers
- c36433 BUG-7062: add revision awareness to Yin/YangTextSchemaSource
- 83fe83 BUG-8039: Enforce binary/string type length

1.1.11 Project-Specific Installation Guides

Centinel Installation Guide

This document is for the user to install the artifacts that are needed for using Centinel functionality in the OpenDaylight by enabling the default Centinel feature. Centinel is a distributed reliable framework for collection, aggregation and analysis of streaming data which is added in this OpenDaylight release.

Overview

The Centinel project aims at providing a distributed, reliable framework for efficiently collecting, aggregating and sinking streaming data across Persistence DB and stream analyzers (e.g., Graylog, Elasticsearch, Spark, Hive). This framework enables SDN applications/services to receive events from multiple streaming sources (e.g., Syslog, Thrift, Avro, AMQP, Log4j, HTTP/REST).

In this release, we develop a "Log Service" and plug-in for log analyzer (e.g., Graylog). The Log service process real time events coming from log analyzer. Additionally, we provide stream collector (Flume- and Sqoop-based) that collects logs from OpenDaylight and sinks them to persistence service (integrated with TSDR). Centinel also includes a RESTCONF interface to inject events to north bound applications for real-time analytic/network configuration. Further, a Centinel User Interface (web interface) will be available to operators to enable rules/alerts/dashboard etc.

Pre Requisites for Installing Centinel

- Recent Linux distribution - 64bit/16GB RAM
- Java Virtual Machine 1.7 or above
- Apache Maven 3.1.1 or above

Preparing for Installation

There are some additional pre-requisites for Centinel, which can be done by integrate Graylog server, Apache Drill, Apache Flume and HBase.

Graylog server2 Installation

- Install MongoDB
  - import the MongoDB public GPG key into apt:

```
sudo apt-key adv --keyserver keyserver.ubuntu.com --recv 7F0CEB10
```

  - Create the MongoDB source list:

```
echo 'deb http://downloads-distro.mongodb.org/repo/debian-sysvinit dist 10gen
↪' | sudo tee /etc/apt/sources.list.d/mongodb.list
```

  - Update your apt package database:

```
sudo apt-get update
```

  - Install the latest stable version of MongoDB with this command:

```
sudo apt-get install mongodb-org
```

- Install Elasticsearch
  - Graylog2 v0.20.2 requires Elasticsearch v.0.90.10. Download and install it with these commands:

```
cd ~; wget https://download.elasticsearch.org/elasticsearch/elasticsearch/
↪elasticsearch-0.90.10.deb
sudo dpkg -i elasticsearch-0.90.10.deb
```

  - We need to change the Elasticsearch cluster.name setting. Open the Elasticsearch configuration file:

```
sudo vi /etc/elasticsearch/elasticsearch.yml
```

  - Find the section that specifies cluster.name. Uncomment it, and replace the default value with graylog2:

```
cluster.name: graylog2
```

  - Find the line that specifies network.bind_host and uncomment it so it looks like this:

```
network.bind_host: localhost
script.disable_dynamic: true
```

  - Save and quit. Next, restart Elasticsearch to put our changes into effect:

```
sudo service elasticsearch restart
```

  - After a few seconds, run the following to test that Elasticsearch is running properly:

```
curl -XGET 'http://localhost:9200/_cluster/health?pretty=true'
```

- Install Graylog2 server
  - Download the Graylog2 archive to /opt with this command:

```
cd /opt; sudo wget https://github.com/Graylog2/graylog2-server/releases/
↪download/0.20.2/graylog2-server-0.20.2.tgz
```

  - Then extract the archive:

```
sudo tar xvf graylog2-server-0.20.2.tgz
```

- Let's create a symbolic link to the newly created directory, to simplify the directory name:

```
sudo ln -s graylog2-server-0.20.2 graylog2-server
```

- Copy the example configuration file to the proper location, in /etc:

```
sudo cp /opt/graylog2-server/graylog2.conf.example /etc/graylog2.conf
```

- Install pwgen, which we will use to generate password secret keys:

```
sudo apt-get install pwgen
```

- Now must configure the admin password and secret key. The password secret key is configured in graylog2.conf, by the password_secret parameter. Generate a random key and insert it into the Graylog2 configuration with the following two commands:

```
SECRET=$(pwgen -s 96 1)
sudo -E sed -i -e 's/password_secret =.*/password_secret = '$SECRET'/' /etc/
↪graylog2.conf

PASSWORD=$(echo -n password | shasum -a 256 | awk '{print $1}')
sudo -E sed -i -e 's/root_password_sha2 =.*/root_password_sha2 = '$PASSWORD'/
↪' /etc/graylog2.conf
```

- Open the Graylog2 configuration to make a few changes: (sudo vi /etc/graylog2.conf):

```
rest_transport_uri = http://127.0.0.1:12900/
elasticsearch_shards = 1
```

- Now let's install the Graylog2 init script. Copy graylog2ctl to /etc/init.d:

```
sudo cp /opt/graylog2-server/bin/graylog2ctl /etc/init.d/graylog2
```

- Update the startup script to put the Graylog2 logs in /var/log and to look for the Graylog2 server JAR file in /opt/graylog2-server by running the two following sed commands:

```
sudo sed -i -e 's/GRAYLOG2_SERVER_JAR=\${GRAYLOG2_SERVER_JAR:=graylog2-server.
↪jar}/GRAYLOG2_SERVER_JAR=\${GRAYLOG2_SERVER_JAR:=\/opt\/graylog2-server\/
↪graylog2-server.jar}/' /etc/init.d/graylog2
sudo sed -i -e 's/LOG_FILE=\${LOG_FILE:=log\/graylog2-server.log}/LOG_FILE=\$
↪{LOG_FILE:=\/var\/log\/graylog2-server.log}/' /etc/init.d/graylog2
```

- Install the startup script:

```
sudo update-rc.d graylog2 defaults
```

- Start the Graylog2 server with the service command:

```
sudo service graylog2 start
```

Install Graylog Server using Virtual Machine

- Download the OVA image from link given below and save it to your disk locally: https://github.com/Graylog2/graylog2-images/tree/master/ova

- Run the OVA in many systems like VMware or VirtualBox.

HBase Installation

- Download hbase-0.98.15-hadoop2.tar.gz
- Unzip the tar file using below command:

```
tar -xvf hbase-0.98.15-hadoop2.tar.gz
```

- Create directory using below command:

```
sudo mkdir /usr/lib/hbase
```

- Move hbase-0.98.15-hadoop2 to hbase using below command:

```
mv hbase-0.98.15-hadoop2/usr/lib/hbase/hbase-0.98.15-hadoop2 hbase
```

- Configuring HBase with java
  - Open your hbase/conf/hbase-env.sh and set the path to the java installed in your system:

    ```
    export JAVA_HOME=/usr/lib/jvm/jdk1.7.0_25
    ```

  - Set the HBASE_HOME path in bashrc file
    - Open bashrc file using this command:

      ```
      gedit ~/.bashrc
      ```

    - In bashrc file append the below 2 statements:

      ```
      export HBASE_HOME=/usr/lib/hbase/hbase-0.98.15-hadoop2

      export PATH=$PATH:$HBASE_HOME/bin
      ```

- To start HBase issue following commands:

```
HBASE_PATH$ bin/start-hbase.sh

HBASE_PATH$ bin/hbase shell
```

- Create centinel table in HBase with stream,alert,dashboard and stringdata as column families using below command:

```
create 'centinel','stream','alert','dashboard','stringdata'
```

- To stop HBase issue following command:

```
HBASE_PATH$ bin/stop-hbase.sh
```

Apache Flume Installation

- Download apache-flume-1.6.0.tar.gz
- Copy the downloaded file to the directory where you want to install Flume.

- Extract the contents of the apache-flume-1.6.0.tar.gz file using below command. Use sudo if necessary:

```
tar -xvzf apache-flume-1.6.0.tar.gz
```

- Starting flume
  - Navigate to the Flume installation directory.
  - Issue the following command to start flume-ng agent:

```
./flume-ng agent --conf conf --conf-file multiplecolumn.conf --name a1 -
↪Dflume.root.logger=INFO,console
```

Apache Drill Installation

- Download apache-drill-1.1.0.tar.gz
- Copy the downloaded file to the directory where you want to install Drill.
- Extract the contents of the apache-drill-1.1.0.tar.gz file using below command:

```
tar -xvzf apache-drill-1.1.0.tar.gz
```

- Starting Drill:
  - Navigate to the Drill installation directory.
  - Issue the following command to launch Drill in embedded mode:

```
bin/drill-embedded
```

- Access the Apache Drill UI on link: http://localhost:8047/
- Go to "Storage" tab and enable "HBase" storage plugin.

Deploying plugins

- Use the following command to download git repository of Centinel:

```
git clone https://git.opendaylight.org/gerrit/p/centinel
```

- Navigate to the installation directory and build the code using maven by running below command:

```
mvn clean install
```

- After building the maven project, a jar file named `centinel-SplittingSerializer-0.0.1-SNAPSHOT.jar` will be created in `centinel/plugins/centinel-SplittingSerializer/target` inside the workspace directory. Copy and rename this jar file to `centinel-SplittingSerializer.jar` (as mentioned in configuration file of flume) and save at location `apache-flume-1.6.0-bin/lib` inside flume directory.
- After successful build, copy the jar files present at below locations to `/opt/graylog/plugin` in graylog server(VM):

```
centinel/plugins/centinel-alertcallback/target/centinel-alertcallback-0.1.0-
↪SNAPSHOT.jar

centinel/plugins/centinel-output/target/centinel-output-0.1.0-SNAPSHOT.jar
```

- Restart the server after adding plugin using below command:

```
sudo graylog-ctl restart graylog-server
```

Configure rsyslog

Make changes to following file:

```
/etc/rsyslog.conf
```

- Uncomment `$InputTCPServerRun 1514`
- Add the following lines:

```
module(load="imfile" PollingInterval="10") #needs to be done just once
input(type="imfile"
File="<karaf.log>" #location of log file
StateFile="statefile1"
Tag="tag1")
*.* @@127.0.0.1:1514 # @@used for TCP
```

  - Use the following format and comment the previous one:

```
$ActionFileDefaultTemplate RSYSLOG_SyslogProtocol23Format
```

- Use the below command to send Centinel logs to a port:

```
tail -f <location of log file>/karaf.log|logger
```

- Restart rsyslog service after making above changes in configuration file:

```
sudo service rsyslog restart
```

Install the following feature

Finally, from the Karaf console install the Centinel feature with this command:

```
feature:install odl-centinel-all
```

Verifying your Installation

If the feature install was successful you should be able to see the following Centinel commands added:

```
centinel:list

centinel:purgeAll
```

Troubleshooting

Check the `../data/log/karaf.log` for any exception related to Centinel related features

Upgrading From a Previous Release

Only fresh installation is supported.

Uninstalling Centinel

To uninstall the Centinel functionality, you need to do the following from Karaf console:

```
feature:uninstall centinel-all
```

Its recommended to restart the Karaf container after uninstallation of the Centinel functionality.

NetVirt Installation Guide

NetVirt Design Specifications

Starting from Carbon, NetVirt uses an RST format Design Specification document for all new features. These specifications are a perfect way to understand various NetVirt features.

Contents:

Table of Contents

- *Title of the feature*
  - *Problem description*
    - *Use Cases*
  - *Proposed change*
    - *Pipeline changes*
    - *Yang changes*
    - *Configuration impact*
    - *Clustering considerations*
    - *Other Infra considerations*
    - *Security considerations*
    - *Scale and Performance Impact*
    - *Targeted Release*
    - *Alternatives*
  - *Usage*
    - *Features to Install*
    - *REST API*
    - *CLI*
  - *Implementation*
    - *Assignee(s)*

        * *Work Items*
    - *Dependencies*
    - *Testing*
        * *Unit Tests*
        * *Integration Tests*
        * *CSIT*
    - *Documentation Impact*
    - *References*

Title of the feature

[link to gerrit patch]

Brief introduction of the feature.

Problem description

Detailed description of the problem being solved by this feature

Use Cases

Use cases addressed by this feature.

Proposed change

Details of the proposed change.

Pipeline changes

Any changes to pipeline must be captured explicitly in this section.

Yang changes

This should detail any changes to yang models.

Configuration impact

Any configuration parameters being added/deprecated for this feature? What will be defaults for these? How will it impact existing deployments?

Note that outright deletion/modification of existing configuration is not allowed due to backward compatibility. They can only be deprecated and deleted in later release(s).

Clustering considerations

This should capture how clustering will be supported. This can include but not limited to use of CDTCL, EOS, Cluster Singleton etc.

Other Infra considerations

This should capture impact from/to different infra components like MDSAL Datastore, karaf, AAA etc.

Security considerations

Document any security related issues impacted by this feature.

Scale and Performance Impact

What are the potential scale and performance impacts of this change? Does it help improve scale and performance or make it worse?

Targeted Release

What release is this feature targeted for?

Alternatives

Alternatives considered and why they were not selected.

Usage

How will end user use this feature? Primary focus here is how this feature will be used in an actual deployment.

e.g. For most netvirt features this will include OpenStack APIs.

This section will be primary input for Test and Documentation teams. Along with above this should also capture REST API and CLI.

Features to Install

odl-netvirt-openstack

Identify existing karaf feature to which this change applies and/or new karaf features being introduced. These can be user facing features which are added to integration/distribution or internal features to be used by other projects.

REST API

Sample JSONS/URIs. These will be an offshoot of yang changes. Capture these for User Guide, CSIT, etc.

CLI

Any CLI if being added.

Implementation

Assignee(s)

Who is implementing this feature? In case of multiple authors, designate a primary assigne and other contributors.

Primary assignee: <developer-a>

Other contributors: <developer-b> <developer-c>

Work Items

Break up work into individual items. This should be a checklist on Trello card for this feature. Give link to trello card or duplicate it.

Dependencies

Any dependencies being added/removed? Dependencies here refers to internal [other ODL projects] as well as external [OVS, karaf, JDK etc.] This should also capture specific versions if any of these dependencies. e.g. OVS version, Linux kernel version, JDK etc.

This should also capture impacts on existing project that depend on Netvirt.

Following projects currently depend on Netvirt: Unimgr

Testing

Capture details of testing that will need to be added.

Unit Tests

Integration Tests

CSIT

Documentation Impact

What is impact on documentation for this change? If documentation change is needed call out one of the <contributors> who will work with Project Documentation Lead to get the changes done.

Don't repeat details already discussed but do reference and call them out.

References

Add any useful references. Some examples:

- Links to Summit presentation, discussion etc.
- Links to mail list discussions
- Links to patches in other projects
- Links to external documentation

[1] OpenDaylight Documentation Guide

[2] https://specs.openstack.org/openstack/nova-specs/specs/kilo/template.html

---

Note: This template was derived from [2], and has been modified to support our project.

This work is licensed under a Creative Commons Attribution 3.0 Unported License. http://creativecommons.org/licenses/by/3.0/legalcode

---

Table of Contents

- *Setup Source-MAC-Address for routed packets destined to virtual endpoints*
  - *Problem description*
    - *Use Cases*
  - *Proposed change*
    - *Pipeline changes*
    - *Yang changes*
    - *Configuration impact*
    - *Clustering considerations*
    - *Other Infra considerations*
    - *Security considerations*
    - *Scale and Performance Impact*
    - *Targeted Release*
    - *Alternatives*
  - *Usage*
    - *Features to Install*
    - *REST API*
    - *CLI*
  - *Implementation*
    - *Assignee(s)*
    - *Work Items*
  - *Dependencies*

- *Testing*
  - *Unit Tests*
  - *Integration Tests*
  - *CSIT*
- *Documentation Impact*
- *References*

Setup Source-MAC-Address for routed packets destined to virtual endpoints

https://git.opendaylight.org/gerrit/#/q/topic:SMAC_virt_endpoints

All L3 Routed packets destined to virtual endpoints in the datacenter managed by ODL do not carry a proper source-mac address in such frames put out to virtual endpoints.

This spec makes sure a proper source-mac is updated in the packet at the point where the packet is delivered to the VM, regardless of the tenant network type. On the actual datapath, there will be no change in the source mac-addresses and packets continue to use the same mechanism that is used today.

Addressing the datapath requires unique MAC allocation per OVS Datapath, so that it can be used as the source MAC for all distributively routed packets of an ODL enabled cloud. It would be handled in some future spec.

Problem description

Today all L3 Routed packets destined to virtual endpoints in the datacenter either

- Incorrectly carry the source mac-address of the originator (regardless of which network the originator is in)
- Incorrectly carry sometimes the reserved source mac address of 00:00:00:00:00:00

This spec is intended to setup a source-mac-address in the frame of L3 Routed packets just before such frames are directed into the virtual endpoints themselves. This enables use-cases where certain virtual endpoints which are VNFs in the datacenter that are source-mac conscious (or mandate that src-mac in frames be valid) can become functional on their instantiation in an OpenDaylight enabled cloud.

Use Cases

- Intra-Datacenter L3 forwarded packets within a hypervisor.
- Intra-Datacenter L3 forwarded packets over Internal VXLAN Tunnels between two hypervisors in the datacenter.
- Inter-Datacenter L3 forwarded packets :
  - Destined to VMs associated floating IP over External VLAN Provider Networks.
  - Destined to VMs associated floating IP over External MPLSOverGRE Tunnels.
  - SNAT traffic from VMs over External MPLSOverGRE Tunnels.
  - SNAT traffic from VMS over External VLAN Provider Networks.

Proposed change

All the L3 Forwarded traffic today reaches the VM via a LocalNextHopGroup managed by the VPN Engine (including FIBManager).

Currently the LocalNextHopGroup sets-up the destination MAC Address of the VM and forwards the traffic to EGRESS_LPORT_DISPATCHER_TABLE (Table 220). In that LocalNextHopGroup we will additionally setup source-mac-address for the frame. There are two cases to decide what source-mac-address should go into the frame:

- If the VM is on a subnet (on a network) for which a subnet gatewayip port exists, then the source-mac address of that subnet gateway port will be setup as the frame's source-mac inside the LocalNextHop group.This is typical of the case when a subnet is added to a router, as the router interface port created by neutron will be representing the subnet's gateway-ip address.
- If the VM is on a subnet (on a network), for which there is no subnet gatewayip port but that network is part of a BGPVPN , then the source-mac address would be that of the connected mac-address of the VM itself. The connected mac-address is nothing but the mac-address on the ovs-datapath for the VMs tapxxx/vhuxxx port on that hypervisor itself.

The implementation also applies to Extra-Routes (on a router) and Discovered Routes as they both use the LocalNextHopGroup in their last mile to send packets into their Nexthop VM.

We need to note that when a network is already part of a BGPVPN, adding a subnet on such a network to a router is disallowed currently by NeutronVPN. And so the need to swap the mac-addresses inside the LocalNextHopGroup to reflect the subnet gatewayip port here does not arise.

For all the use-cases listed in the USE-CASES section above, proper source mac address will be filled-up in the frame before it enters the virtual endpoint.

Pipeline changes

There are no pipeline changes.

The only change is in the NextHopGroup created by VPN Engine (i.e., VRFEntryListener). In the NextHopGroup we will additionally fill up the ethernet source mac address field with proper mac-address as outlined in the 'Proposed change' section.

Currently the LocalNextHopGroup is used in the following tables of VPN Pipeline:

- L3_LFIB_TABLE (Table 20) - Lands all routed packets from MPLSOverGRE tunnel into the virtual endpoint.
- INTERNAL_TUNNEL_TABLE (Table 36) - Lands all routed packets on Internal VXLAN Tunnel within the DC into the virtual end point.
- L3_FIB_TABLE (Table 21) - Lands all routed packets within a specific hypervisor into the virtual endpoint.

```
cookie=0x8000002, duration=50.676s, table=20, n_packets=0, n_bytes=0, priority=10,
↪mpls,mpls_label=70006 actions=write_actions(pop_mpls:0x0800,group:150000)
cookie=0x8000003, duration=50.676s, table=21, n_packets=0, n_bytes=0, priority=42,ip,
↪metadata=0x222f2/0xfffffffe,nw_dst=10.1.1.3 actions=write_actions(group:150000)
cookie=0x9011176, duration=50.676s, table=36, n_packets=0, n_bytes=0, priority=5,tun_
↪id=0x11176 actions=write_actions(group:150000)

NEXTHOP GROUP:
group_id=150000,type=all,bucket=actions=set_field:fa:16:3e:01:1a:40->eth_src,set_
↪field:fa:16:3e:8b:c5:51->eth_dst,load:0x300->NXM_NX_REG6[],resubmit(,220)
```

Yang changes

None.

Configuration impact

None.

Clustering considerations

None.

Other Infra considerations

None.

Security considerations

None.

Scale and Performance Impact

None

Targeted Release

Carbon/Boron

Alternatives

None.

Usage

N/A.

Features to Install

odl-netvirt-openstack

REST API

N/A.

CLI

N/A.

Implementation

Assignee(s)

Primary assignee:

- Achuth Maniyedath (achuth.m@altencalsoftlabs.com)

Other contributors:

- Karthik Prasad (karthik.p@altencalsoftlabs.com)
- Vivekanandan Narasimhan (n.vivekanandan@ericsson.com)

Work Items

https://trello.com/c/IfAmnFFr/110-add-source-macs-in-frames-for-l3-routed-packets-before-such-frames-get-to-the-virtual-endpoint

- Determine the smac address to be used for L3 packets forwarded to VMs.
- Update the LocalNextHopGroup table with proper ethernet source-mac parameter.

Dependencies

No new dependencies.

Testing

Verify the Source-MAC-Address setting on frames forwarded to Virtual endpoints in following cases.

Intra-Datacenter traffic to VMs (Intra/Inter subnet).

- VM to VM traffic within a hypervisor.
- VM to VM traffic across hypervisor over Internal VXLAN tunnel.

Inter-Datacenter traffic to/from VMs.

- External access to VMs using Floating IPs on MPLSOverGRE tunnels.
- External access to VMs using Floating IPs over VLAN provider networks.
- External access from VMs using SNAT over VLAN provider networks.
- External access from VMs using SNAT on MPLSOverGRE tunnels.

Unit Tests

N/A.

Integration Tests

N/A.

CSIT

- Validate that router-interface src-mac is available on received frames within the VM when that VM is on a router-arm.
- Validate that connected-mac as src-mac available on received frames within the VM when that VM is on a network-driven L3 BGPVPN.

Documentation Impact

N/A

References

N/A

OpFlex agent-ovs Install Guide

Required Packages

You'll need to install the following packages and their dependencies:

- libuv
- openvswitch-gbp
- openvswitch-gbp-lib
- openvswitch-gbp-kmod
- libopflex
- libmodelgbp
- agent-ovs

Packages are available for Red Hat Enterprise Linux 7 and Ubuntu 14.04 LTS. Some of the examples below are specific to RHEL7 but you can run the equivalent commands for upstart instead of systemd.

Note that many of these steps may be performed automatically if you're deploying this along with a larger orchestration system.

Host Networking Configuration

You'll need to set up your VM host uplink interface. You should ensure that the MTU of the underlying network is sufficient to handle tunneled traffic. We will use an example of setting up *eth0* as your uplink interface with a vlan of 4093 used for the networking control infrastructure and tunnel data plane.

We just need to set the MTU and disable IPv4 and IPv6 autoconfiguration. The MTU needs to be large enough to allow both the VXLAN header and VLAN tags to pass through without fragmenting for best performance. We'll use 1600

bytes which should be sufficient assuming you are using a default 1500 byte MTU on your virtual machine traffic. If you already have any NetworkManager connections configured for your uplink interface find the connection name and proceed to the next step. Otherwise, create a connection with (be sure to update the variable UPLINK_IFACE as needed):

```
UPLINK_IFACE=eth0
nmcli c add type ethernet ifname $UPLINK_IFACE
```

Now, configure your interface as follows:

```
CONNECTION_NAME="ethernet-$UPLINK_IFACE"
nmcli connection mod "$CONNECTION_NAME" connection.autoconnect yes \
    ipv4.method link-local \
    ipv6.method ignore \
    802-3-ethernet.mtu 9000 \
    ipv4.routes '224.0.0.0/4 0.0.0.0 2000'
```

Then bring up the interface with:

```
nmcli connection up "$CONNECTION_NAME"
```

Next, create the infrastructure interface using the infrastructure VLAN (4093 by default). We'll need to create a vlan subinterface of your uplink interface, the configure DHCP on that interface. Run the following commands. Be sure to replace the variable values if needed. If you're not using NIC teaming, replace the variable team0 below:

```
UPLINK_IFACE=team0
INFRA_VLAN=4093
nmcli connection add type vlan ifname $UPLINK_IFACE.$INFRA_VLAN dev $UPLINK_IFACE id
↪$INFRA_VLAN
nmcli connection mod vlan-$UPLINK_IFACE.$INFRA_VLAN \
    ethernet.mtu 1600 ipv4.routes '224.0.0.0/4 0.0.0.0 1000'
sed "s/CLIENT_ID/01:$(ip link show $UPLINK_IFACE | awk '/ether/ {print $2}')/" \
    > /etc/dhcp/dhclient-$UPLINK_IFACE.$INFRA_VLAN.conf <<EOF
send dhcp-client-identifier CLIENT_ID;
request subnet-mask, domain-name, domain-name-servers, host-name;
EOF
```

Now bring up the new interface with:

```
nmcli connection up vlan-$UPLINK_IFACE.$INFRA_VLAN
```

If you were successful, you should be able to see an IP address when you run:

```
ip addr show dev $UPLINK_IFACE.$INFRA_VLAN
```

OVS Bridge Configuration

We'll need to configure an OVS bridge which will handle the traffic for any virtual machines or containers that are hosted on the VM host. First, enable the openvswitch service and start it:

```
# systemctl enable openvswitch
ln -s '/usr/lib/systemd/system/openvswitch.service' '/etc/systemd/system/multi-user.
↪target.wants/openvswitch.service'
# systemctl start openvswitch
# systemctl status openvswitch
```

(continues on next page)

(continued from previous page)

```
openvswitch.service - Open vSwitch
   Loaded: loaded (/usr/lib/systemd/system/openvswitch.service; enabled)
   Active: active (exited) since Fri 2014-12-12 17:20:13 PST; 3s ago
  Process: 3053 ExecStart=/bin/true (code=exited, status=0/SUCCESS)
 Main PID: 3053 (code=exited, status=0/SUCCESS)
Dec 12 17:20:13 ovs-server.cisco.com systemd[1]: Started Open vSwitch.
```

Next, we can create an OVS bridge (you may wish to use a different bridge name):

```
# ovs-vsctl add-br br0
# ovs-vsctl show
34aa83d7-b918-4e49-bcec-1b521acd1962
    Bridge "br0"
        Port "br0"
            Interface "br0"
                type: internal
    ovs_version: "2.3.90"
```

Next, we configure a tunnel interface on our new bridge as follows:

```
# ovs-vsctl add-port br0 br0_vxlan0 -- \
    set Interface br0_vxlan0 type=vxlan \
    options:remote_ip=flow options:key=flow options:dst_port=8472
# ovs-vsctl show
34aa83d7-b918-4e49-bcec-1b521acd1962
    Bridge "br0"
        Port "br0_vxlan0"
            Interface "br0_vxlan0"
                type: vxlan
                options: {dst_port="8472", key=flow, remote_ip=flow}
        Port "br0"
            Interface "br0"
                type: internal
    ovs_version: "2.3.90"
```

Open vSwitch is now configured and ready.

Agent Configuration

Before enabling the agent, we'll need to edit its configuration file, which is located at "/etc/opflex-agent-ovs/opflex-agent-ovs.conf".

First, we'll configure the Opflex protocol parameters. If you're using an ACI fabric, you'll need the OpFlex domain from the ACI configuration, which is the name of the VMM domain you mapped to the interface for this hypervisor. Set the "domain" field to this value. Next, set the "name" field to a hostname or other unique identifier for the VM host. Finally, set the "peers" list to contain the fixed static anycast peer address of 10.0.0.30 and port 8009. Here is an example of a completed section (bold text shows areas you'll need to modify):

```
"opflex": {
    // The globally unique policy domain for this agent.
    "domain": "[CHANGE ME]",

    // The unique name in the policy domain for this agent.
    "name": "[CHANGE ME]",
```

(continues on next page)

(continued from previous page)

```
    // a list of peers to connect to, by hostname and port.  One
    // peer, or an anycast pseudo-peer, is sufficient to bootstrap
    // the connection without needing an exhaustive list of all
    // peers.
    "peers": [
        {"hostname": "10.0.0.30", "port": 8009}
    ],

    "ssl": {
        // SSL mode.  Possible values:
        // disabled: communicate without encryption
        // encrypted: encrypt but do not verify peers
        // secure: encrypt and verify peer certificates
        "mode": "encrypted",

        // The path to a directory containing trusted certificate
        // authority public certificates, or a file containing a
        // specific CA certificate.
        "ca-store": "/etc/ssl/certs/"
    }
},
```

Next, configure the appropriate policy renderer for the ACI fabric. You'll want to use a stitched-mode renderer. You'll need to configure the bridge name and the uplink interface name. The remote anycast IP address will need to be obtained from the ACI configuration console, but unless the configuration is unusual, it will be 10.0.0.32:

```
// Renderers enforce policy obtained via OpFlex.
"renderers": {
    // Stitched-mode renderer for interoperating with a
    // hardware fabric such as ACI
    "stitched-mode": {
        "ovs-bridge-name": "br0",

        // Set encapsulation type.  Must set either vxlan or vlan.
        "encap": {
            // Encapsulate traffic with VXLAN.
            "vxlan" : {
                // The name of the tunnel interface in OVS
                "encap-iface": "br0_vxlan0",

                // The name of the interface whose IP should be used
                // as the source IP in encapsulated traffic.
                "uplink-iface": "eth0.4093",

                // The vlan tag, if any, used on the uplink interface.
                // Set to zero or omit if the uplink is untagged.
                "uplink-vlan": 4093,

                // The IP address used for the destination IP in
                // the encapsulated traffic.  This should be an
                // anycast IP address understood by the upstream
                // stitched-mode fabric.
                "remote-ip": "10.0.0.32"
            }
        },
        // Configure forwarding policy
```

(continues on next page)

(continued from previous page)

```
        "forwarding": {
            // Configure the virtual distributed router
            "virtual-router": {
                // Enable virtual distributed router.  Set to true
                // to enable or false to disable.  Default true.
                "enabled": true,

                // Override MAC address for virtual router.
                // Default is "00:22:bd:f8:19:ff"
                "mac": "00:22:bd:f8:19:ff",

                // Configure IPv6-related settings for the virtual
                // router
                "ipv6" : {
                    // Send router advertisement messages in
                    // response to router solicitation requests as
                    // well as unsolicited advertisements.
                    "router-advertisement": true
                }
            },

            // Configure virtual distributed DHCP server
            "virtual-dhcp": {
                // Enable virtual distributed DHCP server.  Set to
                // true to enable or false to disable.  Default
                // true.
                "enabled": true,

                // Override MAC address for virtual dhcp server.
                // Default is "00:22:bd:f8:19:ff"
                "mac": "00:22:bd:f8:19:ff"
            }
        },

        // Location to store cached IDs for managing flow state
        "flowid-cache-dir": "DEFAULT_FLOWID_CACHE_DIR"
    }
}
```

Finally, enable the agent service:

```
# systemctl enable agent-ovs
ln -s '/usr/lib/systemd/system/agent-ovs.service' '/etc/systemd/system/multi-user.
↪target.wants/agent-ovs.service'
# systemctl start agent-ovs
# systemctl status agent-ovs
agent-ovs.service - Opflex OVS Agent
   Loaded: loaded (/usr/lib/systemd/system/agent-ovs.service; enabled)
   Active: active (running) since Mon 2014-12-15 10:03:42 PST; 5min ago
 Main PID: 6062 (agent_ovs)
   CGroup: /system.slice/agent-ovs.service
           └─6062 /usr/bin/agent_ovs
```

The agent is now running and ready to enforce policy. You can add endpoints to the local VM hosts using the OpFlex Group-based policy plugin from OpenStack, or manually.

TSDR Installation Guide

This document is for the user to install the artifacts that are needed for using Time Series Data Repository (TSDR) functionality in the ODL Controller by enabling either an HSQLDB, HBase, or Cassandra Data Store.

Overview

The Time Series Data Repository (TSDR) project in OpenDaylight (ODL) creates a framework for collecting, storing, querying, and maintaining time series data in the OpenDaylight SDN controller. Please refer to the User Guide for the detailed description of the functionality of the project and how to use the corresponding features provided in TSDR.

Pre Requisites for Installing TSDR

The software requirements for TSDR HBase Data Store are as follows:

- In the case when the user chooses HBase or Cassandra data store, besides the software that ODL requires, we also require HBase and Cassandra database running in single node deployment scenario.

No additional software is required for the HSQLDB Data Stores.

Preparing for Installation

- When using HBase data store, download HBase from the following website:

  http://archive.apache.org/dist/hbase/hbase-0.94.15/

- When using Cassandra data store, download Cassandra from the following website:

  http://www.eu.apache.org/dist/cassandra/2.1.10/

- No additional steps are required to install the TSDR HSQL Data Store.

Installing TSDR Data Stores

Installing HSQLDB Data Store

Once OpenDaylight distribution is up, from karaf console install the HSQLDB data store using the following command:

```
feature:install odl-tsdr-hsqldb-all
```

This will install hsqldb related dependency features (and can take sometime) as well as OpenFlow statistics collector before returning control to the console.

Installing HBase Data Store

Installing TSDR HBase Data Store contains two steps:

1. Installing HBase server, and
2. Installing TSDR HBase Data Store features from ODL Karaf console.

In this release, we only support HBase single node running together on the same machine as OpenDaylight. Therefore, follow the steps to download and install HBase server onto the same machine as where OpenDaylight is running:

1. Create a folder in Linux operating system for the HBase server. For example, create an hbase directory under `/usr/lib`:

```
mkdir /usr/lib/hbase
```

2. Unzip the downloaded HBase server tar file.

   Run the following command to unzip the installation package:

```
tar xvf <hbase-installer-name>  /usr/lib/hbase
```

3. Make proper changes in hbase-site.xml

   (a) Under <hbase-install-directory>/conf/, there is a hbase-site.xml. Although it is not recommended, an experienced user with HBase can modify the data directory for hbase server to store the data.

   (b) Modify the value of the property with name "hbase.rootdir" in the file to reflect the desired file directory for storing hbase data.

   The following is an example of the file:

```
<configuration>
  <property>
    <name>hbase.rootdir</name>
    <value>file:///usr/lib/hbase/data</value>
  </property>
  <property>
    <name>hbase.zookeeper.property.dataDir</name>
    <value>/usr/lib/hbase/zookeeper</value>
  </property>
</configuration>
```

4. start hbase server:

```
cd <hbase-installation-directory>
./start-hbase.sh
```

5. start hbase shell:

```
cd <hbase-insatllation-directory>
./hbase shell
```

6. start Karaf console
7. install hbase data store feature from Karaf console:

```
feature:install odl-tsdr-hbase
```

Installing Cassandra Data Store

Installing TSDR Cassandra Data Store contains two steps:

1. Installing Cassandra server, and
2. Installing TSDR Cassandra Data Store features from ODL Karaf console.

In this release, we only support Cassadra single node running together on the same machine as OpenDaylight. Therefore, follow these steps to download and install Cassandra server onto the same machine as where OpenDaylight is running:

1. Install Cassandra (latest stable version) by downloading the zip file and untar the tar ball to cassandra/ directory on the testing machine:

```
mkdir cassandra
wget http://www.eu.apache.org/dist/cassandra/2.1.10/apache-cassandra-2.1.10-bin.
↪tar.gz[2.1.10 is current stable version, it can vary]
mv apache-cassandra-2.1.10-bin.tar.gz cassandra/
cd cassandra
tar -xvzf apache-cassandra-2.1.10-bin.tar.gz
```

2. Start Cassandra from cassandra directory by running:

```
./apache-cassandra-2.1.10/bin/cassandra
```

3. Start cassandra shell by running:

```
./apache-cassandra-2.1.10/bin/cqlsh
```

4. Start Karaf according to the instructions above.
5. Install Cassandra data store feature from Karaf console:

```
feature:install odl-tsdr-cassandra
```

Verifying your Installation

After the TSDR data store is installed, no matter whether it is HBase data store, Cassandra data store, or HSQLDB data store, the user can verify the installation with the following steps.

1. Verify if the following two TSDR commands are available from Karaf console:

```
tsdr:list
tsdr:purgeAll
```

2. Verify if OpenFlow statistics data can be received successfully:
   (a) Run "feature:install odl-tsdr-openflow-statistics-collector" from Karaf.
   (b) Run mininet to connect to ODL controller. For example, use the following command to start a three node topology:

```
mn --topo single,3  --controller 'remote,ip=172.17.252.210,port=6653' --
↪switch ovsk,protocols=OpenFlow13
```

   (c) From Karaf console, the user should be able to retrieve the statistics data of OpenFlow statistics data from the console:

```
tsdr:list FLOWSTATS
```

Troubleshooting

Check the `../data/log/karaf.log` for any exception related to TSDR features.

Post Installation Configuration

Post Installation Configuration for HSQLDB Data Store

The feature installation takes care of automated configuration of the datasource by installing a file in <install folder>/etc named org.ops4j.datasource-metric.cfg. This contains the default location of <install folder>/tsdr where the HSQLDB datastore files are stored. If you want to change the default location of the datastore files to some other location update the last portion of the url property in the org.ops4j.datasource-metric.cfg and then restart the Karaf container.

Post Installation Configuration for HBase Data Store

Please refer to HBase Data Store User Guide.

Post Installation Configuration for Cassandra Data Store

There is no post configuration for TSDR Cassandra data store.

Upgrading From a Previous Release

The HBase data store was supported in the previous release as well as in this release. However, we do not support data store upgrade for HBase data store. The user needs to reinstall TSDR and start to collect data in TSDR HBase datastore after the installation.

HSQLDB and Cassandra are new data stores introduced in this release. Therefore, upgrading from previous release does not apply in these two data store scenarios.

Uninstalling TSDR Data Stores

To uninstall TSDR HSQLDB data store

To uninstall the TSDR functionality with the default store, you need to do the following from karaf console:

```
feature:uninstall odl-tsdr-hsqldb-all
feature:uninstall odl-tsdr-core
feature:uninstall odl-tsdr-hsqldb
feature:uninstall odl-tsdr-openflow-statistics-collector
```

It is recommended to restart the Karaf container after the uninstallation of the TSDR functionality with the default store.

To uninstall TSDR HBase Data Store

To uninstall the TSDR functionality with the HBase data store,

- Uninstall HBase data store related features from karaf console:

```
feature:uninstall odl-tsdr-hbase
feature:uninstall odl-tsdr-core
```

- stop hbase server:

```
cd <hbase-installation-directory>
./stop-hbase.sh
```

- remove the file directory that contains the HBase server installation:

```
rm -r <hbase-installation-directory>
```

It is recommended to restart the Karaf container after the uninstallation of the TSDR data store.

To uninstall TSDR Cassandra Data Store

To uninstall the TSDR functionality with the Cassandra store,

- uninstall cassandra data store related features following from karaf console:

```
feature:uninstall odl-tsdr-cassandra
feature:uninstall odl-tsdr-core
```

- stop cassandra database:

```
ps auwx | grep cassandra
sudo kill pid
```

- remove the cassandra installation files:

```
rm <cassandra-installation-directory>
```

It is recommended to restart the Karaf container after uninstallation of the TSDR data store.

VTN Installation Guide

Overview

OpenDaylight Virtual Tenant Network (VTN) is an application that provides multi-tenant virtual network on an SDN controller.

Conventionally, huge investment in the network systems and operating expenses are needed because the network is configured as a silo for each department and system. Therefore various network appliances must be installed for each tenant and those boxes cannot be shared with others. It is a heavy work to design, implement and operate the entire complex network.

The uniqueness of VTN is a logical abstraction plane. This enables the complete separation of logical plane from physical plane. Users can design and deploy any desired network without knowing the physical network topology or bandwidth restrictions.

VTN allows the users to define the network with a look and feel of conventional L2/L3 network. Once the network is designed on VTN, it will automatically be mapped into underlying physical network, and then configured on the individual switch leverage SDN control protocol. The definition of logical plane makes it possible not only to hide the complexity of the underlying network but also to better manage network resources. It achieves reducing reconfiguration time of network services and minimizing network configuration errors. OpenDaylight Virtual Tenant Network (VTN) is an application that provides multi-tenant virtual network on an SDN controller. It provides API for creating a common virtual network irrespective of the physical network.

It is implemented as two major components

- *VTN Manager*
- *VTN Coordinator*

VTN Manager

An OpenDaylight Plugin that interacts with other modules to implement the components of the VTN model. It also provides a REST interface to configure VTN components in OpenDaylight. VTN Manager is implemented as one plugin to the OpenDaylight. This provides a REST interface to create/update/delete VTN components. The user command in VTN Coordinator is translated as REST API to VTN Manager by the OpenDaylight Driver component. In addition to the above mentioned role, it also provides an implementation to the OpenStack L2 Network Functions API.

VTN Coordinator

The VTN Coordinator is an external application that provides a REST interface for an user to use OpenDaylight VTN Virtualization. It interacts with VTN Manager plugin to implement the user configuration. It is also capable of multiple OpenDaylight orchestration. It realizes VTN provisioning in OpenDaylight instances. In the OpenDaylight architecture VTN Coordinator is part of the network application, orchestration and services layer. VTN Coordinator will use the REST interface exposed by the VTN Manger to realize the virtual network using OpenDaylight. It uses OpenDaylight APIs (REST) to construct the virtual network in OpenDaylight instances. It provides REST APIs for northbound VTN applications and supports virtual networks spanning across multiple OpenDaylight by coordinating across OpenDaylight.

Preparing for Installation

VTN Manager

Follow the instructions in *Installing OpenDaylight*.

VTN Coordinator

1. Arrange a physical/virtual server with any one of the supported 64-bit OS environment.
   - RHEL 7
   - CentOS 7
   - Fedora 20 / 21 / 22
2. Install these packages:

```
yum install perl-Digest-SHA uuid libxslt libcurl unixODBC json-c bzip2
rpm -ivh http://yum.postgresql.org/9.3/redhat/rhel-6-x86_64/pgdg-redhat93-9.3-3.
↪noarch.rpm
yum install postgresql93-libs postgresql93 postgresql93-server postgresql93-
↪contrib postgresql93-odbc
```

Installing VTN

VTN Manager

Install Feature:

```
feature:install odl-vtn-manager-neutron odl-vtn-manager-rest
```

Note: The above command will install all features of VTN Manager. You can install only REST or Neutron also.

VTN Coordinator

- To get the Boron distribution for VTN coordinator download the latest "tar.bz2" file from the below link:

```
https://nexus.opendaylight.org/content/repositories/public/org/opendaylight/vtn/
↪distribution.vtn-coordinator/6.3.0-Boron/
```

- Run the below command to extract VTN Coordinator from the tar.bz2 file:

```
tar -C/ -jxvf distribution.vtn-coordinator-6.3.0-Boron-bin.tar.bz2
```

This will install VTN Coordinator to /usr/local/vtn directory. The name of the tar.bz2 file name varies depending on the version. Please give the same tar.bz2 file name which is there in your directory.

- Configuring database for VTN Coordinator:

```
/usr/local/vtn/sbin/db_setup
```

- To start the Coordinator:

```
/usr/local/vtn/bin/vtn_start
```

Using VTN REST API:

Get the version of VTN REST API using the below command, and make sure the setup is working:

```
curl --user admin:adminpass -H 'content-type: application/json' -X GET http://<VTN_
↪COORDINATOR_IP_ADDRESS>:8083/vtn-webapi/api_version.json
```

The response should be like this, but version might differ:

```
{"api_version":{"version":"V1.2"}}
```

Verifying your Installation

VTN Manager

- In the karaf prompt, type the below command to ensure that vtn packages are installed:

```
feature:list | grep vtn
```

- Run any VTN Manager REST API:

```
curl --user "admin":"admin" -H "Accept: application/json" -H "Content-type:␣
↪application/json" -X GET http://localhost:8181/restconf/operational/vtn:vtns
```

VTN Coordinator

```
ps -ef | grep unc will list all the vtn apps
Run any REST API for VTN Coordinator version
```

Uninstalling VTN

VTN Manager

```
feature:uninstall odl-vtnmanager-all
```

VTN Coordinator

1. Stop VTN:

   ```
   /usr/local/vtn/bin/vtn_stop
   ```

2. Remove the `usr/local/vtn` folder

YANG IDE Installation Guide

Overview

The YANG IDE project provides an Eclipse plugin for viewing and editing Yang model files. When you create a "Yang Project" using the plugin, it creates a small Maven project with a POM file (pom.xml) that references the appropriate OpenDaylight dependencies, along with a sample Yang model file (acme-system.yang).

Pre Requisites for Installing YANG IDE

- YANG IDE has the same hardware requirements as the Eclipse IDE, which is about the same as the hardware requirements for Java 7.
- At least Java 7 is required to run Eclipse (also an obvious requirement), but Java 8 will be required if you are building an application using OpenDaylight, and Java 8 is recommended anyway.

Preparing for Installation

As soon as at least Java 7 (Java 8 preferred) and Eclipse are installed, and Eclipse is running, you can install YANG IDE.

You can find the Oracle Java installer at http://www.oracle.com/technetwork/java/javase/downloads/index.html .

The Eclipse installer can be found at http://www.eclipse.org/downloads/ . You should select the "Eclipse IDE for Java Developers", and make sure you select the installer for the correct platform (for instance, 32-bit or 64-bit).

Installing YANG IDE

The YANG IDE plugin can be installed by using the public update site URL provided, which is https://nexus.opendaylight.org/content/sites/p2repos/org.opendaylight.yangide/release/ .

While in Eclipse, select "Help" from the menu bar and select "Install New Software . . . ". On the resulting "Install" dialog, click the "Add. . . " button. In that dialog, enter the update site URL as specified above and give it a name of "YANG IDE". Select the provided plugin and approve the license.

Eclipse will prompt you to restart Eclipse. Do that.

Installation is complete at this point.

Network Connections

If the installation failed with an indication that it could not reach the internet, then your work computer may be behind a firewall. You will need to go to the "Network Connections" section of the Eclipse preferences (Menubar: "Window"->"Preferences"->"General"->"Network Connections").

Before you make these changes, you will need to know the host and port of your outbound proxy server.

On the "Network Connections" page, you should select "Manual" in the "Active Provider" dropdown, then edit the "HTTP" and "HTTPS" rows in the table, setting the host and port of the outbound proxy server.

If the proxy server requires authentication, turn on the "Requires Authentication" checkbox and enter the required userid and password fields. If you do not know whether your proxy server requires authentication, it probably does not.

Verifying your Installation

This is not really a "usage guide", but following these steps will verify that the plugin was properly installed.

When installation is complete, you can select "File" from the menu bar, then "New", then "Other" (you may have a keyboard shortcut for "Ctrl+n" for this).

In the "New" dialog, you can enter "yang" in the field under the "Wizards" label, which starts out with the content of "type filter text". That will limit the list to the "YANG" folder and the two choices of "YANG File" and "YANG Project". Select the "YANG Project" option and click "Next".

On the "New Yang Project" dialog, you may see a wizard page titled "Specify YANG Code Generators Parameters". Do not change anything on that page and click "Next".

On the next wizard page, with the title "Select project name and location", check the "Create a simple project" checkbox and click "Next".

On that dialog, enter anything you want in the "Group Id" field. Enter a project name (again, whatever you want for now) in the "Artifact Id" field and click "Finish". No other fields on the page need to be changed.

The dialog will now go away and Eclipse will create the project, which you should see in either the "Package Explorer" or "Project Explorer" view, on the left side.

Click the arrow just left of the project name to expand the contents of the project.

In that resulting list, there are only two entries that you will ever care about. One is "src/main/yang", which is where you will store the Yang model files, and the "pom.xml" file, which is where you will enter dependencies for Yang model files to import. If you will not be importing any Yang model files, or you will only be importing other Yang model files in your own project, then you will never have to do anything with the "pom.xml" file.

Click the arrow to the left of the "src/main/yang" entry to expand that.

You should see a "acme-system.yang" file, which the plugin created by default. Double-click on that entry to open the file in the editor.

Troubleshooting

If Eclipse fails to start up initially, then there is something wrong with either the Java installation or the Eclipse installation.

You can determine whether Java is installed correctly by opening a shell or command window and entering "java -version" and verifying whether the output corresponds to the version of Java that you installed.

If the Java installation seems fine, but Eclipse still fails to start up, you can ask questions on the #eclipse IRC channel, or post questions on the "Newcomers" forum at http://www.eclipse.org/forums/ .

If Java and Eclipse seem to be fine, but the YANG IDE is having problems, ask questions on the "yangide-dev" mailing list.

Post Installation Configuration

Setting Proxy Used For Maven

If your work computer sits behind a firewall, you will have had to put information about your firewall in the "Network Connections" section of the Eclipse preferences. That would have allowed you to at least obtain the plugin and install it into Eclipse.

Much of the functionality of YANG IDE uses Maven internally. You do not need to be a Maven expert to use this functionality, but you will need to add a few more lines of configuration so that Maven can get through the firewall. Maven, even when running inside Eclipse, as it is when you are using YANG IDE, does not use the Eclipse "Network Connection" settings to reach the internet. You have to set the proxy server information in a different place for Maven.

Maven looks for a file at `$HOME/.m2/settings.xml` (Linux) or `%HOME%\.m2\settings.xml` (Windows). If the `.m2` folder does not exist, you will need to create it. If the "settings.xml" file does not exist, you should create it with the following contents:

```
<?xml version="1.0" encoding="UTF-8"?>
<settings xmlns="http://maven.apache.org/SETTINGS/1.0.0"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://maven.apache.org/SETTINGS/1.0.0 http://maven.apache.org/
↪xsd/settings-1.0.0.xsd">
  <proxies>
    <proxy>
      <id>proxy</id>
      <active>true</active>
      <protocol>http</protocol>
      <host>FULLY QUALIFIED NAME OF PROXY HOST</host>
      <port>PROXY PORT</port>
    </proxy>
    <proxy>
      <id>proxy2</id>
      <active>true</active>
      <protocol>https</protocol>
      <host>FULLY QUALIFIED NAME OF PROXY HOST</host>
      <port>PROXY PORT</port>
    </proxy>
  </proxies>
</settings>
```

Replace "FULLY QUALIFIED NAME OF PROXY HOST" and "PROXY PORT" with the host and port of your proxy server.

If the "settings.xml" file already existed, then you will need to edit it, inserting the "proxies" element from the above sample at an appropriate place.

Upgrading From a Previous Release

If you already had the "YANG IDE" plugin from "Xored", you will need to uninstall that plugin before you install this one.

Uninstalling YANG IDE

Uninstalling the YANG IDE plugin is the same as uninstalling any other Eclipse plugin.

Click on the "Help" menu item and select "Installation Details". That list will have all the plugins you have installed (or that came with the distribution). To uninstall YANG IDE, you will need to select four entries from that list:

- "m2e connector for YANG"
- "m2e connector for YANG Developer Resources"
- "YANG IDE"
- "YANG IDE Developer Resources"

Use the Control key to select multiple entries in this list. When all four entries are selected, click the "Uninstall" button. The next dialog shows what you selected and asks you to confirm with the "Finish" button.

It will then uninstall the plugin and prompt you to restart Eclipse. When Eclipse restarts, the uninstall process is complete.

1.1.12 Common OpenDaylight Features

OpenDaylight User Interface (DLUX)

This section introduces you to the OpenDaylight User Experience (DLUX) application.

Getting Started with DLUX

DLUX provides a number of different Karaf features, which you can enable and disable separately. In Boron they are:

1. odl-dlux-core
2. odl-dlux-node
3. odl-dlux-yangui
4. odl-dlux-yangvisualizer

Logging In

To log in to DLUX, after installing the application:

1. Open a browser and enter the login URL http://<your-karaf-ip>:8181/index.html in your browser (Chrome is recommended).
2. Login to the application with your username and password credentials.

Note: OpenDaylight's default credentials are *admin* for both the username and password.

Working with DLUX

After you login to DLUX, if you enable only odl-dlux-core feature, you will see only topology application available in the left pane.

Note: To make sure topology displays all the details, enable the odl-l2switch-switch feature in Karaf.

DLUX has other applications such as node, yang UI and those apps won't show up, until you enable their features odl-dlux-node and odl-dlux-yangui respectively in the Karaf distribution.

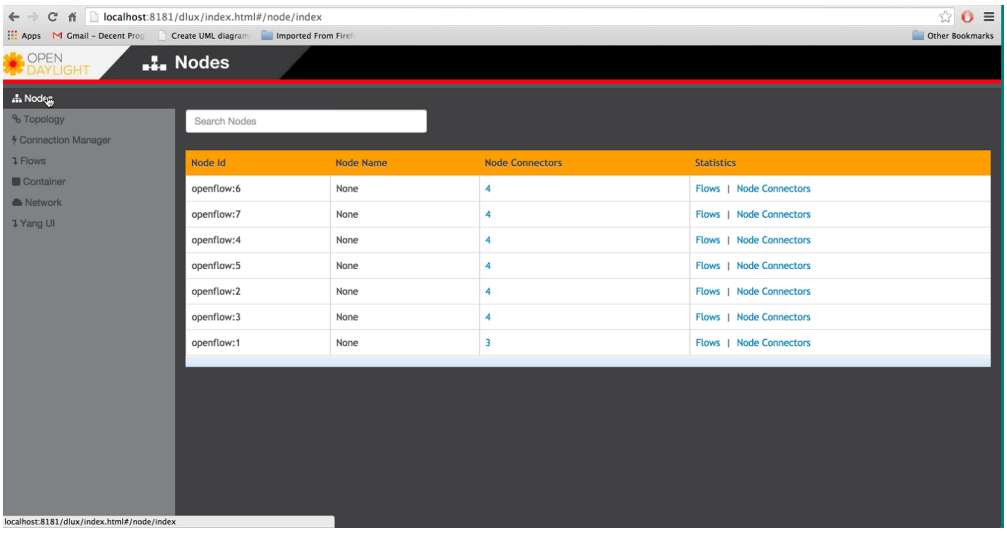

Fig. 1: DLUX Modules

Note: If you install your application in dlux, they will also show up on the left hand navigation after browser page refresh.

Viewing Network Statistics

The *Nodes* module on the left pane enables you to view the network statistics and port information for the switches in the network.

To use the *Nodes* module:

1. Select *Nodes* on the left pane. The right pane displays atable that lists all the nodes, node connectors and the statistics.

2. Enter a node ID in the *Search Nodes* tab to search by node connectors.
3. Click on the *Node Connector* number to view details such as port ID, port name, number of ports per switch, MAC Address, and so on.
4. Click *Flows* in the Statistics column to view Flow Table Statistics for the particular node like table ID, packet match, active flows and so on.
5. Click *Node Connectors* to view Node Connector Statistics for the particular node ID.

Viewing Network Topology

The Topology tab displays a graphical representation of network topology created.

Note: DLUX does not allow for editing or adding topology information. The topology is generated and edited in other modules, e.g., the OpenFlow plugin. OpenDaylight stores this information in the MD-SAL datastore where DLUX can read and display it.

To view network topology:

1. Select *Topology* on the left pane. You will view the graphical representation on the right pane. In the diagram blue boxes represent the switches, the black represents the hosts available, and lines represents how the switches and hosts are connected.
2. Hover your mouse on hosts, links, or switches to view source and destination ports.
3. Zoom in and zoom out using mouse scroll to verify topology for larger topologies.

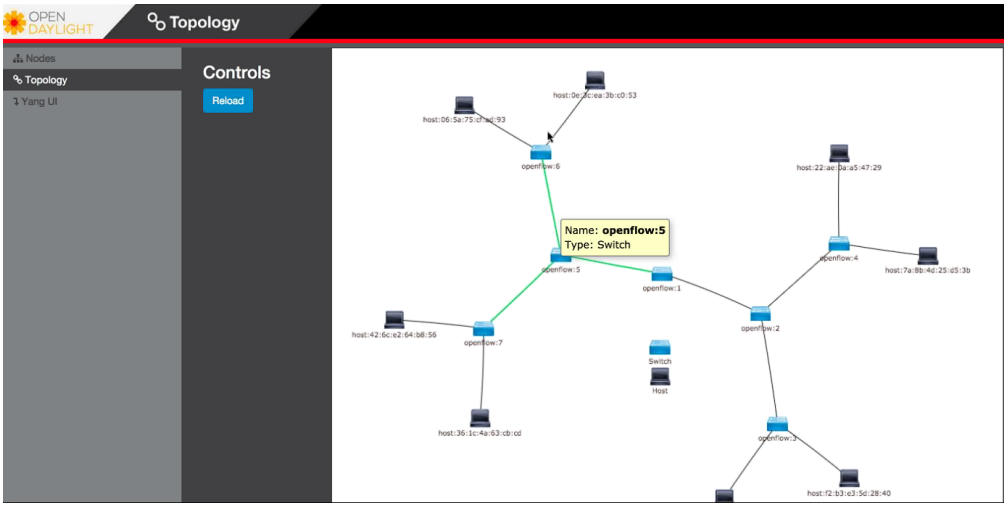

Fig. 2: Topology Module

Interacting with the YANG-based MD-SAL datastore

The *Yang UI* module enables you to interact with the YANG-based MD-SAL datastore. For more information about YANG and how it interacts with the MD-SAL datastore, see the *Controller* and *YANG Tools* section of the *OpenDaylight Developer Guide*.

To use Yang UI:

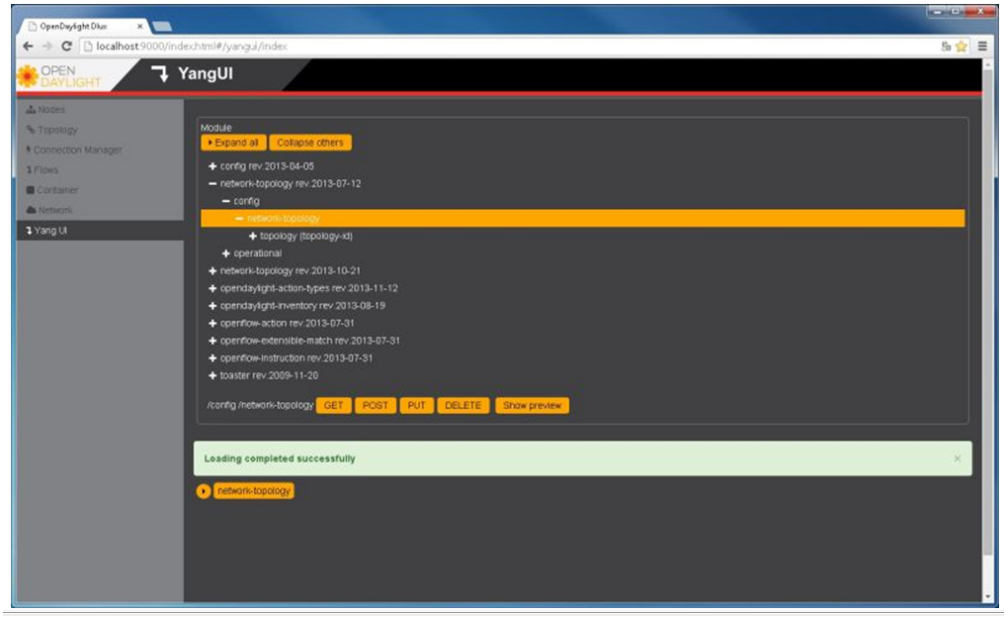

Fig. 3: Yang UI

1. Select *Yang UI* on the left pane. The right pane is divided in two parts.
2. The top part displays a tree of APIs, subAPIs, and buttons to call possible functions (GET, POST, PUT, and DELETE).

   Note: Not every subAPI can call every function. For example, subAPIs in the *operational* store have GET functionality only.

   Inputs can be filled from OpenDaylight when existing data from OpenDaylight is displayed or can be filled by user on the page and sent to OpenDaylight.

   Buttons under the API tree are variable. It depends on subAPI specifications. Common buttons are:

   - GET to get data from OpenDaylight,
   - PUT and POST for sending data to OpenDaylight for saving
   - DELETE for sending data to OpenDaylight for deleting.

   You must specify the xpath for all these operations. This path is displayed in the same row before buttons and it may include text inputs for specific path element identifiers.
3. The bottom part of the right pane displays inputs according to the chosen subAPI.
   - Lists are handled as a special case. For example, a device can store multiple flows. In this case "flow" is name of the list and every list element is identified by a unique key value. Elements of a list can, in turn, contain other lists.
   - In Yang UI, each list element is rendered with the name of the list it belongs to, its key, its value, and a button for removing it from the list.
4. After filling in the relevant inputs, click the *Show Preview* button under the API tree to display request that will be sent to OpenDaylight. A pane is displayed on the right side with text of request when some input is filled.

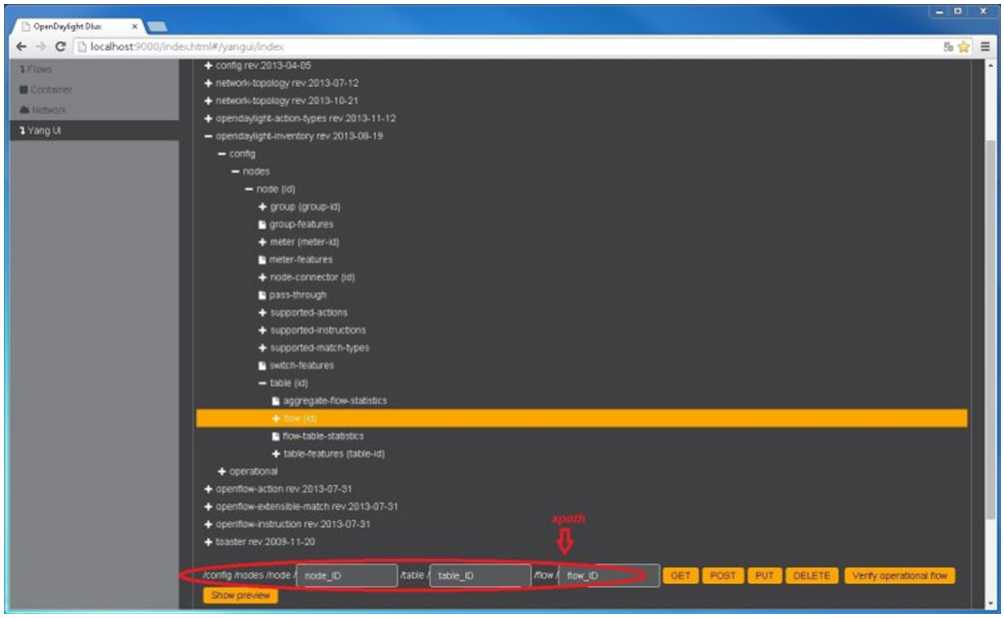

Fig. 4: Yang API Specification

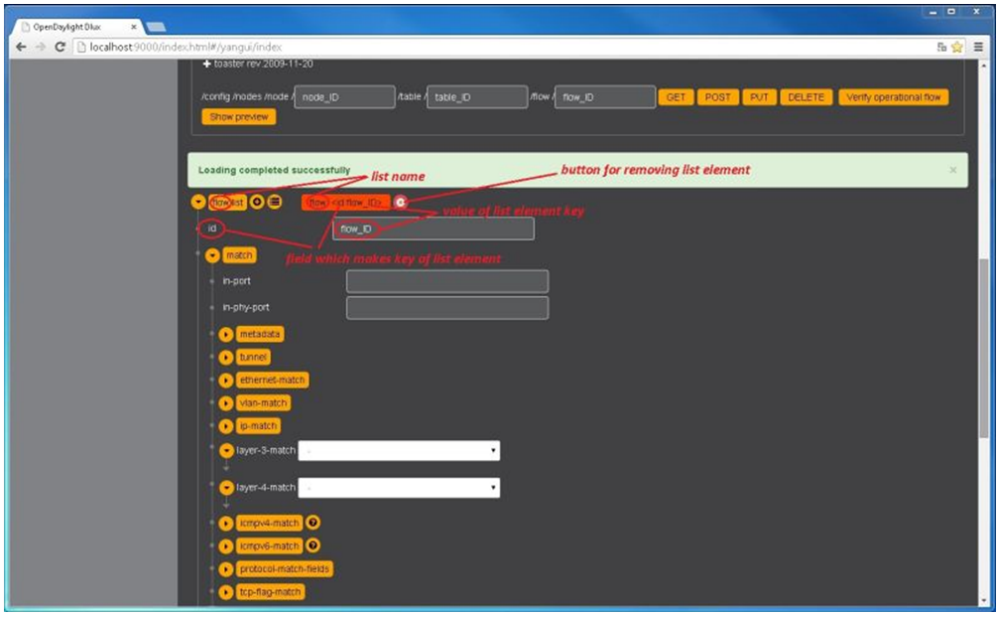

Fig. 5: Yang UI API Specification

Displaying Topology on the Yang UI

To display topology:

1. Select subAPI network-topology <topology revision number> == > operational == > network-topology.
2. Get data from OpenDaylight by clicking on the "GET" button.
3. Click *Display Topology*.

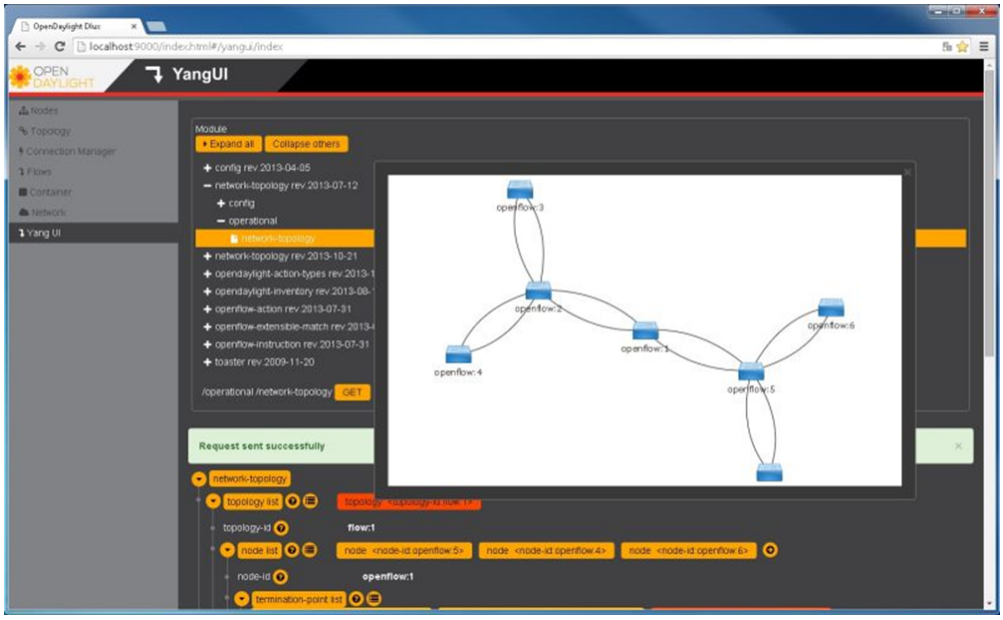

Fig. 6: DLUX Yang Topology

Configuring List Elements on the Yang UI

Lists in Yang UI are displayed as trees. To expand or collapse a list, click the arrow before name of the list. To configure list elements in Yang UI:

1. To add a new list element with empty inputs use the plus icon-button **+** that is provided after list name.
2. To remove several list elements, use the *X* button that is provided after every list element.
3. In the YANG-based data store all elements of a list must have a unique key. If you try to assign two or more elements the same key, a warning icon *!* is displayed near their name buttons.
4. When the list contains at least one list element, after the + icon, there are buttons to select each individual list element. You can choose one of them by clicking on it. In addition, to the right of the list name, there is a button which will display a vertically scrollable pane with all the list elements.

Setting Up Clustering

Clustering Overview

Clustering is a mechanism that enables multiple processes and programs to work together as one entity. For example, when you search for something on google.com, it may seem like your search request is processed by only one web

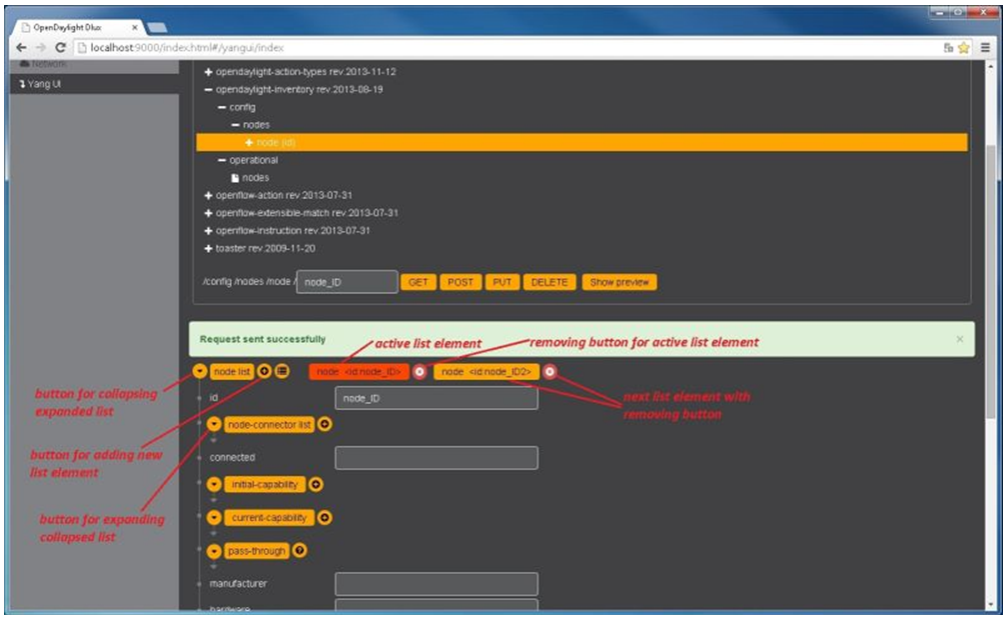

Fig. 7: DLUX List Elements

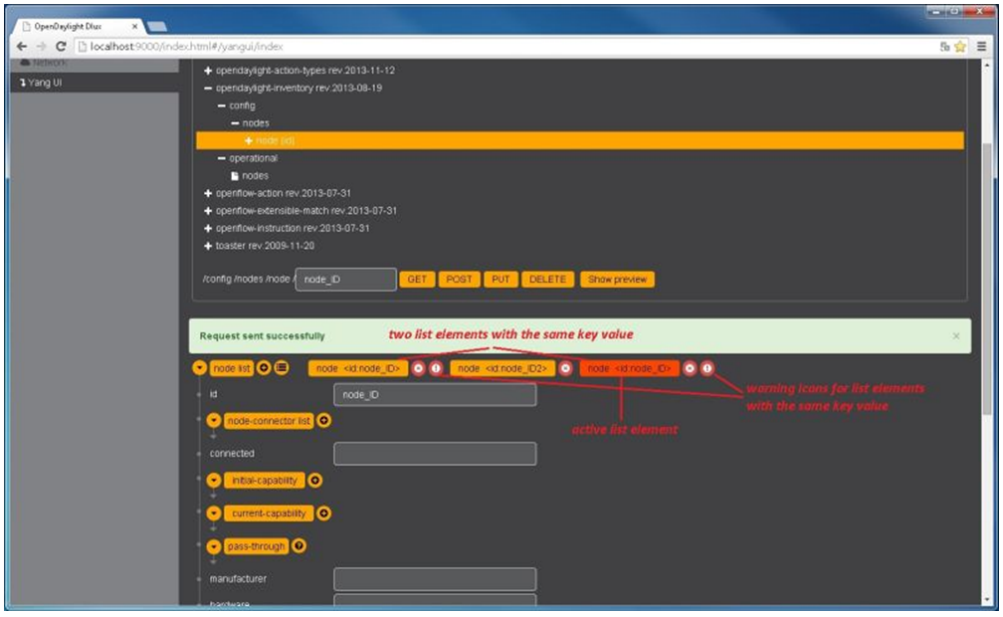

Fig. 8: DLUX List Warnings

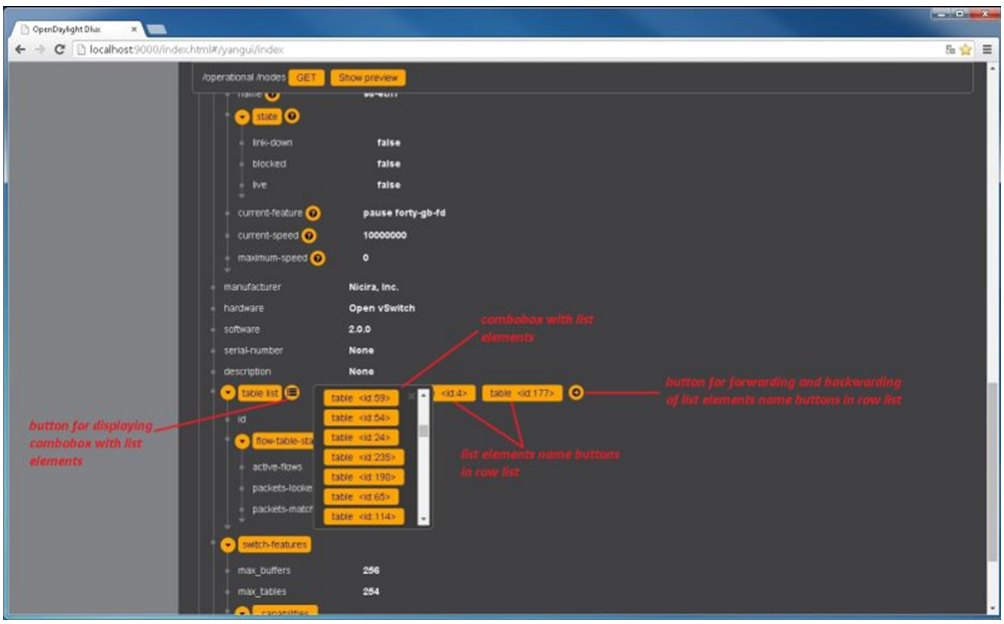

Fig. 9: DLUX List Button

server. In reality, your search request is processed by may web servers connected in a cluster. Similarly, you can have multiple instances of OpenDaylight working together as one entity.

Advantages of clustering are:

- Scaling: If you have multiple instances of OpenDaylight running, you can potentially do more work and store more data than you could with only one instance. You can also break up your data into smaller chunks (shards) and either distribute that data across the cluster or perform certain operations on certain members of the cluster.
- High Availability: If you have multiple instances of OpenDaylight running and one of them crashes, you will still have the other instances working and available.
- Data Persistence: You will not lose any data stored in OpenDaylight after a manual restart or a crash.

The following sections describe how to set up clustering on both individual and multiple OpenDaylight instances.

Multiple Node Clustering

The following sections describe how to set up multiple node clusters in OpenDaylight.

Deployment Considerations

To implement clustering, the deployment considerations are as follows:

- To set up a cluster with multiple nodes, we recommend that you use a minimum of three machines. You can set up a cluster with just two nodes. However, if one of the two nodes fail, the cluster will not be operational.

Note: This is because clustering in OpenDaylight requires a majority of the nodes to be up and one node cannot be a majority of two nodes.

- Every device that belongs to a cluster needs to have an identifier. OpenDaylight uses the node's `role` for this purpose. After you define the first node's role as *member-1* in the `akka.conf` file, OpenDaylight uses *member-1* to identify that node.
- Data shards are used to contain all or a certain segment of a OpenDaylight's MD-SAL datastore. For example, one shard can contain all the inventory data while another shard contains all of the topology data.

  If you do not specify a module in the `modules.conf` file and do not specify a shard in `module-shards.conf`, then (by default) all the data is placed in the default shard (which must also be defined in `module-shards.conf` file). Each shard has replicas configured. You can specify the details of where the replicas reside in `module-shards.conf` file.
- If you have a three node cluster and would like to be able to tolerate any single node crashing, a replica of every defined data shard must be running on all three cluster nodes.

  Note: This is because OpenDaylight's clustering implementation requires a majority of the defined shard replicas to be running in order to function. If you define data shard replicas on two of the cluster nodes and one of those nodes goes down, the corresponding data shards will not function.

- If you have a three node cluster and have defined replicas for a data shard on each of those nodes, that shard will still function even if only two of the cluster nodes are running. Note that if one of those remaining two nodes goes down, the shard will not be operational.
- It is recommended that you have multiple seed nodes configured. After a cluster member is started, it sends a message to all of its seed nodes. The cluster member then sends a join command to the first seed node that responds. If none of its seed nodes reply, the cluster member repeats this process until it successfully establishes a connection or it is shut down.
- After a node is unreachable, it remains down for configurable period of time (10 seconds, by default). Once a node goes down, you need to restart it so that it can rejoin the cluster. Once a restarted node joins a cluster, it will synchronize with the lead node automatically.

Clustering Scripts

OpenDaylight includes some scripts to help with the clustering configuration.

Note: Scripts are stored in the OpenDaylight distribution/bin folder, and maintained in the distribution project repository in the folder distribution-karaf/src/main/assembly/bin/.

Configure Cluster Script

This script is used to configure the cluster parameters (e.g. akka.conf, module-shards.conf) on a member of the controller cluster. The user should restart the node to apply the changes.

Note: The script can be used at any time, even before the controller is started for the first time.

Usage:

```
bin/configure_cluster.sh <index> <seed_nodes_list>
```

- index: Integer within 1..N, where N is the number of seed nodes. This indicates which controller node (1..N) is configured by the script.
- seed_nodes_list: List of seed nodes (IP address), separated by comma or space.

The IP address at the provided index should belong to the member executing the script. When running this script on multiple seed nodes, keep the seed_node_list the same, and vary the index from 1 through N.

Optionally, shards can be configured in a more granular way by modifying the file "custom_shard_configs.txt" in the same folder as this tool. Please see that file for more details.

Example:

```
bin/configure_cluster.sh 2 192.168.0.1 192.168.0.2 192.168.0.3
```

The above command will configure the member 2 (IP address 192.168.0.2) of a cluster made of 192.168.0.1 192.168.0.2 192.168.0.3.

Setting Up a Multiple Node Cluster

To run OpenDaylight in a three node cluster, perform the following:

First, determine the three machines that will make up the cluster. After that, do the following on each machine:

1. Copy the OpenDaylight distribution zip file to the machine.
2. Unzip the distribution.
3. Open the following .conf files:
   - configuration/initial/akka.conf
   - configuration/initial/module-shards.conf
4. In each configuration file, make the following changes:

   Find every instance of the following lines and replace _127.0.0.1_ with the hostname or IP address of the machine on which this file resides and OpenDaylight will run:

   ```
   netty.tcp {
     hostname = "127.0.0.1"
   ```

   Note: The value you need to specify will be different for each node in the cluster.

5. Find the following lines and replace _127.0.0.1_ with the hostname or IP address of any of the machines that will be part of the cluster:

   ```
   cluster {
     seed-nodes = ["akka.tcp://opendaylight-cluster-data@${IP_OF_MEMBER1}:2550",
                   <url-to-cluster-member-2>,
                   <url-to-cluster-member-3>]
   ```

6. Find the following section and specify the role for each member node. Here we assign the first node with the *member-1* role, the second node with the *member-2* role, and the third node with the *member-3* role:

   ```
   roles = [
     "member-1"
   ]
   ```

Note: This step should use a different role on each node.

7. Open the configuration/initial/module-shards.conf file and update the replicas so that each shard is replicated to all three nodes:

```
replicas = [
    "member-1",
    "member-2",
    "member-3"
]
```

   For reference, view a sample config files <<_sample_config_files,below>>.

8. Move into the +<karaf-distribution-directory>/bin+ directory.

9. Run the following command:

```
JAVA_MAX_MEM=4G JAVA_MAX_PERM_MEM=512m ./karaf
```

10. Enable clustering by running the following command at the Karaf command line:

```
feature:install odl-mdsal-clustering
```

OpenDaylight should now be running in a three node cluster. You can use any of the three member nodes to access the data residing in the datastore.

Sample Config Files

Sample `akka.conf` file:

```
odl-cluster-data {
  bounded-mailbox {
    mailbox-type = "org.opendaylight.controller.cluster.common.actor.
↪MeteredBoundedMailbox"
    mailbox-capacity = 1000
    mailbox-push-timeout-time = 100ms
  }

  metric-capture-enabled = true

  akka {
    loglevel = "DEBUG"
    loggers = ["akka.event.slf4j.Slf4jLogger"]

    actor {

      provider = "akka.cluster.ClusterActorRefProvider"
      serializers {
                java = "akka.serialization.JavaSerializer"
                proto = "akka.remote.serialization.ProtobufSerializer"
              }

              serialization-bindings {
                  "com.google.protobuf.Message" = proto
```

(continues on next page)

(continued from previous page)

```
            }
    }
    remote {
      log-remote-lifecycle-events = off
      netty.tcp {
        hostname = "10.194.189.96"
        port = 2550
        maximum-frame-size = 419430400
        send-buffer-size = 52428800
        receive-buffer-size = 52428800
      }
    }

    cluster {
      seed-nodes = ["akka.tcp://opendaylight-cluster-data@10.194.189.96:2550",
                    "akka.tcp://opendaylight-cluster-data@10.194.189.98:2550",
                    "akka.tcp://opendaylight-cluster-data@10.194.189.101:2550"]

      auto-down-unreachable-after = 10s

      roles = [
        "member-2"
      ]

    }
  }
}

odl-cluster-rpc {
  bounded-mailbox {
    mailbox-type = "org.opendaylight.controller.cluster.common.actor.
↪MeteredBoundedMailbox"
    mailbox-capacity = 1000
    mailbox-push-timeout-time = 100ms
  }

  metric-capture-enabled = true

  akka {
    loglevel = "INFO"
    loggers = ["akka.event.slf4j.Slf4jLogger"]

    actor {
      provider = "akka.cluster.ClusterActorRefProvider"

    }
    remote {
      log-remote-lifecycle-events = off
      netty.tcp {
        hostname = "10.194.189.96"
        port = 2551
      }
    }

    cluster {
      seed-nodes = ["akka.tcp://opendaylight-cluster-rpc@10.194.189.96:2551"]
```

(continues on next page)

(continued from previous page)

```
      auto-down-unreachable-after = 10s
    }
  }
}
```

Sample `module-shards.conf` file:

```
module-shards = [
    {
        name = "default"
        shards = [
            {
                name="default"
                replicas = [
                    "member-1",
                    "member-2",
                    "member-3"
                ]
            }
        ]
    },
    {
        name = "topology"
        shards = [
            {
                name="topology"
                replicas = [
                    "member-1",
                    "member-2",
                    "member-3"
                ]
            }
        ]
    },
    {
        name = "inventory"
        shards = [
            {
                name="inventory"
                replicas = [
                    "member-1",
                    "member-2",
                    "member-3"
                ]
            }
        ]
    },
    {
         name = "toaster"
         shards = [
             {
                 name="toaster"
                 replicas = [
                    "member-1",
                    "member-2",
                    "member-3"
                 ]
```

(continues on next page)

(continued from previous page)

```
                }
            ]
        }
]
```

Cluster Monitoring

OpenDaylight exposes shard information via MBeans, which can be explored with JConsole, VisualVM, or other JMX clients, or exposed via a REST API using Jolokia, provided by the `odl-jolokia` Karaf feature. This is convenient, due to a significant focus on REST in OpenDaylight.

The basic URI that lists a schema of all available MBeans, but not their content itself is:

```
GET  /jolokia/list
```

To read the information about the shards local to the queried OpenDaylight instance use the following REST calls. For the config datastore:

```
GET  /jolokia/read/org.opendaylight.controller:type=DistributedConfigDatastore,
↪Category=ShardManager,name=shard-manager-config
```

For the operational datastore:

```
GET  /jolokia/read/org.opendaylight.controller:type=DistributedOperationalDatastore,
↪Category=ShardManager,name=shard-manager-operational
```

The output contains information on shards present on the node:

```
{
  "request": {
    "mbean": "org.opendaylight.controller:Category=ShardManager,name=shard-manager-
↪operational,type=DistributedOperationalDatastore",
    "type": "read"
  },
  "value": {
    "LocalShards": [
      "member-1-shard-default-operational",
      "member-1-shard-entity-ownership-operational",
      "member-1-shard-topology-operational",
      "member-1-shard-inventory-operational",
      "member-1-shard-toaster-operational"
    ],
    "SyncStatus": true,
    "MemberName": "member-1"
  },
  "timestamp": 1483738005,
  "status": 200
}
```

The exact names from the "LocalShards" lists are needed for further exploration, as they will be used as part of the URI to look up detailed info on a particular shard. An example output for the `member-1-shard-default-operational` looks like this:

```
{
  "request": {
    "mbean": "org.opendaylight.controller:Category=Shards,name=member-1-shard-default-
↪operational,type=DistributedOperationalDatastore",
    "type": "read"
  },
  "value": {
    "ReadWriteTransactionCount": 0,
    "SnapshotIndex": 4,
    "InMemoryJournalLogSize": 1,
    "ReplicatedToAllIndex": 4,
    "Leader": "member-1-shard-default-operational",
    "LastIndex": 5,
    "RaftState": "Leader",
    "LastCommittedTransactionTime": "2017-01-06 13:19:00.135",
    "LastApplied": 5,
    "LastLeadershipChangeTime": "2017-01-06 13:18:37.605",
    "LastLogIndex": 5,
    "PeerAddresses": "member-3-shard-default-operational: akka.tcp://opendaylight-
↪cluster-data@192.168.16.3:2550/user/shardmanager-operational/member-3-shard-default-
↪operational, member-2-shard-default-operational: akka.tcp://opendaylight-cluster-
↪data@192.168.16.2:2550/user/shardmanager-operational/member-2-shard-default-
↪operational",
    "WriteOnlyTransactionCount": 0,
    "FollowerInitialSyncStatus": false,
    "FollowerInfo": [
      {
        "timeSinceLastActivity": "00:00:00.320",
        "active": true,
        "matchIndex": 5,
        "voting": true,
        "id": "member-3-shard-default-operational",
        "nextIndex": 6
      },
      {
        "timeSinceLastActivity": "00:00:00.320",
        "active": true,
        "matchIndex": 5,
        "voting": true,
        "id": "member-2-shard-default-operational",
        "nextIndex": 6
      }
    ],
    "FailedReadTransactionsCount": 0,
    "StatRetrievalTime": "810.5 μs",
    "Voting": true,
    "CurrentTerm": 1,
    "LastTerm": 1,
    "FailedTransactionsCount": 0,
    "PendingTxCommitQueueSize": 0,
    "VotedFor": "member-1-shard-default-operational",
    "SnapshotCaptureInitiated": false,
    "CommittedTransactionsCount": 6,
    "TxCohortCacheSize": 0,
    "PeerVotingStates": "member-3-shard-default-operational: true, member-2-shard-
↪default-operational: true",
    "LastLogTerm": 1,
```


(continues on next page)

(continued from previous page)

```
    "StatRetrievalError": null,
    "CommitIndex": 5,
    "SnapshotTerm": 1,
    "AbortTransactionsCount": 0,
    "ReadOnlyTransactionCount": 0,
    "ShardName": "member-1-shard-default-operational",
    "LeadershipChangeCount": 1,
    "InMemoryJournalDataSize": 450
  },
  "timestamp": 1483740350,
  "status": 200
}
```

The output helps identifying shard state (leader/follower, voting/non-voting), peers, follower details if the shard is a leader, and other statistics/counters.

The Integration team is maintaining a Python based tool, that takes advantage of the above MBeans exposed via Jolokia, and the *systemmetrics* project offers a DLUX based UI to display the same information.

Geo-distributed Active/Backup Setup

An OpenDaylight cluster works best when the latency between the nodes is very small, which practically means they should be in the same datacenter. It is however desirable to have the possibility to fail over to a different datacenter, in case all nodes become unreachable. To achieve that, the cluster can be expanded with nodes in a different datacenter, but in a way that doesn't affect latency of the primary nodes. To do that, shards in the backup nodes must be in "non-voting" state.

The API to manipulate voting states on shards is defined as RPCs in the cluster-admin.yang file in the *controller* project, which is well documented. A summary is provided below.

Note: Unless otherwise indicated, the below POST requests are to be sent to any single cluster node.

To create an active/backup setup with a 6 node cluster (3 active and 3 backup nodes in two locations) there is an RPC to set voting states of all shards on a list of nodes to a given state:

```
POST  /restconf/operations/cluster-admin:change-member-voting-states-for-all-shards
```

This RPC needs the list of nodes and the desired voting state as input. For creating the backup nodes, this example input can be used:

```
{
  "input": {
    "member-voting-state": [
      {
        "member-name": "member-4",
        "voting": false
      },
      {
        "member-name": "member-5",
        "voting": false
      },
      {
        "member-name": "member-6",
        "voting": false
```

(continues on next page)

(continued from previous page)

```
        }
      ]
    }
}
```

When an active/backup deployment already exists, with shards on the backup nodes in non-voting state, all that is needed for a fail-over from the active "sub-cluster" to backup "sub-cluster" is to flip the voting state of each shard (on each node, active AND backup). That can be easily achieved with the following RPC call (no parameters needed):

```
POST  /restconf/operations/cluster-admin:flip-member-voting-states-for-all-shards
```

If it's an unplanned outage where the primary voting nodes are down, the "flip" RPC must be sent to a backup non-voting node. In this case there are no shard leaders to carry out the voting changes. However there is a special case whereby if the node that receives the RPC is non-voting and is to be changed to voting and there's no leader, it will apply the voting changes locally and attempt to become the leader. If successful, it persists the voting changes and replicates them to the remaining nodes.

When the primary site is fixed and you want to fail back to it, care must be taken when bringing the site back up. Because it was down when the voting states were flipped on the secondary, its persisted database won't contain those changes. If brought back up in that state, the nodes will think they're still voting. If the nodes have connectivity to the secondary site, they should follow the leader in the secondary site and sync with it. However if this does not happen then the primary site may elect its own leader thereby partitioning the 2 clusters, which can lead to undesirable results. Therefore it is recommended to either clean the databases (i.e., `journal` and `snapshots` directory) on the primary nodes before bringing them back up or restore them from a recent backup of the secondary site (see section *Backing Up and Restoring the Datastore*).

If is also possible to gracefully remove a node from a cluster, with the following RPC:

```
POST  /restconf/operations/cluster-admin:remove-all-shard-replicas
```

and example input:

```
{
  "input": {
    "member-name": "member-1"
  }
}
```

or just one particular shard:

```
POST  /restconf/operations/cluster-admin:remove-shard-replica
```

with example input:

```
{
  "input": {
    "shard-name": "default",
    "member-name": "member-2",
    "data-store-type": "config"
  }
}
```

Now that a (potentially dead/unrecoverable) node was removed, another one can be added at runtime, without changing the configuration files of the healthy nodes (requiring reboot):

```
POST  /restconf/operations/cluster-admin:add-replicas-for-all-shards
```

No input required, but this RPC needs to be sent to the new node, to instruct it to replicate all shards from the cluster.

Note: While the cluster admin API allows adding and removing shards dynamically, the `module-shard.conf` and `modules.conf` files are still used on startup to define the initial configuration of shards. Modifications from the use of the API are not stored to those static files, but to the journal.

Persistence and Backup

Set Persistence Script

This script is used to enable or disable the config datastore persistence. The default state is enabled but there are cases where persistence may not be required or even desired. The user should restart the node to apply the changes.

Note: The script can be used at any time, even before the controller is started for the first time.

Usage:

```
bin/set_persistence.sh <on/off>
```

Example:

```
bin/set_persistence.sh off
```

The above command will disable the config datastore persistence.

Backing Up and Restoring the Datastore

The same cluster-admin API that is used above for managing shard voting states has an RPC allowing backup of the datastore in a single node, taking only the file name as a parameter:

```
POST  /restconf/operations/cluster-admin:backup-datastore
```

RPC input JSON:

```
{
  "input": {
    "file-path": "/tmp/datastore_backup"
  }
}
```

Note: This backup can only be restored if the YANG models of the backed-up data are identical in the backup OpenDaylight instance and restore target instance.

To restore the backup on the target node the file needs to be placed into the `$KARAF_HOME/clustered-datastore-restore` directory, and then the node restarted. If the directory does not exist (which is quite likely if this is a first-time restore) it needs to be created. On startup, ODL checks if the `journal` and `snapshots` directories in `$KARAF_HOME` are empty, and only then tries to read the contents of

the `clustered-datastore-restore` directory, if it exists. So for a successful restore, those two directories should be empty. The backup file name itself does not matter, and the startup process will delete it after a successful restore.

The backup is node independent, so when restoring a 3 node cluster, it is best to restore it on each node for consistency. For example, if restoring on one node only, it can happen that the other two empty nodes form a majority and the cluster comes up with no data.

Running XSQL Console Commands and Queries

XSQL Overview

XSQL is an XML-based query language that describes simple stored procedures which parse XML data, query or update database tables, and compose XML output. XSQL allows you to query tree models like a sequential database. For example, you could run a query that lists all of the ports configured on a particular module and their attributes.

The following sections cover the XSQL installation process, supported XSQL commands, and the way to structure queries.

Installing XSQL

To run commands from the XSQL console, you must first install XSQL on your system:

1. Navigate to the directory in which you unzipped OpenDaylight
2. Start Karaf:

```
./bin/karaf
```

3. Install XSQL:

```
feature:install odl-mdsal-xsql
```

XSQL Console Commands

To enter a command in the XSQL console, structure the command as follows:

```
odl:xsql _<XSQL command>_
```

The following table describes the commands supported in this OpenDaylight release.

Supported XSQL Console Commands

Command	*Description*
r	Repeats the last command you executed.
list vtables	Lists the schema node containers that are currently installed. Whenever an OpenDaylight module is installed, its YANG model is placed in the schema context. At that point, the XSQL receives a notification, confirms that the module's YANG model resides in the schema context and then maps the model to XSQL by setting up the necessary vtables and vfields. This command is useful when you need to determine vtable information for a query.
list vfields *<vtable name>*	Lists the vfields present in a specific vtable. This command is useful when you need to determine vfields information for a query.
jdbc *<ip address>*	When the ODL server is behind a firewall, and the JDBC client cannot connect to the JDBC server, run this command to start the client as a server and establish a connection.
exit	Closes the console.
tocsv	Enables or disables the forwarding of query output as a .csv file.
filename *<filename>*	Specifies the .tocsv file to which the query data is exported. If you do not specify a value for this option when the toccsv option is enabled, the filename for the query data file is generated automatically.

XSQL Queries

You can run a query to extract information that meets the criteria you specify using the information provided by the *list vtables* and *list vfields* _<vtable name>_ commands. Any query you run should be structured as follows:

select _<vfields you want to search for, separated by a comma and a space>_ *from* _<vtables you want to search in, separated by a comma and a space>_ *where* _<criteria>_ '*_<*criteria operator*>_';*

For example, if you want to search the nodes/node ID field in the nodes/node-connector table and find every instance of the Hardware-Address object that contains _BA_ in its text string, enter the following query:

```
select nodes/node.ID from nodes/node-connector where Hardware-Address like '%BA%';
```

The following criteria operators are supported:

Supported XSQL Query Criteria Operators

Criteria Operators	Description
=	Lists results that equal the value you specify.
!=	Lists results that do not equal the value you specify.
like	Lists results that contain the substring you specify. For example, if you specify *like %BC%*, every string that contains that particular substring is displayed.
<	Lists results that are less than the value you specify.
>	Lists results that are more than the value you specify.
and	Lists results that match both values you specify.
or	Lists results that match either of the two values you specify.
>=	Lists results that are more than or equal to the value you specify.
<=	Lists results that are less than or equal to the value you specify.
is null	Lists results for which no value is assigned.
not null	Lists results for which any value is assigned.
skip	Use this operator to list matching results from a child node, even if its parent node does not meet the specified criteria. See the following example for more information.

Example: Skip Criteria Operator

If you are looking at the following structure and want to determine all of the ports that belong to a YY type module:

- Network Element 1
  - Module 1, Type XX
    - Module 1.1, Type YY
      - Port 1
      - Port 2
  - Module 2, Type YY
    - Port 1
    - Port 2

If you specify *Module.Type='YY'* in your query criteria, the ports associated with module 1.1 will not be returned since its parent module is type XX. Instead, enter *Module.Type='YY' or skip Module!='YY'*. This tells XSQL to disregard any parent module data that does not meet the type YY criteria and collect results for any matching child modules. In this example, you are instructing the query to skip module 1 and collect the relevant data from module 1.1.

OpenDaylight Version

Overview

This feature allows NETCONF/RESTCONF users to determine the version of OpenDaylight they are communicating with.

Install the Version Feature

Follow these steps to install the version feature:

1. Navigate to the directory in which you installed OpenDaylight

2. Start Karaf:

```
./bin/karaf
```

3. Install Version feature:

```
feature:install odl-distribution-version
```

Note: For RESTCONF access, it is recommended to install odl-restconf and odl-netconf-connector-ssh.

Version Feature Usage

Example of RESTCONF request using curl from bash:

```
$ curl -u 'admin:admin' localhost:8181/restconf/config/network-topology:network-
↪topology/topology/topology-netconf/node/controller-config/yang-ext:mount/
↪config:modules/module/odl-distribution-version:odl-version/odl-distribution-version
```

Example response (formatted):

```
{
 "module": [
  {
   "type": "odl-distribution-version:odl-version",
   "name": "odl-distribution-version",
   "odl-distribution-version:version": "0.5.0-SNAPSHOT"
  }
 ]
}
```

1.1.13 Security Considerations

This document discusses the various security issues that might affect OpenDaylight. The document also lists specific recommendations to mitigate security risks.

This document also contains information about the corrective steps you can take if you discover a security issue with OpenDaylight, and if necessary, contact the Security Response Team, which is tasked with identifying and resolving security threats.

Overview of OpenDaylight Security

There are many different kinds of security vulnerabilities that could affect an OpenDaylight deployment, but this guide focuses on those where (a) the servers, virtual machines or other devices running OpenDaylight have been properly physically (or virtually in the case of VMs) secured against untrusted individuals and (b) individuals who have access, either via remote logins or physically, will not attempt to attack or subvert the deployment intentionally or otherwise.

While those attack vectors are real, they are out of the scope of this document.

What remains in scope is attacks launched from a server, virtual machine, or device other than the one running OpenDaylight where the attack does not have valid credentials to access the OpenDaylight deployment.

The rest of this document gives specific recommendations for deploying OpenDaylight in a secure manner, but first we highlight some high-level security advantages of OpenDaylight.

- Separating the control and management planes from the data plane (both logically and, in many cases, physically) allows possible security threats to be forced into a smaller attack surface.
- Having centralized information and network control gives network administrators more visibility and control over the entire network, enabling them to make better decisions faster. At the same time, centralization of network control can be an advantage only if access to that control is secure.

Note: While both previous advantages improve security, they also make an OpenDaylight deployment an attractive target for attack making understanding these security considerations even more important.

- The ability to more rapidly evolve southbound protocols and how they are used provides more and faster mechanisms to enact appropriate security mitigations and remediations.
- OpenDaylight is built from OSGi bundles and the Karaf Java container. Both Karaf and OSGi provide some level of isolation with explicit code boundaries, package imports, package exports, and other security-related features.
- OpenDaylight has a history of rapidly addressing known vulnerabilities and a well-defined process for reporting and dealing with them.

OpenDaylight Security Resources

- If you have any security issues, you can send a mail to *security@lists.opendaylight.org*.
- For the list of current OpenDaylight security issues that are either being fixed or resolved, refer to https://wiki.opendaylight.org/view/Security:Advisories.
- To learn more about the OpenDaylight security issues policies and procedure, refer to https://wiki.opendaylight.org/view/Security:Main

Deployment Recommendations

We recommend that you follow the deployment guidelines in setting up OpenDaylight to minimize security threats.

- The default credentials should be changed before deploying OpenDaylight.
- OpenDaylight should be deployed in a private network that cannot be accessed from the internet.
- Separate the data network (that connects devices using the network) from the management network (that connects the network devices to OpenDaylight).

Note: Deploying OpenDaylight on a separate, private management network does not eliminate threats, but only mitigates them. By construction, some messages must flow from the data network to the management network, e.g., OpenFlow *packet_in* messages, and these create an attack surface even if it is a small one.

- Implement an authentication policy for devices that connect to both the data and management network. These are the devices which bridge, likely untrusted, traffic from the data network to the management network.

Securing OSGi bundles

OSGi is a Java-specific framework that improves the way that Java classes interact within a single JVM. It provides an enhanced version of the *java.lang.SecurityManager* (ConditionalPermissionAdmin) in terms of security.

Java provides a security framework that allows a security policy to grant permissions, such as reading a file or opening a network connection, to specific code. The code maybe classes from the jarfile loaded from a specific URL, or a class signed by a specific key. OSGi builds on the standard Java security model to add the following features:

- A set of OSGi-specific permission types, such as one that grants the right to register an OSGi service or get an OSGi service from the service registry.
- The ability to dynamically modify permissions at runtime. This includes the ability to specify permissions by using code rather than a text configuration file.
- A flexible predicate-based approach to determining which rules are applicable to which *ProtectionDomain*. This approach is much more powerful than the standard Java security policy which can only grant rights based on a jarfile URL or class signature. A few standard predicates are provided, including selecting rules based upon bundle symbolic-name.
- Support for bundle *local permissions* policies with optional further constraints such as *DENY* operations. Most of this functionality is accessed by using the *OSGi ConditionalPermissionAdmin* service which is part of the OSGi core and can be obtained from the OSGi service registry. The *ConditionalPermissionAdmin* API replaces the earlier *PermissionAdmin* API.

For more information, refer to http://www.osgi.org/Main/HomePage.

Securing the Karaf container

Apache Karaf is a OSGi-based runtime platform which provides a lightweight container for OpenDaylight and applications. Apache Karaf uses either Apache Felix Framework or Eclipse Equinox OSGi frameworks, and provide additional features on top of the framework.

Apache Karaf provides a security framework based on Java Authentication and Authorization Service (JAAS) in compliance with OSGi recommendations, while providing RBAC (Role-Based Access Control) mechanism for the console and Java Management Extensions (JMX).

The Apache Karaf security framework is used internally to control the access to the following components:

- OSGi services
- console commands
- JMX layer
- WebConsole

The remote management capabilities are present in Apache Karaf by default, however they can be disabled by using various configuration alterations. These configuration options may be applied to the OpenDaylight Karaf distribution.

Note: Refer to the following list of publications for more information on implementing security for the Karaf container.

- For role-based JMX administration, refer to http://karaf.apache.org/manual/latest/users-guide/monitoring.html.
- For remote SSH access configuration, refer to http://karaf.apache.org/manual/latest/users-guide/remote.html.
- For WebConsole access, refer to http://karaf.apache.org/manual/latest/users-guide/webconsole.html.
- For Karaf security features, refer to http://karaf.apache.org/manual/latest/developers-guide/security-framework.html.

Disabling the remote shutdown port

You can lock down your deployment post installation. Set `karaf.shutdown.port=-1` in `etc/custom.properties` or `etc/config.properties` to disable the remote shutdown port.

Securing Southbound Plugins

Many individual southbound plugins provide mechanisms to secure their communication with network devices. For example, the OpenFlow plugin supports TLS connections with bi-directional authentication and the NETCONF plugin supports connecting over SSH. Meanwhile, the Unified Secure Channel plugin provides a way to form secure, remote connections for supported devices.

When deploying OpenDaylight, you should carefully investigate the secure mechanisms to connect to devices using the relevant plugins.

Securing OpenDaylight using AAA

AAA stands for Authentication, Authorization, and Accounting. All three of can help improve the security posture of and OpenDaylight deployment. In this release, only authentication is fully supported, while authorization is an experimental feature and accounting remains a work in progress.

The vast majority of OpenDaylight's northbound APIs (and all RESTCONF APIs) are protected by AAA by default when installing the +odl-restconf+ feature. In the cases that APIs are *not* protected by AAA, this will be noted in the per-project release notes.

By default, OpenDaylight has only one user account with the username and password *admin*. This should be changed before deploying OpenDaylight.

Security Considerations for Clustering

While OpenDaylight clustering provides many benefits including high availability, scale-out performance, and data durability, it also opens a new attack surface in the form of the messages exchanged between the various instances of OpenDaylight in the cluster. In the current OpenDaylight release, these messages are neither encrypted nor authenticated meaning that anyone with access to the management network where OpenDaylight exchanges these clustering messages can forge and/or read the messages. This means that if clustering is enabled, it is even more important that the management network be kept secure from any untrusted entities.

1.2 OpenDaylight User Guide

1.2.1 Overview

This first part of the user guide covers the basic user operations of the OpenDaylight Release using the generic base functionality.

OpenDaylight Controller Overview

The OpenDaylight controller is JVM software and can be run from any operating system and hardware as long as it supports Java. The controller is an implementation of the Software Defined Network (SDN) concept and makes use of the following tools:

- **Maven**: OpenDaylight uses Maven for easier build automation. Maven uses pom.xml (Project Object Model) to script the dependencies between bundle and also to describe what bundles to load and start.
- **OSGi**: This framework is the back-end of OpenDaylight as it allows dynamically loading bundles and packages JAR files, and binding bundles together for exchanging information.
- **JAVA interfaces**: Java interfaces are used for event listening, specifications, and forming patterns. This is the main way in which specific bundles implement call-back functions for events and also to indicate awareness of specific state.
- **REST APIs**: These are northbound APIs such as topology manager, host tracker, flow programmer, static routing, and so on.

The controller exposes open northbound APIs which are used by applications. The OSGi framework and bidirectional REST are supported for the northbound APIs. The OSGi framework is used for applications that run in the same address space as the controller while the REST (web-based) API is used for applications that do not run in the same address space (or even the same system) as the controller. The business logic and algorithms reside in the applications. These applications use the controller to gather network intelligence, run its algorithm to do analytics, and then orchestrate the new rules throughout the network. On the southbound, multiple protocols are supported as plugins, e.g. OpenFlow 1.0, OpenFlow 1.3, BGP-LS, and so on. The OpenDaylight controller starts with an OpenFlow 1.0 southbound plugin. Other OpenDaylight contributors begin adding to the controller code. These modules are linked dynamically into a **Service Abstraction Layer** (SAL).

The SAL exposes services to which the modules north of it are written. The SAL figures out how to fulfill the requested service irrespective of the underlying protocol used between the controller and the network devices. This provides investment protection to the applications as OpenFlow and other protocols evolve over time. For the controller to control devices in its domain, it needs to know about the devices, their capabilities, reachability, and so on. This information is stored and managed by the **Topology Manager**. The other components like ARP handler, Host Tracker, Device Manager, and Switch Manager help in generating the topology database for the Topology Manager.

For a more detailed overview of the OpenDaylight controller, see the *OpenDaylight Developer Guide*.

Using the OpenDaylight User Interface (DLUX)

This section introduces you to the OpenDaylight User Experience (DLUX) application.

Getting Started with DLUX

DLUX provides a number of different Karaf features, which you can enable and disable separately. In Beryllum they are: . odl-dlux-core . odl-dlux-node . odl-dlux-yangui . odl-dlux-yangvisualizer

Logging In

To log in to DLUX, after installing the application:

1. Open a browser and enter the login URL http://<your-karaf-ip>:8181/index.html in your browser (Chrome is recommended).
2. Login to the application with your username and password credentials.

Note: OpenDaylight's default credentials are *admin* for both the username and password.

Working with DLUX

After you login to DLUX, if you enable only odl-dlux-core feature, you will see only topology application available in the left pane.

Note: To make sure topology displays all the details, enable the odl-l2switch-switch feature in Karaf.

DLUX has other applications such as node, yang UI and those apps won't show up, until you enable their features odl-dlux-node and odl-dlux-yangui respectively in the Karaf distribution.

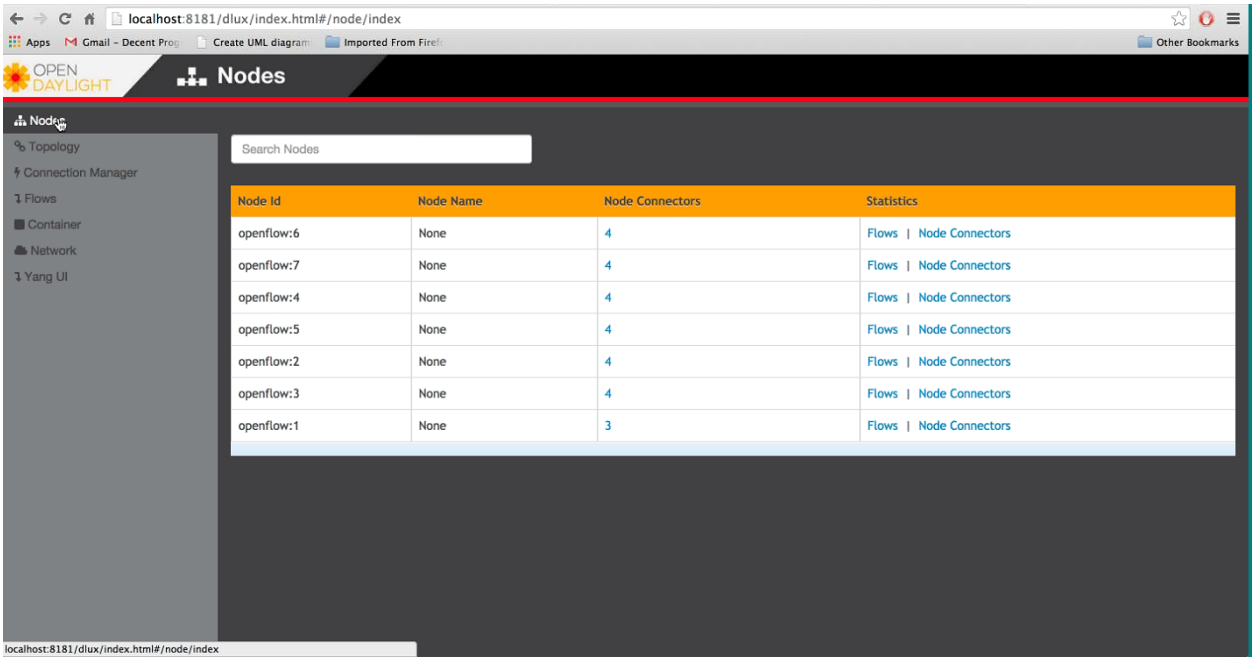

Fig. 10: DLUX Modules

Note: If you install your application in dlux, they will also show up on the left hand navigation after browser page refresh.

Viewing Network Statistics

The **Nodes** module on the left pane enables you to view the network statistics and port information for the switches in the network.

To use the **Nodes** module:

1. Select **Nodes** on the left pane. The right pane displays atable that lists all the nodes, node connectors and the statistics.
2. Enter a node ID in the **Search Nodes** tab to search by node connectors.
3. Click on the **Node Connector** number to view details such as port ID, port name, number of ports per switch, MAC Address, and so on.

4. Click **Flows** in the Statistics column to view Flow Table Statistics for the particular node like table ID, packet match, active flows and so on.
5. Click **Node Connectors** to view Node Connector Statistics for the particular node ID.

Viewing Network Topology

The Topology tab displays a graphical representation of network topology created.

Note: DLUX does not allow for editing or adding topology information. The topology is generated and edited in other modules, e.g., the OpenFlow plugin. OpenDaylight stores this information in the MD-SAL datastore where DLUX can read and display it.

To view network topology:

1. Select **Topology** on the left pane. You will view the graphical representation on the right pane. In the diagram blue boxes represent the switches, the black represents the hosts available, and lines represents how the switches and hosts are connected.
2. Hover your mouse on hosts, links, or switches to view source and destination ports.
3. Zoom in and zoom out using mouse scroll to verify topology for larger topologies.

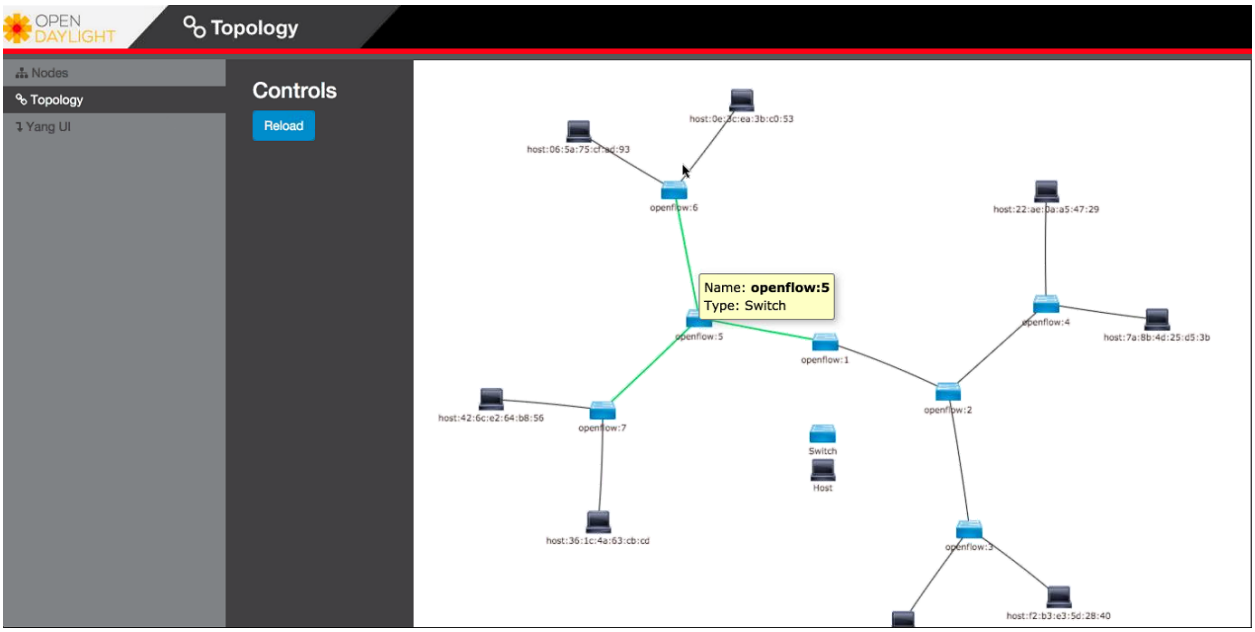

Fig. 11: Topology Module

Interacting with the YANG-based MD-SAL datastore

The **Yang UI** module enables you to interact with the YANG-based MD-SAL datastore. For more information about YANG and how it interacts with the MD-SAL datastore, see the *Controller* and *YANG Tools* section of the *OpenDaylight Developer Guide*.

To use Yang UI:

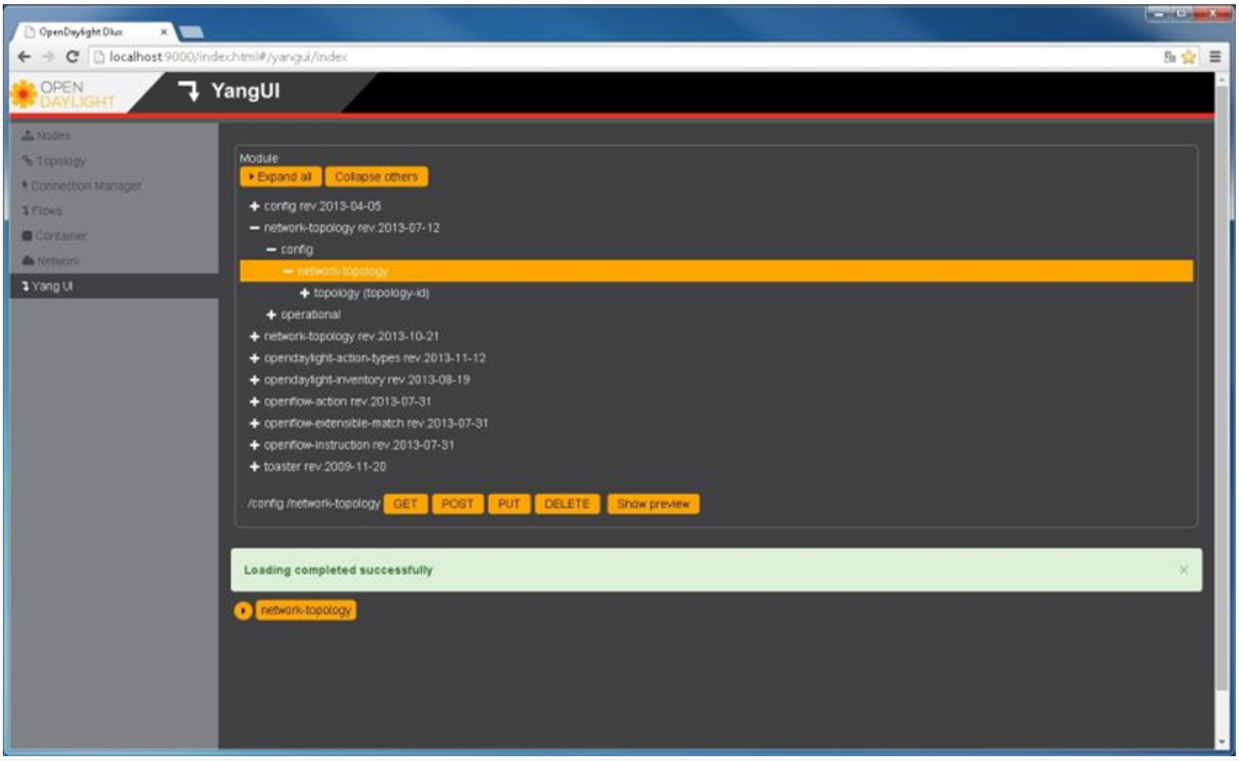

Fig. 12: Yang UI

1. Select **Yang UI** on the left pane. The right pane is divided in two parts.
2. The top part displays a tree of APIs, subAPIs, and buttons to call possible functions (GET, POST, PUT, and DELETE).

   Note: every subAPI can call every function. For example, subAPIs in the *operational* store have GET functionality only.

   Inputs can be filled from OpenDaylight when existing data from OpenDaylight is displayed or can be filled by user on the page and sent to OpenDaylight.

   Buttons under the API tree are variable. It depends on subAPI specifications. Common buttons are:

   - GET to get data from OpenDaylight,
   - PUT and POST for sending data to OpenDaylight for saving
   - DELETE for sending data to OpenDaylight for deleting.

     You must specify the xpath for all these operations. This path is displayed in the same row before buttons and it may include text inputs for specific path element identifiers.
3. The bottom part of the right pane displays inputs according to the chosen subAPI.
   - Lists are handled as a special case. For example, a device can store multiple flows. In this case "flow" is name of the list and every list element is identified by a unique key value. Elements of a list can, in turn, contain other lists.
   - In Yang UI, each list element is rendered with the name of the list it belongs to, its key, its value, and a button for removing it from the list.

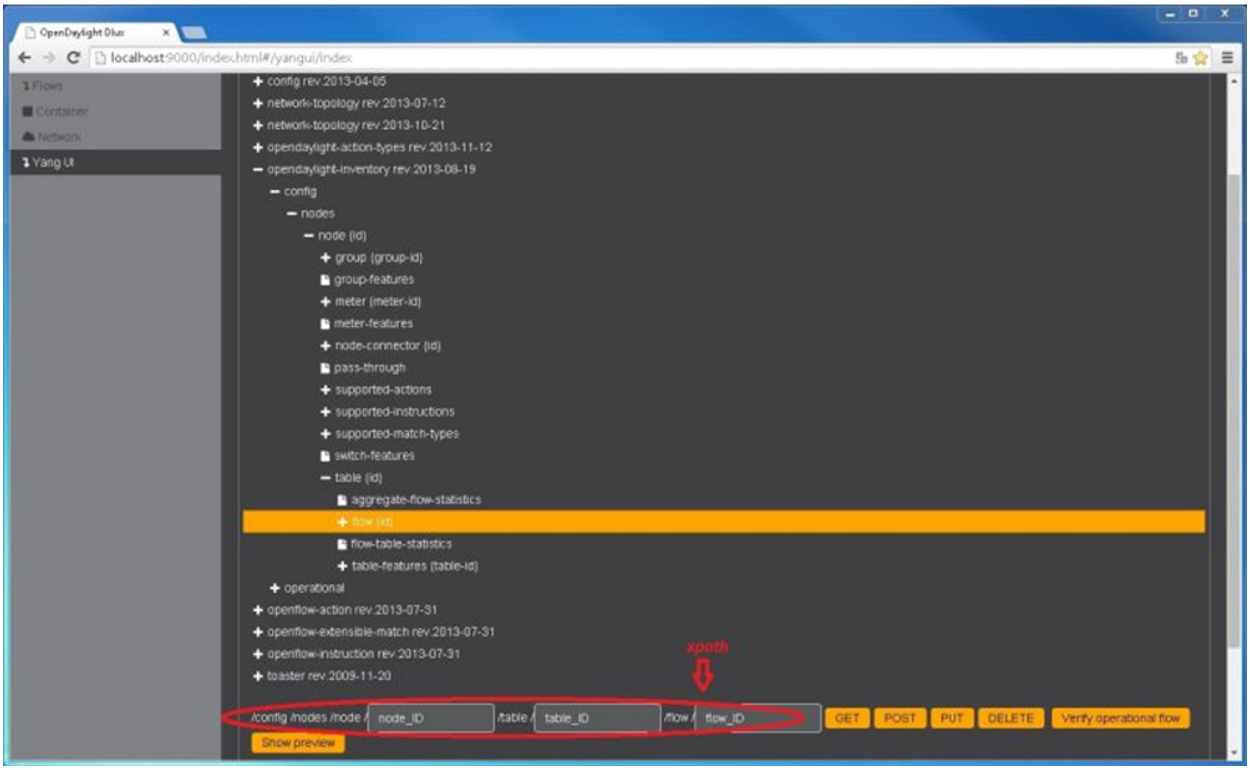

Fig. 13: Yang API Specification

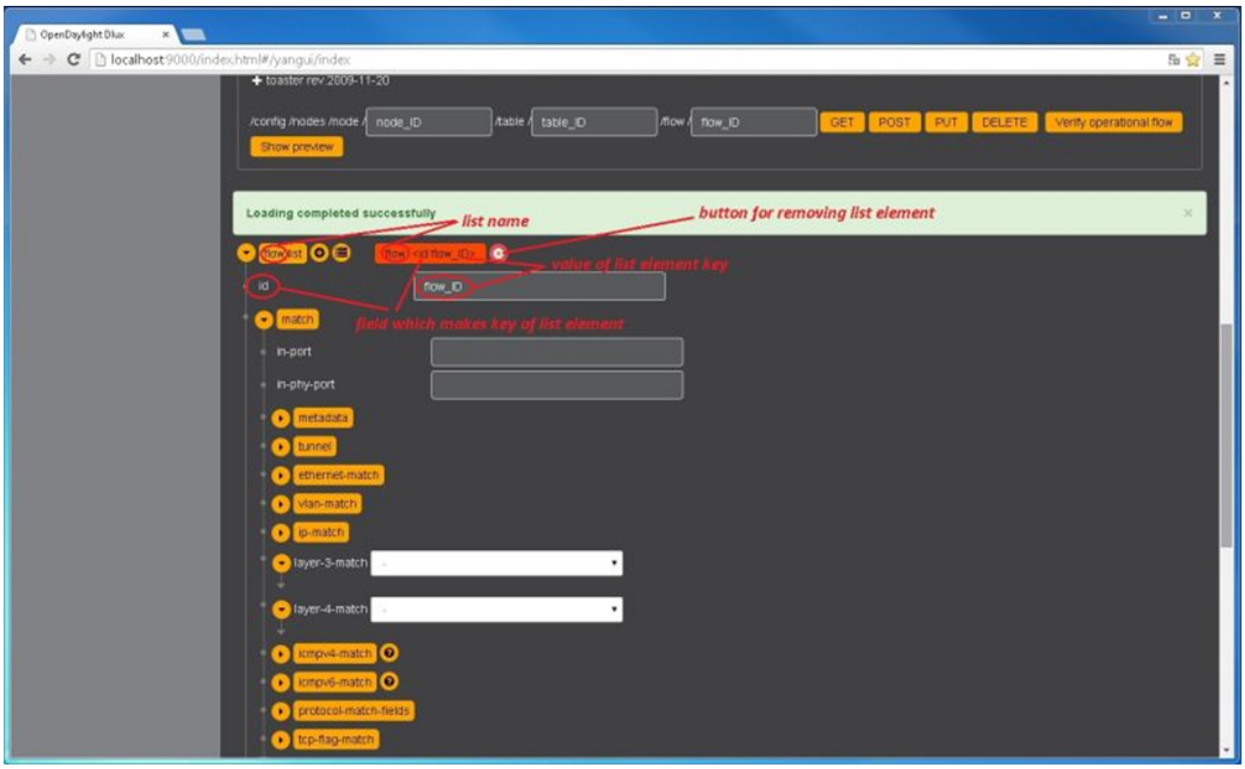

Fig. 14: Yang UI API Specification

4. After filling in the relevant inputs, click the **Show Preview** button under the API tree to display request that will be sent to OpenDaylight. A pane is displayed on the right side with text of request when some input is filled.

Displaying Topology on the Yang UI

To display topology:

1. Select subAPI network-topology <topology revision number> == > operational == > network-topology.
2. Get data from OpenDaylight by clicking on the "GET" button.
3. Click **Display Topology**.

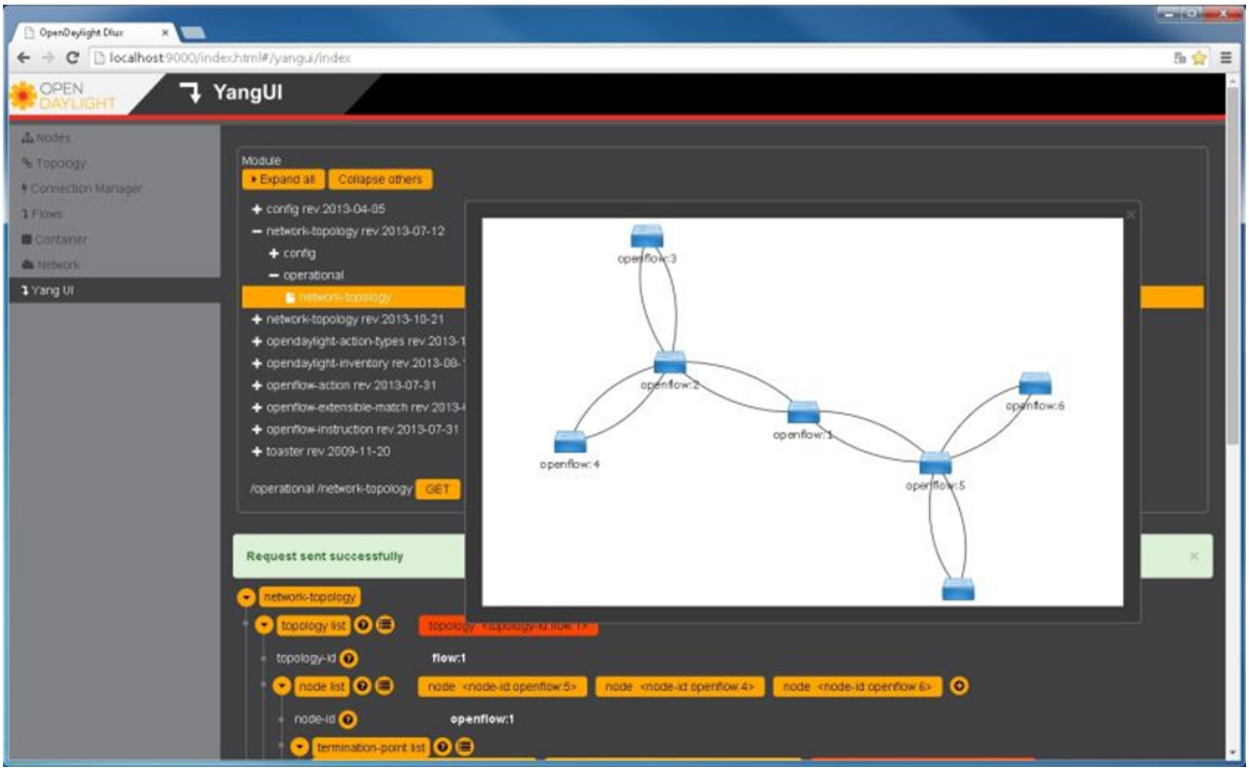

Fig. 15: DLUX Yang Topology

Configuring List Elements on the Yang UI

Lists in Yang UI are displayed as trees. To expand or collapse a list, click the arrow before name of the list. To configure list elements in Yang UI:

1. To add a new list element with empty inputs use the plus icon-button **+** that is provided after list name.
2. To remove several list elements, use the **X** button that is provided after every list element.
3. In the YANG-based data store all elements of a list must have a unique key. If you try to assign two or more elements the same key, a warning icon **!** is displayed near their name buttons.
4. When the list contains at least one list element, after the **+** icon, there are buttons to select each individual list element. You can choose one of them by clicking on it. In addition, to the right of the list name, there is a button which will display a vertically scrollable pane with all the list elements.

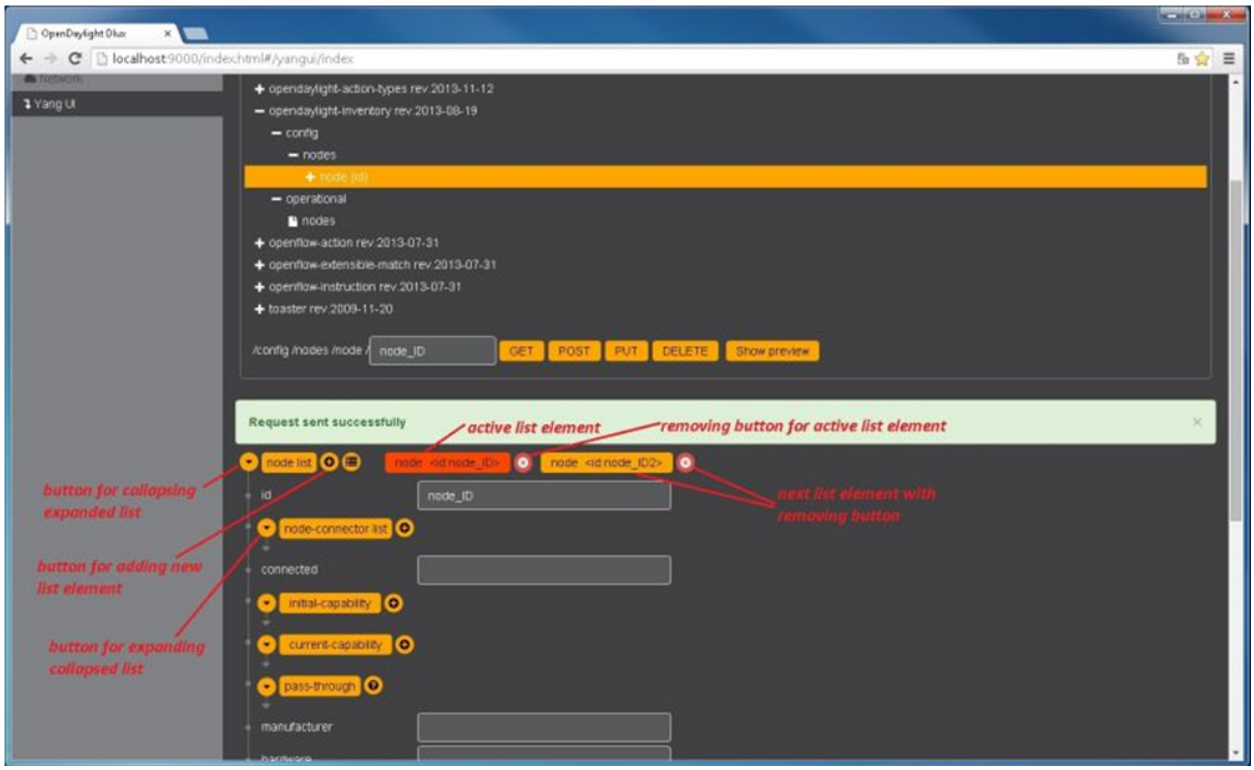

Fig. 16: DLUX List Elements

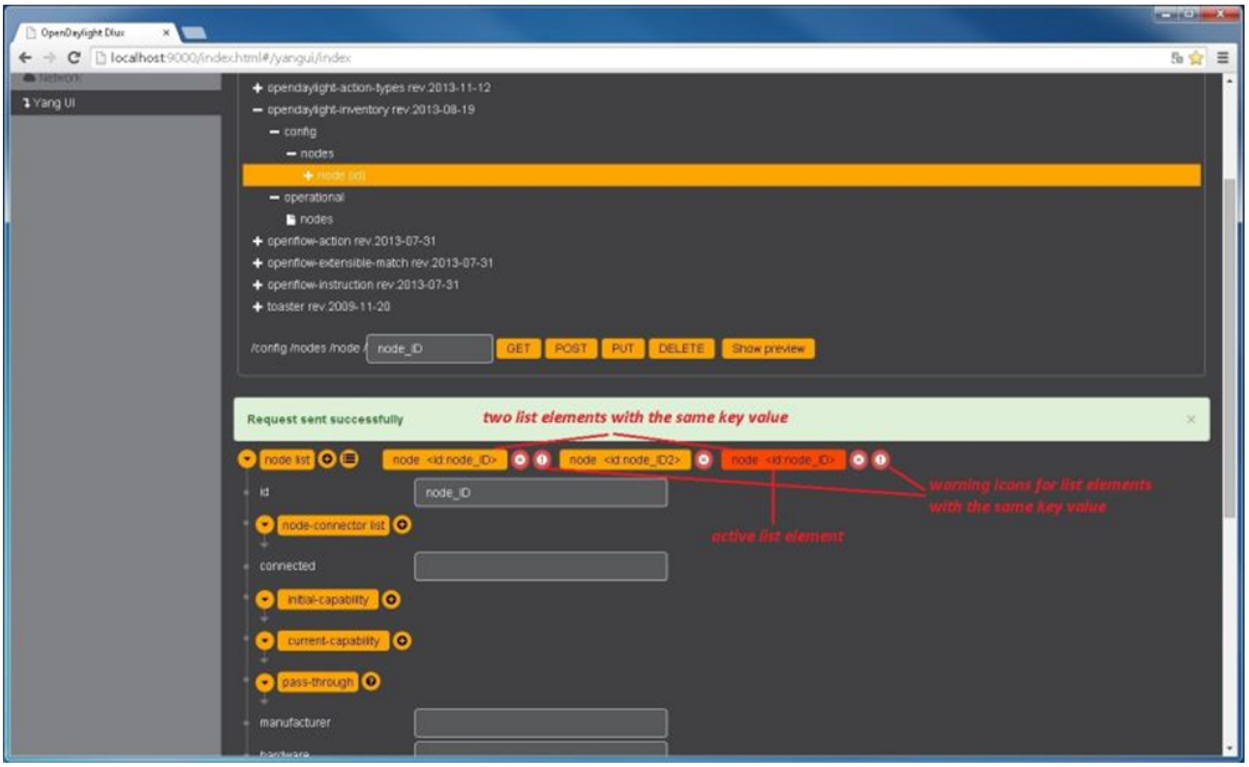

Fig. 17: DLUX List Warnings

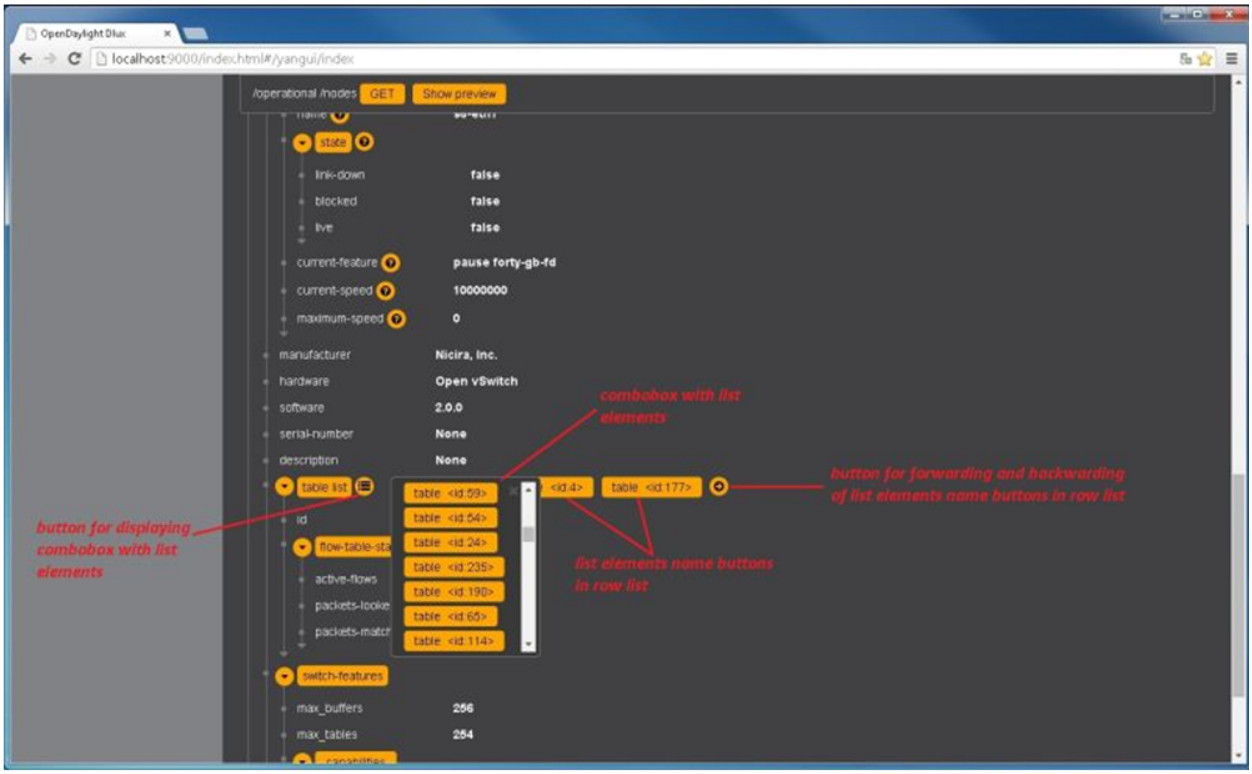

Fig. 18: DLUX List Button1

Running XSQL Console Commands and Queries

XSQL Overview

XSQL is an XML-based query language that describes simple stored procedures which parse XML data, query or update database tables, and compose XML output. XSQL allows you to query tree models like a sequential database. For example, you could run a query that lists all of the ports configured on a particular module and their attributes.

The following sections cover the XSQL installation process, supported XSQL commands, and the way to structure queries.

Installing XSQL

To run commands from the XSQL console, you must first install XSQL on your system:

1. Navigate to the directory in which you unzipped OpenDaylight
2. Start Karaf:

```
./bin/karaf
```

3. Install XSQL:

```
feature:install odl-mdsal-xsql
```

XSQL Console Commands

To enter a command in the XSQL console, structure the command as follows: **odl:xsql** <*XSQL command*>

The following table describes the commands supported in this OpenDaylight release.

Command	Description
r	Repeats the last command you executed.
list vtables	Lists the schema node containers that are currently installed. Whenever an OpenDaylight module is installed, its YANG model is placed in the schema context. At that point, the XSQL receives a notification, confirms that the module's YANG model resides in the schema context and then maps the model to XSQL by setting up the necessary vtables and vfields. This command is useful when you need to determine vtable information for a query.
list vfields <*vtable name*>	Lists the vfields present in a specific vtable. This command is useful when you need to determine vfields information for a query.
jdbc <*ip address*>	When the ODL server is behind a firewall, and the JDBC client cannot connect to the JDBC server, run this command to start the client as a server and establish a connection.
exit	Closes the console.
tocsv	Enables or disables the forwarding of query output as a .csv file.
filename <*filename*>	Specifies the .tocsv file to which the query data is exported. If you do not specify a value for this option when the toccsv option is enabled, the filename for the query data file is generated automatically.

Table: Supported XSQL Console Commands

XSQL Queries

You can run a query to extract information that meets the criteria you specify using the information provided by the **list vtables** and **list vfields** <*vtable name*> commands. Any query you run should be structured as follows:

select <*vfields you want to search for, separated by a comma and a space*> **from** <*vtables you want to search in, separated by a comma and a space*> **where** <*criteria*> ***<criteria operator>**;*

For example, if you want to search the nodes/node ID field in the nodes/node-connector table and find every instance of the Hardware-Address object that contains *BA* in its text string, enter the following query:

```
select nodes/node.ID from nodes/node-connector where Hardware-Address like '%BA%';
```

The following criteria operators are supported:

Criteria Operators	Description
=	Lists results that equal the value you specify.
!=	Lists results that do not equal the value you specify.
like	Lists results that contain the substring you specify. For example, if you specify **like %BC%**, every string that contains that particular substring is displayed.
<	Lists results that are less than the value you specify.
>	Lists results that are more than the value you specify.
and	Lists results that match both values you specify.
or	Lists results that match either of the two values you specify.
>=	Lists results that are more than or equal to the value you specify.
	Lists results that are less than or equal to the value you specify.
is null	Lists results for which no value is assigned.
not null	Lists results for which any value is assigned.
skip	Use this operator to list matching results from a child node, even if its parent node does not meet the specified criteria. See the following example for more information.

Table: Supported XSQL Query Criteria Operators

Example: Skip Criteria Operator

If you are looking at the following structure and want to determine all of the ports that belong to a YY type module:

- Network Element 1
  - Module 1, Type XX
    - Module 1.1, Type YY
      - Port 1
      - Port 2
  - Module 2, Type YY
    - Port 1
    - Port 2

If you specify **Module.Type=*YY*** in your query criteria, the ports associated with module 1.1 will not be returned since its parent module is type XX. Instead, enter **Module.Type=*YY* or skip Module!=*YY***. This tells XSQL to disregard any parent module data that does not meet the type YY criteria and collect results for any matching child modules. In this example, you are instructing the query to skip module 1 and collect the relevant data from module 1.1.

1.2.2 Project-specific User Guides

ALTO User Guide

Overview

The ALTO project is aimed to provide support for **Application Layer Traffic Optimization** services defined in RFC 7285 in OpenDaylight.

This user guide will introduce the three basic services (namely `simple-ird`, `manual-maps` and `host-tracker`) which are implemented since the Beryllium release, and give instructions on how to configure them to provide corresponding ALTO services.

A new feature named `simple-pce` (**Simple Path Computation Engine**) is added into Boron release as an ALTO extension service.

How to Identify ALTO Resources

Each ALTO resource can be uniquely identified by a tuple . For each resource, a *version-tag* is used to support historical look-ups.

The formats of *resource-id* and *version-tag* are defined in section 10.2 and section 10.3 respectively. The *context-id* is not part of the protocol and we choose the same format as a *universal unique identifier* (UUID) which is defined in RFC 4122.

A context is like a namespace for ALTO resources, which eliminates *resource-id* collisions. For simplicity, we also provide a default context with the id **00000000000-0000-0000-0000-00000000**.

How to Use Simple IRD

The simple IRD feature provides a simple *information resource directory* (IRD) service defined in RFC 7285.

Install the Feature

To enable simple IRD, run the following command in the karaf CLI:

```
karaf > feature:install odl-alto-simpleird
```

After the feature is successfully installed, a special context will be created for all simple IRD resources. The id for this context can be seen by executing the following command in a terminal:

```
curl -X GET -u admin:admin http://localhost:8181/restconf/operational/alto-simple-
↪ird:information/
```

Create a new IRD

To create a new IRD resource, two fields MUST be provided:

- Field **instance-id**: the *resource-id* of the IRD resource;
- Field **entry-context**: the context-id for non-IRD entries managed by this IRD resource.

Using the following script, one can create an empty IRD resource:

```
#!/bin/bash
# filename: ird-create
INSTANCE_ID=$1
if [ $2 ]; then
    CONTEXT_ID=$2
else
    CONTEXT_ID="00000000-0000-0000-0000-000000000000"
fi
```

(continues on next page)

(continued from previous page)

```
URL="`http://localhost:8181/restconf/config/alto-simple-ird:ird-instance-
↪configuration/"$INSTANCE_ID"/[`http://localhost:8181/restconf/config/alto-simple-
↪ird:ird-instance-configuration/"$INSTANCE_ID"/`]`"
DATA=$(cat template \
  | sed 's/\$1/'$CONTEXT_ID'/g' \
  | sed 's/\$2/'$INSTANCE_ID'/g')
curl -4 -D - -X PUT -u admin:admin \
  -H "Content-Type: application/json" -d "$(echo $DATA)"\
  $URL
```

For example, the following command will create a new IRD named *ird* which can accept entries with the default context-id:

```
$ ./ird-create ird 000000000000-0000-0000-0000-00000000
```

And below is the what the template file looks like:

```
{
    "ird-instance-configuration": {
        "entry-context": "/alto-resourcepool:context[alto-resourcepool:context-id='$1
↪']",
        "instance-id": "$2"
    }
}
```

Remove an IRD

To remove an existing IRD (and all the entries in it), one can use the following command in a terminal:

```
curl -X DELETE -u admin:admin http://localhost:8181/restconf/config/alto-simple-
↪ird:ird-instance-configuration/$INSTANCE_ID
```

Add a new entry

There are several ways to add entries to an IRD and in this section we introduce only the simplest method. Using the following script, one can add a new entry to the target IRD.

For each new entry, four parameters MUST be provided:

- Parameter *ird-id*: the *resource-id* of the target IRD;
- Parameter *entry-id*: the *resource-id* of the ALTO service to be added;
- Parameter *context-id*: the *context-id* of the ALTO service to be added, which MUST be identical to the target IRD's *entry-context*;
- Parameter *location*: either a URI or a relative path to the ALTO service.

The following script can be used to add one entry to the target IRD, where the relative path is used:

```
#!/bin/bash
# filename: ird-add-entry
IRD_ID=$1
ENTRY_ID=$2
CONTEXT_ID=$3
```

(continues on next page)

(continued from previous page)

```
BASE_URL=$4
URL="`http://localhost:8181/restconf/config/alto-simple-ird:ird-instance-
↪configuration/"$IRD_ID"/ird-configuration-entry/"$ENTRY_ID"/"
DATA=$(cat template \
  | sed 's/\$1/'$ENTRY_ID'/g' \
  | sed 's/\$2/'$CONTEXT_ID'/g' \
  | sed 's/\$3/'$BASE_URL'/g' )
curl -4 -D - -X PUT -u admin:admin \
  -H "Content-Type: application/json" -d "$(echo $DATA)" \
  $URL
```

For example, the following command will add a new resource named *networkmap*, whose context-id is the default context-id and the base URL is /alto/networkmap, to the IRD named *ird*:

```
$ ./ird-add-entry ird networkmap 000000000000-0000-0000-0000-00000000 /alto/networkmap
```

And below is the template file:

```
{
    "ird-configuration-entry": {
        "entry-id": "$1",
        "instance": "/alto-resourcepool:context[alto-resourcepool:context-id='$2']/
↪alto-resourcepool:resource[alto-resourcepool:resource-id='$1']",
        "path": "$3/$1"
    }
}
```

Remove an entry

To remove an entry from an IRD, one can use the following one-line command:

```
curl -X DELETE -u admin:admin http://localhost:8181/restconf/config/alto-simple-
↪ird:ird-instance-configuration/$IRD_ID/ird-configuration-entry/$ENTRY_ID/
```

How to Use Host-tracker-based ECS

As a real instance of ALTO services, ***alto-hosttracker*** reads data from ***l2switch*** and generates a network map with resource id ***hosttracker-network-map*** and a cost map with resource id ***hostracker-cost-map***. It can only work with OpenFlow-enabled networks.

After installing the ***odl-alto-hosttracker*** feature, the corresponding network map and cost map will be inserted into the data store.

Managing Resource with `alto-resourcepool`

After installing `odl-alto-release` feature in Karaf, `alto-resourcepool` feature will be installed automatically. And you can manage all resources in ALTO via RESTCONF APIs provided by `alto-resourcepool`.

With the example bash script below you can get any resource infomation in a given context.

```
#!/bin/bash
RESOURCE_ID=$1
if [ $2 ] ; then
    CONTEXT_ID=$2
else
    CONTEXT_ID="00000000-0000-0000-0000-000000000000"
fi
URL="http://localhost:8181/restconf/operational/alto-resourcepool:context/"$CONTEXT_ID
↪"/alto-resourcepool:resource/"$RESOURCE_ID
curl -X GET -u admin:admin $URL | python -m json.tool | sed -n '/default-tag/p' | sed
↪'s/.*:.*\"\(.*\)\".*/\1/g'
```

Manual Configuration

Using RESTCONF API

After installing `odl-alto-release` feature in Karaf, it is possible to manage network-maps and cost-maps using RESTCONF. Take a look at all the operations provided by `resource-config` at the API service page which can be found at `http://localhost:8181/apidoc/explorer/index.html`.

The easiest method to operate network-maps and cost-maps is to modify data broker via RESTCONF API directly.

Using RPC

The `resource-config` package also provides a query RPC to config the resources. You can CREATE, UPDATE and DELETE **network-maps** and **cost-maps** via query RPC.

Simple Path Computation Engine

The `simple-pce` module provides a simple path computation engine for ALTO and other projects. It supports basic CRUD (create, read, update, delete) operations to manage L2 and L3 routing with/without rate limitation. This module is an independent feature, so you can follow the instruction below to install it independently.

```
karaf > feature:install odl-alto-extenstion
```

Note: The rate limitation meter requires OpenFlow 1.3 support.

Basic Usage with RESTCONF API

You can use the simple path computation engine with RESTCONF API, which is defined in the YANG model here.

Use Case

Server Selection

One of the key use case for ALTO is server selection. For example, a client (with IP address 10.0.0.1) sends a data transferring request to Data Transferring Service (DTS). And there are three data replica servers (with IP address

10.60.0.1, 10.60.0.2 and 10.60.0.3) which can response the request. In this case, DTS can send a query request to ALTO server to make server selection decision.

Following is an example ALTO query:

```
POST /alto/endpointcost HTTP/1.1
Host: localhost:8080
Content-Type: application/alto-endpointcostparams+json
Accept: application/alto-endpointcost+json,application/alto-error+json
{
  "cost-type": {
    "cost-mode": "ordinal",
    "cost-metric": "hopcount"
  },
  "endpoints": {
    "srcs": [ "ipv4:10.0.0.1" ],
    "dsts": [
      "ipv4:10.60.0.1",
      "ipv4:10.60.0.2",
      "ipv4:10.60.0.3"
  ]
  }
}
```

Authentication, Authorization and Accounting (AAA) Services

The Boron AAA services are based on the Apache Shiro Java Security Framework. The main configuration file for AAA is located at "etc/shiro.ini" relative to the ODL karaf home directory.

Terms And Definitions

Token A claim of access to a group of resources on the controller

Domain A group of resources, direct or indirect, physical, logical, or virtual, for the purpose of access control. ODL recommends using the default "sdn" domain in the Boron release.

User A person who either owns or has access to a resource or group of resources on the controller

Role Opaque representation of a set of permissions, which is merely a unique string as admin or guest

Credential Proof of identity such as username and password, OTP, biometrics, or others

Client A service or application that requires access to the controller

Claim A data set of validated assertions regarding a user, e.g. the role, domain, name, etc.

How to enable AAA

AAA is enabled through installing the odl-aaa-shiro feature. odl-aaa-shiro is automatically installed as part of the odl-restconf offering.

How to disable AAA

Edit the "etc/shiro.ini" file and replace the following:

```
/** = authcBasic
```

with

```
/** = anon
```

Then restart the karaf process.

How application developers can leverage AAA to provide servlet security

In order to provide security to a servlet, add the following to the servlet's web.xml file as the first filter definition:

```
<context-param>
  <param-name>shiroEnvironmentClass</param-name>
  <param-value>org.opendaylight.aaa.shiro.web.env.KarafIniWebEnvironment</param-value>
</context-param>

<listener>
    <listener-class>org.apache.shiro.web.env.EnvironmentLoaderListener</listener-
↪class>
</listener>

<filter>
    <filter-name>ShiroFilter</filter-name>
    <filter-class>org.opendaylight.aaa.shiro.filters.AAAShiroFilter</filter-class>
</filter>

<filter-mapping>
    <filter-name>AAAShiroFilter</filter-name>
    <url-pattern>/*</url-pattern>
</filter-mapping>
```

Note: It is very important to place this AAAShiroFilter as the first javax.servlet.Filter, as Jersey applies Filters in the order they appear within web.xml. Placing the AAAShiroFilter first ensures incoming HTTP/HTTPS requests have proper credentials before any other filtering is attempted.

AAA Realms

AAA plugin utilizes realms to support pluggable authentication & authorization schemes. There are two parent types of realms:

- AuthenticatingRealm
  - Provides no Authorization capability.
  - Users authenticated through this type of realm are treated equally.
- AuthorizingRealm
  - AuthorizingRealm is a more sophisticated AuthenticatingRealm, which provides the additional mechanisms to distinguish users based on roles.
  - Useful for applications in which roles determine allowed cabilities.

ODL Contains Four Implementations

- TokenAuthRealm
  - An AuthorizingRealm built to bridge the Shiro-based AAA service with the h2-based AAA implementation.
  - Exposes a RESTful web service to manipulate IdM policy on a per-node basis. If identical AAA policy is desired across a cluster, the backing data store must be synchronized using an out of band method.
  - A python script located at "etc/idmtool" is included to help manipulate data contained in the TokenAuthRealm.
  - Enabled out of the box.
- ODLJndiLdapRealm
  - An AuthorizingRealm built to extract identity information from IdM data contained on an LDAP server.
  - Extracts group information from LDAP, which is translated into ODL roles.
  - Useful when federating against an existing LDAP server, in which only certain types of users should have certain access privileges.
  - Disabled out of the box.
- ODLJndiLdapRealmAuthNOnly
  - The same as ODLJndiLdapRealm, except without role extraction. Thus, all LDAP users have equal authentication and authorization rights.
  - Disabled out of the box.
- ActiveDirectoryRealm

Note: More than one Realm implementation can be specified. Realms are attempted in order until authentication succeeds or all realm sources are exhausted.

TokenAuthRealm Configuration

TokenAuthRealm stores IdM data in an h2 database on each node. Thus, configuration of a cluster currently requires configuring the desired IdM policy on each node. There are two supported methods to manipulate the TokenAuthRealm IdM configuration:

- idmtool Configuration
- RESTful Web Service Configuration

idmtool Configuration

A utility script located at "etc/idmtool" is used to manipulate the TokenAuthRealm IdM policy. idmtool assumes a single domain (sdn), since multiple domains are not leveraged in the Boron release. General usage information for idmtool is derived through issuing the following command:

```
$ python etc/idmtool -h
usage: idmtool [-h] [--target-host TARGET_HOST]
               user
               {list-users,add-user,change-password,delete-user,list-domains,list-
↪roles,add-role,delete-role,add-grant,get-grants,delete-grant}
               ...
```

(continues on next page)

(continued from previous page)

```
positional arguments:
  user                  username for BSC node
  {list-users,add-user,change-password,delete-user,list-domains,list-roles,add-role,
↪delete-role,add-grant,get-grants,delete-grant}
                        sub-command help
    list-users          list all users
    add-user            add a user
    change-password     change a password
    delete-user         delete a user
    list-domains        list all domains
    list-roles          list all roles
    add-role            add a role
    delete-role         delete a role
    add-grant           add a grant
    get-grants          get grants for userid on sdn
    delete-grant        delete a grant

optional arguments:
  -h, --help            show this help message and exit
  --target-host TARGET_HOST
                        target host node
```

Add a user

```
python etc/idmtool admin add-user newUser
Password:
Enter new password:
Re-enter password:
add_user(admin)

command succeeded!

json:
{
    "description": "",
    "domainid": "sdn",
    "email": "",
    "enabled": true,
    "name": "newUser",
    "password": "**********",
    "salt": "**********",
    "userid": "newUser@sdn"
}
```

Note: AAA redacts the password and salt fields for security purposes.

Delete a user

```
$ python etc/idmtool admin delete-user newUser@sdn
Password:
delete_user(newUser@sdn)

command succeeded!
```

List all users

```
$ python etc/idmtool admin list-users
Password:
list_users

command succeeded!

json:
{
    "users": [
        {
            "description": "user user",
            "domainid": "sdn",
            "email": "",
            "enabled": true,
            "name": "user",
            "password": "**********",
            "salt": "**********",
            "userid": "user@sdn"
        },
        {
            "description": "admin user",
            "domainid": "sdn",
            "email": "",
            "enabled": true,
            "name": "admin",
            "password": "**********",
            "salt": "**********",
            "userid": "admin@sdn"
        }
    ]
}
```

Change a user's password

```
$ python etc/idmtool admin change-password admin@sdn
Password:
Enter new password:
Re-enter password:
change_password(admin)

command succeeded!

json:
{
    "description": "admin user",
```

(continues on next page)

(continued from previous page)

```
    "domainid": "sdn",
    "email": "",
    "enabled": true,
    "name": "admin",
    "password": "**********",
    "salt": "**********",
    "userid": "admin@sdn"
}
```

Add a role

```
$ python etc/idmtool admin add-role network-admin
Password:
add_role(network-admin)

command succeeded!

json:
{
    "description": "",
    "domainid": "sdn",
    "name": "network-admin",
    "roleid": "network-admin@sdn"
}
```

Delete a role

```
$ python etc/idmtool admin delete-role network-admin@sdn
Password:
delete_role(network-admin@sdn)

command succeeded!
```

List all roles

```
$ python etc/idmtool admin list-roles
Password:
list_roles

command succeeded!

json:
{
    "roles": [
        {
            "description": "a role for admins",
            "domainid": "sdn",
            "name": "admin",
            "roleid": "admin@sdn"
        },
```

(continues on next page)

(continued from previous page)

```
        {
            "description": "a role for users",
            "domainid": "sdn",
            "name": "user",
            "roleid": "user@sdn"
        }
    ]
}
```

List all domains

```
$ python etc/idmtool admin list-domains
Password:
list_domains

command succeeded!

json:
{
    "domains": [
        {
            "description": "default odl sdn domain",
            "domainid": "sdn",
            "enabled": true,
            "name": "sdn"
        }
    ]
}
```

Add a grant

```
$ python etc/idmtool admin add-grant user@sdn admin@sdn
Password:
add_grant(userid=user@sdn,roleid=admin@sdn)

command succeeded!

json:
{
    "domainid": "sdn",
    "grantid": "user@sdn@admin@sdn@sdn",
    "roleid": "admin@sdn",
    "userid": "user@sdn"
}
```

Delete a grant

```
$ python etc/idmtool admin delete-grant user@sdn admin@sdn
Password:
http://localhost:8181/auth/v1/domains/sdn/users/user@sdn/roles/admin@sdn
```

(continues on next page)

(continued from previous page)

```
delete_grant(userid=user@sdn,roleid=admin@sdn)

command succeeded!
```

Get grants for a user

```
python etc/idmtool admin get-grants admin@sdn
Password:
get_grants(admin@sdn)

command succeeded!

json:
{
    "roles": [
        {
            "description": "a role for users",
            "domainid": "sdn",
            "name": "user",
            "roleid": "user@sdn"
        },
        {
            "description": "a role for admins",
            "domainid": "sdn",
            "name": "admin",
            "roleid": "admin@sdn"
        }
    ]
}
```

RESTful Web Service

The TokenAuthRealm IdM policy is fully configurable through a RESTful web service. Full documentation for manipulating AAA IdM data is located online (https://wiki.opendaylight.org/images/0/00/AAA_Test_Plan.docx), and a few examples are included in this guide:

Get All Users

```
curl -u admin:admin http://localhost:8181/auth/v1/users
OUTPUT:
{
    "users": [
        {
            "description": "user user",
            "domainid": "sdn",
            "email": "",
            "enabled": true,
            "name": "user",
            "password": "**********",
            "salt": "**********",
            "userid": "user@sdn"
```

(continues on next page)

(continued from previous page)

```
        },
        {
            "description": "admin user",
            "domainid": "sdn",
            "email": "",
            "enabled": true,
            "name": "admin",
            "password": "**********",
            "salt": "**********",
            "userid": "admin@sdn"
        }
    ]
}
```

Create a User

```
curl -u admin:admin -X POST -H "Content-Type: application/json" --data-binary @./user.
↪json http://localhost:8181/auth/v1/users
PAYLOAD:
{
    "name": "ryan",
    "userid": "ryan@sdn",
    "password": "ryan",
    "domainid": "sdn",
    "description": "Ryan's User Account",
    "email": "ryandgoulding@gmail.com"
}

OUTPUT:
{
    "userid":"ryan@sdn",
    "name":"ryan",
    "description":"Ryan's User Account",
    "enabled":true,
    "email":"ryandgoulding@gmail.com",
    "password":"**********",
    "salt":"**********",
    "domainid":"sdn"
}
```

Create an OAuth2 Token For Admin Scoped to SDN

```
curl -d 'grant_type=password&username=admin&password=a&scope=sdn' http://
↪localhost:8181/oauth2/token

OUTPUT:
{
    "expires_in":3600,
    "token_type":"Bearer",
    "access_token":"5a615fbc-bcad-3759-95f4-ad97e831c730"
}
```

Use an OAuth2 Token

```
curl -H "Authorization: Bearer 5a615fbc-bcad-3759-95f4-ad97e831c730" http://
↪localhost:8181/auth/v1/domains
{
    "domains":
    [
        {
            "domainid":"sdn",
            "name":"sdn",
            "description":"default odl sdn domain",
            "enabled":true
        }
    ]
}
```

ODLJndiLdapRealm Configuration

LDAP integration is provided in order to externalize identity management. To configure LDAP parameters, modify "etc/shiro.ini" parameters to include the ODLJndiLdapRealm:

```
# ODL provides a few LDAP implementations, which are disabled out of the box.
# ODLJndiLdapRealm includes authorization functionality based on LDAP elements
# extracted through and LDAP search.  This requires a bit of knowledge about
# how your LDAP system is setup.  An example is provided below:
ldapRealm = org.opendaylight.aaa.shiro.realm.ODLJndiLdapRealm
ldapRealm.userDnTemplate = uid={0},ou=People,dc=DOMAIN,dc=TLD
ldapRealm.contextFactory.url = ldap://<URL>:389
ldapRealm.searchBase = dc=DOMAIN,dc=TLD
ldapRealm.ldapAttributeForComparison = objectClass
ldapRealm.groupRolesMap = "Person":"admin"
# ...
# further down in the file...
# Stacked realm configuration;  realms are round-robbined until authentication␣
↪succeeds or realm sources are exhausted.
securityManager.realms = $tokenAuthRealm, $ldapRealm
```

This configuration allows federation with an external LDAP server, and the user's ODL role parameters are mapped to corresponding LDAP attributes as specified by the groupRolesMap. Thus, an LDAP operator can provision attributes for LDAP users that support different ODL role structures.

ODLJndiLdapRealmAuthNOnly Configuration

Edit the "etc/shiro.ini" file and modify the following:

```
ldapRealm = org.opendaylight.aaa.shiro.realm.ODLJndiLdapRealm
ldapRealm.userDnTemplate = uid={0},ou=People,dc=DOMAIN,dc=TLD
ldapRealm.contextFactory.url = ldap://<URL>:389
# ...
# further down in the file...
# Stacked realm configuration;  realms are round-robbined until authentication␣
↪succeeds or realm sources are exhausted.
securityManager.realms = $tokenAuthRealm, $ldapRealm
```

This is useful for setups where all LDAP users are allowed equal access.

Token Store Configuration Parameters

Edit the file "etc/opendaylight/karaf/08-authn-config.xml" and edit the following: .**timeToLive**: Configure the maximum time, in milliseconds, that tokens are to be cached. Default is 360000. Save the file.

Authorization Configuration

Shiro-Based Authorization

OpenDaylight AAA has support for Role Based Access Control based on the Apache Shiro permissions system. Configuration of the authorization system is done offline; authorization currently cannot be configured after the controller is started. Thus, Authorization in this release is aimed towards supporting coarse-grained security policies, with the aim to provide more robust configuration capabilities in the future. Shiro-based Authorization is documented on the Apache Shiro website (http://shiro.apache.org/web.html#Web-%7B%7B%5Curls%5C%7D%7D).

Enable "admin" Role Based Access to the IdMLight RESTful web service

Edit the "etc/shiro.ini" configuration file and add "/auth/v1/**= authcBasic, roles[admin]" above the line "/** = authcBasic" within the "urls" section.

```
/auth/v1/** = authcBasic, roles[admin]
/** = authcBasic
```

This will restrict the idmlight rest endpoints so that a grant for admin role must be present for the requesting user.

Note: The ordering of the authorization rules above is important!

AuthZ Broker Facade

ODL includes an experimental Authorization Broker Facade, which allows finer grained access control for REST endpoints. Since this feature was not well tested in the Boron release, it is recommended to use the Shiro-based mechanism instead, and rely on the Authorization Broker Facade for POC use only.

AuthZ Broker Facade Feature Installation

To install the authorization broker facade, please issue the following command in the karaf shell:

```
feature:install odl-restconf odl-aaa-authz
```

Add an Authorization Rule

The following shows how one might go about securing the controller so that only admins can access restconf.

```
curl -u admin:admin -H "Content-Type: application/xml" --data-binary @./rule.json␣
↪http://localhost:8181/restconf/config/authorization-schema:simple-authorization/
↪policies/RestConfService/
cat ./rule.json
{
    "policies": {
        "resource": "*",
        "service":"RestConfService",
        "role": "admin"
    }
}
```

Accounting Configuration

All AAA logging is output to the standard karaf.log file.

```
log:set TRACE org.opendaylight.aaa
```

This command enables the most verbose level of logging for AAA components.

Atrium User Guide

Overview

Project Atrium is an open source SDN distribution - a vertically integrated set of open source components which together form a complete SDN stack. It's goals are threefold:

- Close the large integration-gap of the elements that are needed to build an SDN stack - while there are multiple choices at each layer, there are missing pieces with poor or no integration.
- Overcome a massive gap in interoperability - This exists both at the switch level, where existing products from different vendors have limited compatibility, making it difficult to connect an arbitrary switch and controller and at an API level, where its difficult to write a portable application across multiple controller platforms.
- Work closely with network operators on deployable use-cases, so that they could download near production quality code from one location, and get started with functioning software defined networks on real hardware.

Architecture

The key components of Atrium BGP Peering Router Application are as follows:

- Data Plane Switch - Data plane switch is the entity that uses flow table entries installed by BGP Routing Application through SDN controller. In the simplest form data plane switch with the installed flows act like a BGP Router.
- OpenDaylight Controller - OpenDaylight SDN controller has many utility applications or plugins which are leveraged by the BGP Router application to manage the control plane information.
- BGP Routing Application - An application running within the OpenDaylight runtime environment to handle I-BGP updates.
- *DIDM* - DIDM manages the drivers specific to each data plane switch connected to the controller. The drivers are created primarily to hide the underlying complexity of the devices and to expose a uniform API to applications.

- Flow Objectives API - The driver implementation provides a pipeline abstraction and exposes Flow Objectives API. This means applications need to be aware of only the Flow Objectives API without worrying about the Table IDs or the pipelines.
- Control Plane Switch - This component is primarily used to connect the OpenDaylight SDN controller with the Quagga Soft-Router and establish a path for forwarding E-BGP packets to and from Quagga.
- Quagga soft router - An open source routing software that handles E-BGP updates.

Running Atrium

- To run the Atrium BGP Routing Application in OpenDaylight distribution, simply install the `odl-atrium-all` feature.

```
feature:install odl-atrium-all
```

BGP User Guide

This guide contains information on how to use OpenDaylight Border Gateway Protocol (BGP) plugin. The user should learn about BGP basic concepts, supported capabilities, configuration and usage.

Contents

- *Overview*
- *Running BGP*
- *Basic Configuration & Concepts*
- *IP Unicast Family*
- *IP Labeled Unicast Family*
- *IP L3VPN Family*
- *Link-State Family*
- *Flow Specification Family*
- *EVPN Family*
- *Additional Path*
- *Route Refresh*
- *High Availability*
- *Topology Provider*
- *Test Tools*
- *Troubleshooting*

Overview

This section provides high-level overview of the Border Gateway Protocol, OpenDaylight implementation and BGP usage in SDN era.

Contents

- *Border Gateway Protocol*
- *BGP in SDN*
- *OpenDaylight BGP plugin*
- *List of supported capabilities*

Border Gateway Protocol

The Border Gateway Protocol (BGP) is an inter-Autonomous System (AS) routing protocol. The primary role of the BGP is an exchange of routes among other BGP systems. The route is an unit of information which pairs destination (IP address prefix) with attributes to the path with the destination. One of the most interesting attributes is a list of ASes that the route traversed - essential when avoiding loop routing. Advertised routes are stored in the Routing Information Bases (RIBs). Routes are later used to forward packets, stored in Routing Table for this purpose. The main advantage of the BGP over other routing protocols is its scalability, thus it has become the standardized Internet routing protocol (Internet is a set of ASes).

BGP in SDN

However BGP evolved long time before SDN was born, it plays a significant role in many SDN use-cases. Also, continuous evolution of the protocol brings extensions that are very well suited for SDN. Nowadays, BGP can carry various types of routing information - L3VPN, L2VPN, IP multicast, linkstate, etc. Here is a brief list of software-based/legacy-network technologies where BGP-based SDN solution get into an action:

- SDN WAN - WAN orchestration and optimization
- SDN router - Turns switch into an Internet router
- Virtual Route Reflector - High-performance server-based BGP Route Reflector
- SDX - A Software Defined Internet Exchange controller
- Large-Scale Data Centers - BGP Data Center Routing, MPLS/SR in DCs, DC interconnection
- DDoS mitigation - Traffic Filtering distribution with BGP

OpenDaylight BGP plugin

The OpenDaylight controller provides an implementation of BGP (RFC 4271) as a south-bound protocol plugin. The implementation renders all basic *BGP speaker capabilities*:

- inter/intra-AS peering
- routes advertising
- routes originating
- routes storage

The plugin's **north-bound API** (`REST/Java`) provides to user:

- fully dynamic runtime standardized BGP configuration
- read-only access to all RIBs

- read-write programmable RIBs
- read-only reachability/linkstate topology view

Note: The BGP plugin is NOT a virtual router - does not construct Routing Tables, nor forward traffic.

List of supported capabilities

In addition to the base protocol implementation, the plugin provides many extensions to BGP, all based on IETF standards.

- RFC4271 - A Border Gateway Protocol 4 (BGP-4)
- RFC4456 - BGP Route Reflection: An Alternative to Full Mesh Internal BGP (IBGP)
- RFC1997 - BGP Communities Attribute
- RFC4360 - BGP Extended Communities Attribute
- RFC4486 - Subcodes for BGP Cease Notification Message
- RFC5492 - Capabilities Advertisement with BGP-4
- RFC5004 - Avoid BGP Best Path Transitions from One External to Another
- RFC6286 - Autonomous-System-Wide Unique BGP Identifier for BGP-4
- RFC6793 - BGP Support for Four-Octet Autonomous System (AS) Number Space
- RFC7311 - The Accumulated IGP Metric Attribute for BGP
- RFC5668 - 4-Octet AS Specific BGP Extended Community
- draft-ietf-idr-link-bandwidth - BGP Link Bandwidth Extended Community
- draft-ietf-idr-bgp-extended-messages - Extended Message support for BGP
- **RFC4760 - Multiprotocol Extensions for BGP-4**
  - **RFC7752 - North-Bound Distribution of Link-State and TE Information using BGP**
    - draft-gredler-idr-bgp-ls-segment-routing-ext - BGP Link-State extensions for Segment Routing
    - draft-ietf-idr-bgpls-segment-routing-epe - Segment Routing Egress Peer Engineering BGP-LS Extensions
  - **RFC5575 - Dissemination of Flow Specification Rules**
    - RFC7674 - Clarification of the Flowspec Redirect Extended Community
    - draft-ietf-idr-flow-spec-v6 - Dissemination of Flow Specification Rules for IPv6
    - draft-ietf-idr-flowspec-redirect-ip - BGP Flow-Spec Redirect to IP Action
  - **RFC3107 - Carrying Label Information in BGP-4**
    - draft-ietf-idr-bgp-prefix-sid - Segment Routing Prefix SID extensions for BGP
  - **RFC4364 - BGP/MPLS IP Virtual Private Networks (VPNs)**
    - RFC4659 - BGP-MPLS IP Virtual Private Network (VPN) Extension for IPv6 VPN
  - **RFC7432 - BGP MPLS-Based Ethernet VPN**
    - draft-ietf-bess-evpn-overlay - A Network Virtualization Overlay Solution using EVPN

        * draft-ietf-bess-evpn-vpws - VPWS support in EVPN
- RFC7911 - Advertisement of Multiple Paths in BGP
- RFC2918 - Route Refresh Capability for BGP-4

Running BGP

This section explains how to install BGP plugin.

1. Install BGP feature - `odl-bgpcep-bgp`. Also, for sake of this sample, it is required to install RESTCONF. In the Karaf console, type command:

```
feature:install odl-restconf odl-bgpcep-bgp
```

2. The BGP plugin contains a default configuration, which is applied after the feature starts up. One instance of BGP plugin is created (named *example-bgp-rib*), and its presence can be verified via REST:

   URL: `/restconf/operational/bgp-rib:bgp-rib`

   Method: `GET`

   Response Body:

```
<bgp-rib xmlns="urn:opendaylight:params:xml:ns:yang:bgp-rib">
   <rib>
       <id>example-bgp-rib</id>
       <loc-rib>
       ....
       </loc-rib>
   </rib>
</bgp-rib>
```

Basic Configuration & Concepts

The following section shows how to configure BGP basics, how to verify functionality and presents essential components of the plugin. Next samples demonstrate the plugin's runtime configuration capability. It shows the way to configure the plugin via REST, using standardized OpenConfig BGP APIs.

Contents

- *BGP RIB API*
- *Protocol Configuration*
- *BGP Peering*
  - *External peering configuration*
  - *Route reflector configuration*
  - *MD5 authentication configuration*
  - *Simple Routing Policy configuration*
- *BGP Application Peer and programmable RIB*
  - *Application Peer configuration*

  - *Programmable RIB*
- *BGP pipeline*
- *References*

BGP RIB API

This tree illustrates the BGP RIBs organization in datastore.

```
bgp-rib
  +--ro rib* [id]
     +--ro id          rib-id
     +--ro peer* [peer-id]
     |  +--ro peer-id                  peer-id
     |  +--ro peer-role                peer-role
     |  +--ro simple-routing-policy?   simple-routing-policy
     |  +--ro supported-tables* [afi safi]
     |  |  +--ro afi             identityref
     |  |  +--ro safi            identityref
     |  |  +--ro send-receive?   send-receive
     |  +--ro adj-rib-in
     |  |  +--ro tables* [afi safi]
     |  |     +--ro afi           identityref
     |  |     +--ro safi          identityref
     |  |     +--ro attributes
     |  |     |  +--ro uptodate?   boolean
     |  |     +--ro (routes)?
     |  +--ro effective-rib-in
     |  |  +--ro tables* [afi safi]
     |  |     +--ro afi           identityref
     |  |     +--ro safi          identityref
     |  |     +--ro attributes
     |  |     |  +--ro uptodate?   boolean
     |  |     +--ro (routes)?
     |  +--ro adj-rib-out
     |     +--ro tables* [afi safi]
     |        +--ro afi           identityref
     |        +--ro safi          identityref
     |        +--ro attributes
     |        |  +--ro uptodate?   boolean
     |        +--ro (routes)?
     +--ro loc-rib
        +--ro tables* [afi safi]
           +--ro afi           identityref
           +--ro safi          identityref
           +--ro attributes
           |  +--ro uptodate?   boolean
           +--ro (routes)?
```

Protocol Configuration

As a first step, a new protocol instance needs to be configured. It is a very basic configuration conforming with RFC4271.

URL: `/restconf/config/openconfig-network-instance:network-instances/network-instance/global-bgp/openconfig-network-instance:protocols`

Method: `POST`

Content-Type: `application/xml`

Request Body:

```
1  <protocol xmlns="http://openconfig.net/yang/network-instance">
2      <name>bgp-example</name>
3      <identifier xmlns:x="http://openconfig.net/yang/policy-types">x:BGP</identifier>
4      <bgp xmlns="urn:opendaylight:params:xml:ns:yang:bgp:openconfig-extensions">
5          <global>
6              <config>
7                  <router-id>192.0.2.2</router-id>
8                  <as>65000</as>
9              </config>
10         </global>
11     </bgp>
12 </protocol>
```

@line 2: The unique protocol instance identifier.

@line 7: BGP Identifier of the speaker.

@line 8: Local autonomous system number of the speaker. Note that, OpenDaylight BGP implementation supports four-octet AS numbers only.

The new instance presence can be verified via REST:

URL: `/restconf/operational/bgp-rib:bgp-rib/rib/bgp-example`

Method: `GET`

Response Body:

```
1  <rib xmlns="urn:opendaylight:params:xml:ns:yang:bgp-rib">
2      <id>bgp-example</id>
3      <loc-rib>
4          <tables>
5              <afi xmlns:x="urn:opendaylight:params:xml:ns:yang:bgp-types">x:ipv4-
   ↪address-family</afi>
6              <safi xmlns:x="urn:opendaylight:params:xml:ns:yang:bgp-types">x:unicast-
   ↪subsequent-address-family</safi>
7              <ipv4-routes xmlns="urn:opendaylight:params:xml:ns:yang:bgp-inet"></ipv4-
   ↪routes>
8              <attributes>
9                  <uptodate>true</uptodate>
10             </attributes>
11         </tables>
12     </loc-rib>
13 </rib>
```

@line 3: Loc-RIB - Per-protocol instance RIB, which contains the routes that have been selected by local BGP speaker's decision process.

@line 4: The BGP-4 supports carrying IPv4 prefixes, such routes are stored in *ipv4-address-family/unicast-subsequent-address-family* table.

BGP Peering

To exchange routing information between two BGP systems (peers), it is required to configure a peering on both BGP speakers first. This mean that each BGP speaker has a white list of neighbors, representing remote peers, with which the peering is allowed. BGP uses TCP as its transport protocol, by default listens on port 179.

Important: OpenDaylight BGP plugin is configured to listen on port *1790*, due to privileged ports restriction for non-root users. One of the workarounds is to use port redirection.

The TCP connection is established between two peers and they exchange messages to open and confirm the connection parameters followed by routes exchange.

Here is a sample basic neighbor configuration:

URL: `/restconf/config/openconfig-network-instance:network-instances/network-instance/global-bgp/openconfig-network-instance:protocols/protocol/openconfig-policy-types:BGP/bgp-example/bgp/neighbors`

Method: `POST`

Content-Type: `application/xml`

Request Body:

```
1  <neighbor xmlns="urn:opendaylight:params:xml:ns:yang:bgp:openconfig-extensions">
2      <neighbor-address>192.0.2.1</neighbor-address>
3      <timers>
4          <config>
5              <hold-time>90</hold-time>
6              <connect-retry>10</connect-retry>
7          </config>
8      </timers>
9      <transport>
10         <config>
11             <remote-port>179</remote-port>
12             <passive-mode>false</passive-mode>
13         </config>
14     </transport>
15     <config>
16         <peer-type>INTERNAL</peer-type>
17     </config>
18 </neighbor>
```

@line 2: IP address of the remote BGP peer. Also serves as an unique identifier of a neighbor in a list of neighbors.

@line 5: Proposed number of seconds for value of the Hold Timer. Default value is **90**.

@line 6: Time interval in seconds between attempts to establish session with the peer. Effective in active mode only. Default value is **30**.

@line 11: Remote port number to which the local BGP is connecting. Effective in active mode only. Default value **179**.

@line 12: Wait for peers to issue requests to open a BGP session, rather than initiating sessions from the local router. Default value is **false**.

@line 16: Explicitly designate the peer as internal or external. Default value is **INTERNAL**.

Once the remote peer is connected and it advertised routes to local BGP system, routes are stored in peer's RIBs. The RIBs can be checked via REST:

URL: `/restconf/operational/bgp-rib:bgp-rib/rib/bgp-example/peer/bgp:%2F%2F192.0.2.1`

Method: `GET`

Response Body:

```
<peer xmlns="urn:opendaylight:params:xml:ns:yang:bgp-rib">
    <peer-id>bgp://192.0.2.1</peer-id>
    <supported-tables>
        <afi xmlns:x="urn:opendaylight:params:xml:ns:yang:bgp-types">x:ipv4-address-family</afi>
        <safi xmlns:x="urn:opendaylight:params:xml:ns:yang:bgp-types">x:unicast-subsequent-address-family</safi>
    </supported-tables>
    <peer-role>ibgp</peer-role>
    <adj-rib-in>
        <tables>
            <afi xmlns:x="urn:opendaylight:params:xml:ns:yang:bgp-types">x:ipv4-address-family</afi>
            <safi xmlns:x="urn:opendaylight:params:xml:ns:yang:bgp-types">x:unicast-subsequent-address-family</safi>
            <ipv4-routes xmlns="urn:opendaylight:params:xml:ns:yang:bgp-inet">
                <ipv4-route>
                    <path-id>0</path-id>
                    <prefix>10.0.0.10/32</prefix>
                    <attributes>
                        <as-path></as-path>
                        <origin>
                            <value>igp</value>
                        </origin>
                        <local-pref>
                            <pref>100</pref>
                        </local-pref>
                        <ipv4-next-hop>
                            <global>10.10.1.1</global>
                        </ipv4-next-hop>
                    </attributes>
                </ipv4-route>
            </ipv4-routes>
            <attributes>
                <uptodate>true</uptodate>
            </attributes>
        </tables>
    </adj-rib-in>
    <effective-rib-in>
        <tables>
            <afi xmlns:x="urn:opendaylight:params:xml:ns:yang:bgp-types">x:ipv4-address-family</afi>
            <safi xmlns:x="urn:opendaylight:params:xml:ns:yang:bgp-types">x:unicast-subsequent-address-family</safi>
            <ipv4-routes xmlns="urn:opendaylight:params:xml:ns:yang:bgp-inet">
                <ipv4-route>
                    <path-id>0</path-id>
                    <prefix>10.0.0.10/32</prefix>
                    <attributes>
```

(continues on next page)

(continued from previous page)

```
44                          <as-path></as-path>
45                          <origin>
46                              <value>igp</value>
47                          </origin>
48                          <local-pref>
49                              <pref>100</pref>
50                          </local-pref>
51                          <ipv4-next-hop>
52                              <global>10.10.1.1</global>
53                          </ipv4-next-hop>
54                      </attributes>
55                  </ipv4-route>
56              </ipv4-routes>
57              <attributes>
58                  <uptodate>true</uptodate>
59              </attributes>
60          </tables>
61      </effective-rib-in>
62      <adj-rib-out>
63          <tables>
64              <afi xmlns:x="urn:opendaylight:params:xml:ns:yang:bgp-types">x:ipv4-
    ↪address-family</afi>
65              <safi xmlns:x="urn:opendaylight:params:xml:ns:yang:bgp-types">x:unicast-
    ↪subsequent-address-family</safi>
66              <ipv4-routes xmlns="urn:opendaylight:params:xml:ns:yang:bgp-inet"></ipv4-
    ↪routes>
67              <attributes></attributes>
68          </tables>
69      </adj-rib-out>
70  </peer>
```

@line 8: **Adj-RIB-In** - Per-peer RIB, which contains unprocessed routes that has been advertised to local BGP speaker by the remote peer.

@line 13: Here is the reported route with destination *10.0.0.10/32* in Adj-RIB-In.

@line 35: **Effective-RIB-In** - Per-peer RIB, which contains processed routes as a result of applying inbound policy to Adj-RIB-In routes.

@line 40: Here is the reported route with destination *10.0.0.10/32*, same as in Adj-RIB-In, as it was not touched by import policy.

@line 62: **Adj-RIB-Out** - Per-peer RIB, which contains routes for advertisement to the peer by means of the local speaker's UPDATE message.

@line 66: The peer's Adj-RIB-Out is empty as there are no routes to be advertise from local BGP speaker.

Also the same route should appeared in Loc-RIB now:

URL: `/restconf/operational/bgp-rib:bgp-rib/rib/bgp-example/loc-rib/tables/bgp-types:ipv4-address-family/bgp-types:unicast-subsequent-address-family/ipv4-routes`

Method: `GET`

Response Body:

```
1  <ipv4-routes xmlns="urn:opendaylight:params:xml:ns:yang:bgp-inet">
2      <ipv4-route>
3          <path-id>0</path-id>
4          <prefix>10.0.0.10/32</prefix>
5          <attributes>
6              <as-path></as-path>
7              <origin>
8                  <value>igp</value>
9              </origin>
10             <local-pref>
11                 <pref>100</pref>
12             </local-pref>
13             <ipv4-next-hop>
14                 <global>10.10.1.1</global>
15             </ipv4-next-hop>
16         </attributes>
17     </ipv4-route>
18 </ipv4-routes>
```

@line 4: **Destination** - IPv4 Prefix Address.

@line 6: **AS_PATH** - mandatory attribute, contains a list of the autonomous system numbers through that routing information has traversed.

@line 8: **ORIGIN** - mandatory attribute, indicates an origin of the route - **ibgp**, **egp**, **incomplete**.

@line 11: **LOCAL_PREF** - indicates a degree of preference for external routes, higher value is preferred.

@line 14: **NEXT_HOP** - mandatory attribute, defines IP address of the router that should be used as the next hop to the destination.

There are much more attributes that may be carried along with the destination:

BGP-4 Path Attributes

- **MULTI_EXIT_DISC (MED)** Optional attribute, to be used to discriminate among multiple exit/entry points on external links, lower number is preferred.

```
<multi-exit-disc>
 <med>0</med>
</multi-exit-disc>
```

- **ATOMIC_AGGREGATE** Indicates whether AS_SET was excluded from AS_PATH due to routes aggregation.

```
<atomic-aggregate/>
```

- **AGGREGATOR** Optional attribute, contains AS number and IP address of a BGP speaker which performed routes aggregation.

```
<aggregator>
    <as-number>65000</as-number>
    <network-address>192.0.2.2</network-address>
</aggregator>
```

- **Unrecognised** Optional attribute, used to store optional attributes, unrecognized by a local BGP speaker.

```
<unrecognized-attributes>
    <partial>true</partial>
    <transitive>true</transitive>
    <type>101</type>
    <value>0101010101010101</value>
</unrecognized-attributes>
```

Route Reflector Attributes

- **ORIGINATOR_ID** Optional attribute, carries BGP Identifier of the originator of the route.

```
<originator-id>
    <originator>41.41.41.41</originator>
</originator-id>
```

- **CLUSTER_LIST** Optional attribute, contains a list of CLUSTER_ID values representing the path that the route has traversed.

```
<cluster-id>
    <cluster>40.40.40.40</cluster>
</cluster-id>
```

- **Communities** Optional attribute, may be used for policy routing.

```
<communities>
    <as-number>65000</as-number>
    <semantics>30740</semantics>
</communities>
```

Extended Communities

- **Route Target** Identifies one or more routers that may receive a route.

```
<extended-communities>
    <transitive>true</transitive>
    <route-target-ipv4>
        <global-administrator>192.0.2.2</global-administrator>
        <local-administrator>123</local-administrator>
    </route-target-ipv4>
</extended-communities>
<extended-communities>
    <transitive>true</transitive>
    <as-4-route-target-extended-community>
            <as-4-specific-common>
            <as-number>65000</as-number>
            <local-administrator>123</local-administrator>
        </as-4-specific-common>
    </as-4-route-target-extended-community>
</extended-communities>
```

- **Route Origin** Identifies one or more routers that injected a route.

```
<extended-communities>
    <transitive>true</transitive>
    <route-origin-ipv4>
        <global-administrator>192.0.2.2</global-administrator>
        <local-administrator>123</local-administrator>
    </route-origin-ipv4>
```

(continues on next page)

(continued from previous page)

```
</extended-communities>
<extended-communities>
    <transitive>true</transitive>
    <as-4-route-origin-extended-community>
        <as-4-specific-common>
            <as-number>65000</as-number>
            <local-administrator>123</local-administrator>
        </as-4-origin-common>
    </as-4-route-target-extended-community>
</extended-communities>
```

- **Link Bandwidth** Carries the cost to reach external neighbor.

```
<extended-communities>
    <transitive>true</transitive>
    <link-bandwidth-extended-community>
        <bandwidth>BH9CQAA=</bandwidth>
    </link-bandwidth-extended-community>
</extended-communities>
```

- **AIGP** Optional attribute, carries accumulated IGP metric.

```
<aigp>
    <aigp-tlv>
        <metric>120</metric>
    </aigp-tlv>
</aigp>
```

Note: When the remote peer disconnects, it disappear from operational state of local speaker instance and advertised routes are removed too.

External peering configuration

An example above provided configuration for internal peering only. Following configuration sample is intended for external peering:

URL: `/restconf/config/openconfig-network-instance:network-instances/network-instance/global-bgp/openconfig-network-instance:protocols/protocol/openconfig-policy-types:BGP/bgp-example/bgp/neighbors`

Method: `POST`

Content-Type: `application/xml`

Request Body:

```
1 <neighbor xmlns="urn:opendaylight:params:xml:ns:yang:bgp:openconfig-extensions">
2     <neighbor-address>192.0.2.3</neighbor-address>
3     <config>
4         <peer-type>EXTERNAL</peer-type>
5         <peer-as>64999</peer-as>
6     </config>
7 </neighbor>
```

@line 5: AS number of the remote peer.

Route reflector configuration

The local BGP speaker can be configured with a specific *cluster ID*. Following example adds the cluster ID to the existing speaker instance:

URL: `/restconf/config/openconfig-network-instance:network-instances/network-instance/global-bgp/openconfig-network-instance:protocols/protocol/openconfig-policy-types:BGP/bgp-example/bgp/global/config`

Method: `PUT`

Content-Type: `application/xml`

Request Body:

```
1  <config>
2      <router-id>192.0.2.2</router-id>
3      <as>65000</as>
4      <route-reflector-cluster-id>192.0.2.1</route-reflector-cluster-id>
5  </config>
```

@line 4: Route-reflector cluster id to use when local router is configured as a route reflector. The *router-id* is used as a default value.

Following configuration sample is intended for route reflector client peering:

URL: `/restconf/config/openconfig-network-instance:network-instances/network-instance/global-bgp/openconfig-network-instance:protocols/protocol/openconfig-policy-types:BGP/bgp-example/bgp/neighbors`

Method: `POST`

Content-Type: `application/xml`

Request Body:

```
1  <neighbor xmlns="urn:opendaylight:params:xml:ns:yang:bgp:openconfig-extensions">
2      <neighbor-address>192.0.2.4</neighbor-address>
3      <config>
4          <peer-type>INTERNAL</peer-type>
5      </config>
6      <route-reflector>
7          <config>
8              <route-reflector-client>true</route-reflector-client>
9          </config>
10     </route-reflector>
11 </neighbor>
```

@line 8: Configure the neighbor as a route reflector client. Default value is *false*.

MD5 authentication configuration

The OpenDaylight BGP implementation is supporting TCP MD5 for authentication. Sample configuration below shows how to set authentication password for a peer:

URL: `/restconf/config/openconfig-network-instance:network-instances/network-instance/global-bgp/openconfig-network-instance:protocols/protocol/openconfig-policy-types:BGP/bgp-example/bgp/neighbors`

Method: `POST`

Content-Type: `application/xml`

Request Body:

```
1 <neighbor xmlns="urn:opendaylight:params:xml:ns:yang:bgp:openconfig-extensions">
2     <neighbor-address>192.0.2.5</neighbor-address>
3     <config>
4         <auth-password>topsecret</auth-password>
5     </config>
6 </neighbor>
```

@line 4: Configures an MD5 authentication password for use with neighboring devices.

Simple Routing Policy configuration

The OpenDaylight BGP implementation is supporting *Simple Routing Policy*. Sample configuration below shows how to set *Simple Routing Policy* for a peer:

URL: `/restconf/config/openconfig-network-instance:network-instances/network-instance/global-bgp/openconfig-network-instance:protocols/protocol/openconfig-policy-types:BGP/bgp-example/bgp/neighbors`

Method: `POST`

Content-Type: `application/xml`

Request Body:

```
1 <neighbor xmlns="urn:opendaylight:params:xml:ns:yang:bgp:openconfig-extensions">
2     <neighbor-address>192.0.2.7</neighbor-address>
3     <config>
4         <simple-routing-policy>learn-none</simple-routing-policy>
5     </config>
6 </neighbor>
```

@line 4: *Simple Routing Policy*:

- `learn-none` - routes advertised by the peer are not propagated to Effective-RIB-In and Loc-RIB
- `announce-none` - routes from local Loc-RIB are not advertised to the peer

---

Note: Existing neighbor configuration can be reconfigured (change configuration parameters) anytime. As a result, established connection is dropped, peer instance is recreated with a new configuration settings and connection re-established.

Note: The BGP configuration is persisted on OpendDaylight shutdown and restored after the re-start.

BGP Application Peer and programmable RIB

The OpenDaylight BGP implementation also supports routes injection via *Application Peer*. Such peer has its own programmable RIB, which can be modified by user. This concept allows user to originate new routes and advertise them to all connected peers.

Application Peer configuration

Following configuration sample show a way to configure the *Application Peer*:

URL: `/restconf/config/openconfig-network-instance:network-instances/network-instance/global-bgp/openconfig-network-instance:protocols/protocol/openconfig-policy-types:BGP/bgp-example/bgp/neighbors`

Method: `POST`

Content-Type: `application/xml`

Request Body:

```
1  <neighbor xmlns="urn:opendaylight:params:xml:ns:yang:bgp:openconfig-extensions">
2      <neighbor-address>10.25.1.9</neighbor-address>
3      <config>
4          <peer-group>application-peers</peer-group>
5      </config>
6  </neighbor>
```

@line 2: IP address is uniquely identifying *Application Peer* and its programmable RIB. Address is also used in local BGP speaker decision process.

@line 4: Indicates that peer is associated with *application-peers* group. It serves to distinguish *Application Peer's* from regular neighbors.

The *Application Peer* presence can be verified via REST:

URL: `/restconf/operational/bgp-rib:bgp-rib/rib/bgp-example/peer/bgp:%2F%2F10.25.1.9`

Method: `GET`

Response Body:

```
1  <peer xmlns="urn:opendaylight:params:xml:ns:yang:bgp-rib">
2      <peer-id>bgp://10.25.1.9</peer-id>
3      <peer-role>internal</peer-role>
4      <adj-rib-in>
5          <tables>
6              <afi xmlns:x="urn:opendaylight:params:xml:ns:yang:bgp-types">x:ipv4-
   ↪address-family</afi>
7              <safi xmlns:x="urn:opendaylight:params:xml:ns:yang:bgp-types">x:unicast-
   ↪subsequent-address-family</safi>
8              <ipv4-routes xmlns="urn:opendaylight:params:xml:ns:yang:bgp-inet"></ipv4-
   ↪routes>
9              <attributes>
10                 <uptodate>false</uptodate>
11             </attributes>
12         </tables>
13     </adj-rib-in>
14     <effective-rib-in>
15         <tables>
16             <afi xmlns:x="urn:opendaylight:params:xml:ns:yang:bgp-types">x:ipv4-
   ↪address-family</afi>
17             <safi xmlns:x="urn:opendaylight:params:xml:ns:yang:bgp-types">x:unicast-
   ↪subsequent-address-family</safi>
18             <ipv4-routes xmlns="urn:opendaylight:params:xml:ns:yang:bgp-inet"></ipv4-
   ↪routes>
```

(continues on next page)

(continued from previous page)

```
            <attributes></attributes>
        </tables>
    </effective-rib-in>
</peer>
```

@line 3: Peer role for *Application Peer* is *internal*.

@line 8: Adj-RIB-In is empty, as no routes were originated yet.

Note: There is no Adj-RIB-Out for *Application Peer*.

Programmable RIB

Next example shows how to inject a route into the programmable RIB.

URL: `/restconf/config/bgp-rib:application-rib/10.25.1.9/tables/bgp-types:ipv4-address-family/bgp-types:unicast-subsequent-address-family/bgp-inet:ipv4-routes`

Method: `POST`

Content-Type: `application/xml`

Request Body:

```
<ipv4-route xmlns="urn:opendaylight:params:xml:ns:yang:bgp-inet">
    <path-id>0</path-id>
    <prefix>10.0.0.11/32</prefix>
    <attributes>
        <as-path></as-path>
        <origin>
            <value>igp</value>
        </origin>
        <local-pref>
            <pref>100</pref>
        </local-pref>
        <ipv4-next-hop>
            <global>10.11.1.1</global>
        </ipv4-next-hop>
    </attributes>
</ipv4-route>
```

Now the injected route appears in *Application Peer's* RIBs and in local speaker's Loc-RIB:

URL: `/restconf/operational/bgp-rib:bgp-rib/rib/bgp-example/peer/bgp:%2F%2F10.25.1.9`

Method: `GET`

Response Body:

```
<peer xmlns="urn:opendaylight:params:xml:ns:yang:bgp-rib">
    <peer-id>bgp://10.25.1.9</peer-id>
    <peer-role>internal</peer-role>
```

(continues on next page)

(continued from previous page)

```
    <adj-rib-in>
        <tables>
            <afi xmlns:x="urn:opendaylight:params:xml:ns:yang:bgp-types">x:ipv4-address-family</afi>
            <safi xmlns:x="urn:opendaylight:params:xml:ns:yang:bgp-types">x:unicast-subsequent-address-family</safi>
            <ipv4-routes xmlns="urn:opendaylight:params:xml:ns:yang:bgp-inet">
                <ipv4-route>
                    <path-id>0</path-id>
                    <prefix>10.0.0.11/32</prefix>
                    <attributes>
                        <as-path></as-path>
                        <origin>
                            <value>igp</value>
                        </origin>
                        <local-pref>
                            <pref>100</pref>
                        </local-pref>
                        <ipv4-next-hop>
                            <global>10.11.1.1</global>
                        </ipv4-next-hop>
                    </attributes>
                </ipv4-route>
            </ipv4-routes>
            <attributes>
                <uptodate>false</uptodate>
            </attributes>
        </tables>
    </adj-rib-in>
    <effective-rib-in>
        <tables>
            <afi xmlns:x="urn:opendaylight:params:xml:ns:yang:bgp-types">x:ipv4-address-family</afi>
            <safi xmlns:x="urn:opendaylight:params:xml:ns:yang:bgp-types">x:unicast-subsequent-address-family</safi>
            <ipv4-routes xmlns="urn:opendaylight:params:xml:ns:yang:bgp-inet">
                <ipv4-route>
                    <path-id>0</path-id>
                    <prefix>10.0.0.11/32</prefix>
                    <attributes>
                        <as-path></as-path>
                        <origin>
                            <value>igp</value>
                        </origin>
                        <local-pref>
                            <pref>100</pref>
                        </local-pref>
                        <ipv4-next-hop>
                            <global>10.11.1.1</global>
                        </ipv4-next-hop>
                    </attributes>
                </ipv4-route>
            </ipv4-routes>
            <attributes></attributes>
        </tables>
    </effective-rib-in>
</peer>
```

@line 9: Injected route is present in *Application Peer's* Adj-RIB-In and Effective-RIB-In.

URL: `/restconf/operational/bgp-rib:bgp-rib/rib/bgp-example/loc-rib/tables/bgp-types:ipv4-address-family/bgp-types:unicast-subsequent-address-family/ipv4-routes`

Method: `GET`

Response Body:

```
<ipv4-routes xmlns="urn:opendaylight:params:xml:ns:yang:bgp-inet">
    <ipv4-route>
        <path-id>0</path-id>
        <prefix>10.0.0.10/32</prefix>
        <attributes>
            <as-path></as-path>
            <origin>
                <value>igp</value>
            </origin>
            <local-pref>
                <pref>100</pref>
            </local-pref>
            <ipv4-next-hop>
                <global>10.11.1.1</global>
            </ipv4-next-hop>
        </attributes>
    </ipv4-route>
    <ipv4-route>
        <path-id>0</path-id>
        <prefix>10.0.0.10/32</prefix>
        <attributes>
            <as-path></as-path>
            <origin>
                <value>igp</value>
            </origin>
            <local-pref>
                <pref>100</pref>
            </local-pref>
            <ipv4-next-hop>
                <global>10.10.1.1</global>
            </ipv4-next-hop>
        </attributes>
    </ipv4-route>
</ipv4-routes>
```

@line 2: The injected route is now present in Loc-RIB along with a route (destination *10.0.0.10/32*) advertised by remote peer.

This route is also advertised to the remote peer (*192.0.2.1*), hence route appears in its Adj-RIB-Out:

URL: `/restconf/operational/bgp-rib:bgp-rib/rib/bgp-example/peer/bgp:%2F%2F192.0.2.1/adj-rib-out/tables/bgp-types:ipv4-address-family/bgp-types:unicast-subsequent-address-family/bgp-inet:ipv4-routes`

Method: `GET`

Response Body:

```
<ipv4-route xmlns="urn:opendaylight:params:xml:ns:yang:bgp-inet">
    <path-id>0</path-id>
    <prefix>10.0.0.11/32</prefix>
    <attributes>
        <as-path></as-path>
        <origin>
            <value>igp</value>
        </origin>
        <local-pref>
            <pref>100</pref>
        </local-pref>
        <ipv4-next-hop>
            <global>10.11.1.1</global>
        </ipv4-next-hop>
    </attributes>
</ipv4-route>
```

The injected route can be modified (i.e. different path attribute):

URL: `/restconf/config/bgp-rib:application-rib/10.25.1.9/tables/bgp-types:ipv4-address-family/bgp-types:unicast-subsequent-address-family/bgp-inet:ipv4-routes/ipv4-route/10.0.0.11%2F32/0`

Method: `PUT`

Content-Type: `application/xml`

Request Body:

```
<ipv4-route xmlns="urn:opendaylight:params:xml:ns:yang:bgp-inet">
    <path-id>0</path-id>
    <prefix>10.0.0.11/32</prefix>
    <attributes>
        <as-path></as-path>
        <origin>
            <value>igp</value>
        </origin>
        <local-pref>
            <pref>50</pref>
        </local-pref>
        <ipv4-next-hop>
            <global>10.11.1.2</global>
        </ipv4-next-hop>
    </attributes>
</ipv4-route>
```

The route can be removed from programmable RIB in a following way:

URL: `/restconf/config/bgp-rib:application-rib/10.25.1.9/tables/bgp-types:ipv4-address-family/bgp-types:unicast-subsequent-address-family/bgp-inet:ipv4-routes/ipv4-route/10.0.0.11%2F32/0`

Method: `DELETE`

Also it is possible to remove all routes from a particular table at once:

URL: `/restconf/config/bgp-rib:application-rib/10.25.1.9/tables/bgp-types:ipv4-address-family/bgp-types:unicast-subsequent-address-family/bgp-inet:ipv4-routes/`

Method: `DELETE`

Consequently, route disappears from programmable RIB, *Application Peer's* RIBs, Loc-RIB and peer's Adj-RIB-Out (UPDATE message with prefix withdrawal is send).

Note: Routes stored in programmable RIB are persisted on OpendDaylight shutdown and restored after the re-start.

BGP pipeline

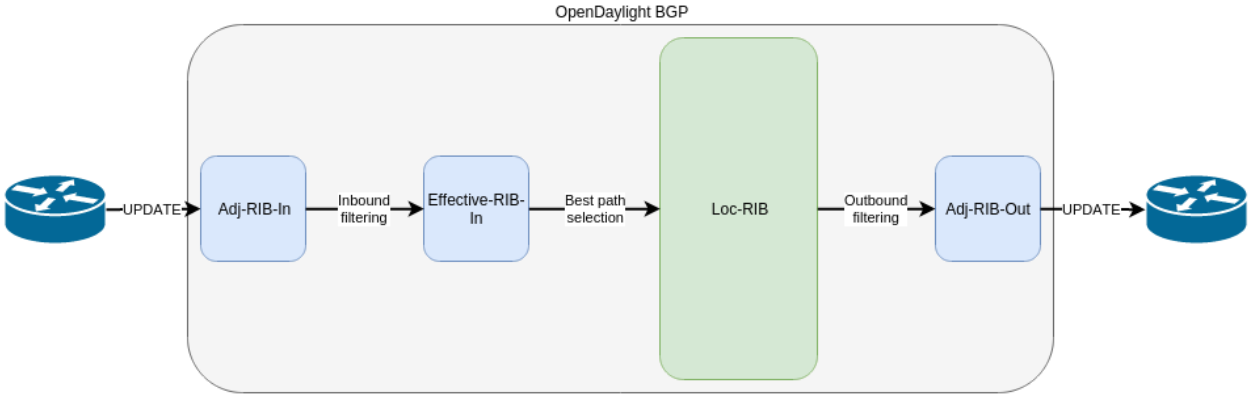

Fig. 19: BGP pipeline - routes re-advertisement.

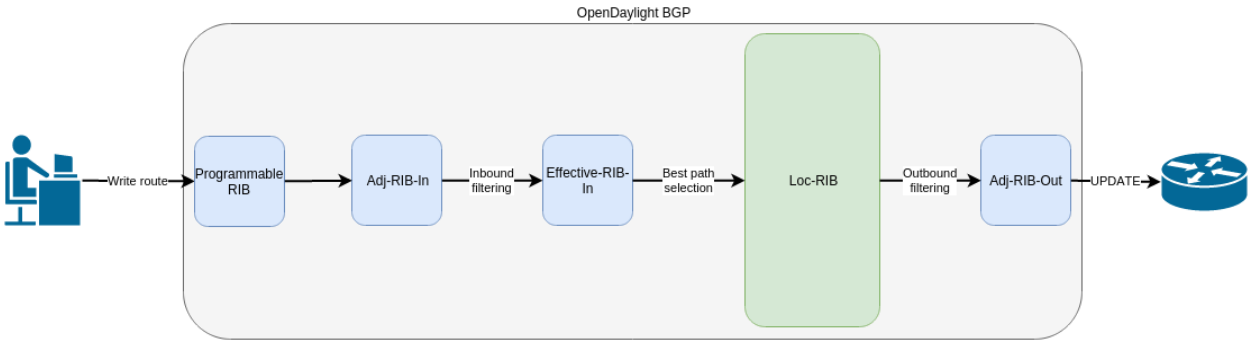

Fig. 20: BGP applcaition peer pipeline - routes injection.

References

- A Border Gateway Protocol 4 (BGP-4)
- BGP Route Reflection
- BGP Communities Attribute

- BGP Support for Four-Octet Autonomous System (AS) Number Space
- The Accumulated IGP Metric Attribute for BGP
- 4-Octet AS Specific BGP Extended Community
- BGP Link Bandwidth Extended Community
- Use of BGP for Routing in Large-Scale Data Centers

IP Unicast Family

The BGP-4 allows to carry IPv4 specific information only. The basic BGP Multiprotocol extension brings *Unicast* Subsequent Address Family (SAFI) - intended to be used for IP unicast forwarding. The combination of IPv4 and IPv6 Address Family (AF) and Unicast SAFI is essential for Internet routing. The IPv4 Unicast routes are interchangeable with BGP-4 routes, as they can carry the same type of routing information.

Contents

- *Configuration*
  - *BGP Speaker*
  - *BGP Peer*
- *IP Unicast API*
  - *IPv4 Unicast Route*
  - *IPv6 Unicast Route*
- *Usage*
  - *IPv4 Unicast*
  - *IPv6 Unicast*
- *Programming*
  - *IPv4 Unicast*
  - *IPv6 Unicast*
- *References*

Configuration

This section shows a way to enable IPv4 and IPv6 Unicast family in BGP speaker and peer configuration.

BGP Speaker

To enable IPv4 and IPv6 Unicast support in BGP plugin, first configure BGP speaker instance:

URL: `/restconf/config/openconfig-network-instance:network-instances/network-instance/global-bgp/openconfig-network-instance:protocols`

Method: `POST`

Content-Type: `application/xml`

Request Body:

```
<protocol xmlns="http://openconfig.net/yang/network-instance">
    <name>bgp-example</name>
    <identifier xmlns:x="http://openconfig.net/yang/policy-types">x:BGP</identifier>
    <bgp xmlns="urn:opendaylight:params:xml:ns:yang:bgp:openconfig-extensions">
        <global>
            <config>
                <router-id>192.0.2.2</router-id>
                <as>65000</as>
            </config>
            <afi-safis>
                <afi-safi>
                    <afi-safi-name xmlns:x="http://openconfig.net/yang/bgp-types">
↪x:IPV4-UNICAST</afi-safi-name>
                </afi-safi>
                <afi-safi>
                    <afi-safi-name xmlns:x="http://openconfig.net/yang/bgp-types">
↪x:IPV6-UNICAST</afi-safi-name>
                </afi-safi>
            </afi-safis>
        </global>
    </bgp>
</protocol>
```

BGP Peer

Here is an example for BGP peer configuration with enabled IPv4 and IPv6 Unicast family.

URL: `/restconf/config/openconfig-network-instance:network-instances/network-instance/global-bgp/openconfig-network-instance:protocols/protocol/openconfig-policy-types:BGP/bgp-example/bgp/neighbors`

Method: `POST`

Content-Type: `application/xml`

Request Body:

```
<neighbor xmlns="urn:opendaylight:params:xml:ns:yang:bgp:openconfig-extensions">
    <neighbor-address>192.0.2.1</neighbor-address>
    <afi-safis>
        <afi-safi>
            <afi-safi-name xmlns:x="http://openconfig.net/yang/bgp-types">x:IPV4-
↪UNICAST</afi-safi-name>
        </afi-safi>
        <afi-safi>
            <afi-safi-name xmlns:x="http://openconfig.net/yang/bgp-types">x:IPV6-
↪UNICAST</afi-safi-name>
        </afi-safi>
    </afi-safis>
</neighbor>
```

IP Unicast API

Following trees illustrate the BGP IP Unicast routes structures.

IPv4 Unicast Route

```
:(ipv4-routes-case)
   +--ro ipv4-routes
     +--ro ipv4-route* [prefix path-id]
        +--ro prefix        inet:ipv4-prefix
        +--ro path-id       path-id
        +--ro attributes
           +--ro origin
           |  +--ro value    bgp-t:bgp-origin
           +--ro as-path
           |  +--ro segments*
           |     +--ro as-sequence*   inet:as-number
           |     +--ro as-set*        inet:as-number
           +--ro (c-next-hop)?
           |  +--:(ipv4-next-hop-case)
           |  |  +--ro ipv4-next-hop
           |  |     +--ro global?   inet:ipv4-address
           |  +--:(ipv6-next-hop-case)
           |  |  +--ro ipv6-next-hop
           |  |     +--ro global?       inet:ipv6-address
           |  |     +--ro link-local?   inet:ipv6-address
           |  +--:(empty-next-hop-case)
           |     +--ro empty-next-hop?            empty
           +--ro multi-exit-disc
           |  +--ro med?   uint32
           +--ro local-pref
           |  +--ro pref?   uint32
           +--ro atomic-aggregate!
           +--ro aggregator
           |  +--ro as-number?         inet:as-number
           |  +--ro network-address?   inet:ipv4-address
           +--ro communities*
           |  +--ro as-number?   inet:as-number
           |  +--ro semantics?   uint16
           +--ro extended-communities*
           |  +--ro transitive?                              boolean
           |  +--ro (extended-community)?
           |     +--:(as-specific-extended-community-case)
           |     |  +--ro as-specific-extended-community
           |     |     +--ro global-administrator?   short-as-number
           |     |     +--ro local-administrator?    binary
           |     +--:(inet4-specific-extended-community-case)
           |     |  +--ro inet4-specific-extended-community
           |     |     +--ro global-administrator?   inet:ipv4-address
           |     |     +--ro local-administrator?    binary
           |     +--:(opaque-extended-community-case)
           |     |  +--ro opaque-extended-community
           |     |     +--ro value?   binary
           |     +--:(route-target-extended-community-case)
           |     |  +--ro route-target-extended-community
           |     |     +--ro global-administrator?   short-as-number
           |     |     +--ro local-administrator?    binary
           |     +--:(route-origin-extended-community-case)
           |     |  +--ro route-origin-extended-community
           |     |     +--ro global-administrator?   short-as-number
           |     |     +--ro local-administrator?    binary
```

(continues on next page)

(continued from previous page)

```
|     +--:(route-target-ipv4-case)
|     |  +--ro route-target-ipv4
|     |     +--ro global-administrator?   inet:ipv4-address
|     |     +--ro local-administrator?    uint16
|     +--:(route-origin-ipv4-case)
|     |  +--ro route-origin-ipv4
|     |     +--ro global-administrator?   inet:ipv4-address
|     |     +--ro local-administrator?    uint16
|     +--:(link-bandwidth-case)
|     |  +--ro link-bandwidth-extended-community
|     |     +--ro bandwidth    netc:bandwidth
|     +--:(as-4-generic-spec-extended-community-case)
|     |  +--ro as-4-generic-spec-extended-community
|     |     +--ro as-4-specific-common
|     |        +--ro as-number              inet:as-number
|     |        +--ro local-administrator    uint16
|     +--:(as-4-route-target-extended-community-case)
|     |  +--ro as-4-route-target-extended-community
|     |     +--ro as-4-specific-common
|     |        +--ro as-number              inet:as-number
|     |        +--ro local-administrator    uint16
|     +--:(as-4-route-origin-extended-community-case)
|     |  +--ro as-4-route-origin-extended-community
|     |     +--ro as-4-specific-common
|     |        +--ro as-number              inet:as-number
|     |        +--ro local-administrator    uint16
|     +--:(encapsulation-case)
|        +--ro encapsulation-extended-community
|           +--ro tunnel-type    encapsulation-tunnel-type
+--ro originator-id
|  +--ro originator?   inet:ipv4-address
+--ro cluster-id
|  +--ro cluster*   bgp-t:cluster-identifier
+--ro aigp
|  +--ro aigp-tlv
|     +--ro metric?   netc:accumulated-igp-metric
+--ro unrecognized-attributes* [type]
   +--ro partial       boolean
   +--ro transitive    boolean
   +--ro type          uint8
   +--ro value         binary
```

IPv6 Unicast Route

```
:(ipv6-routes-case)
   +--ro ipv6-routes
      +--ro ipv6-route* [prefix path-id]
         +--ro prefix        inet:ipv6-prefix
         +--ro path-id       path-id
         +--ro attributes
         ...
```

Usage

IPv4 Unicast

The IPv4 Unicast table in an instance of the speaker's Loc-RIB can be verified via REST:

URL: `/restconf/operational/bgp-rib:bgp-rib/rib/bgp-example/loc-rib/tables/bgp-types:ipv4-address-family/bgp-types:unicast-subsequent-address-family/ipv4-routes`

Method: `GET`

Response Body:

```
<ipv4-routes xmlns="urn:opendaylight:params:xml:ns:yang:bgp-inet">
    <ipv4-route>
        <path-id>0</path-id>
        <prefix>193.0.2.1/32</prefix>
        <attributes>
            <as-path></as-path>
            <origin>
                <value>igp</value>
            </origin>
            <local-pref>
                <pref>100</pref>
            </local-pref>
            <ipv4-next-hop>
                <global>10.0.0.1</global>
            </ipv4-next-hop>
        </attributes>
    </ipv4-route>
</ipv4-routes>
```

IPv6 Unicast

The IPv6 Unicast table in an instance of the speaker's Loc-RIB can be verified via REST:

URL: `/restconf/operational/bgp-rib:bgp-rib/rib/bgp-example/loc-rib/tables/bgp-types:ipv4-address-family/bgp-types:unicast-subsequent-address-family/ipv6-routes`

Method: `GET`

Response Body:

```
<ipv6-routes xmlns="urn:opendaylight:params:xml:ns:yang:bgp-inet">
    <ipv6-route>
        <path-id>0</path-id>
        <prefix>2a02:b80:0:1::/64</prefix>
        <attributes>
            <as-path></as-path>
            <origin>
                <value>igp</value>
            </origin>
            <local-pref>
                <pref>200</pref>
            </local-pref>
```

(continues on next page)

(continued from previous page)

```
            <ipv6-next-hop>
                <global>2a02:b80:0:2::1</global>
            </ipv6-next-hop>
        </attributes>
    </ipv6-route>
</ipv6-routes>
```

Note: IPv4/6 routes mapping to topology nodes is supported by BGP Topology Provider.

Programming

IPv4 Unicast

This examples show how to originate and remove IPv4 route via programmable RIB. Make sure the *Application Peer* is configured first.

URL: `/restconf/config/bgp-rib:application-rib/10.25.1.9/tables/bgp-types:ipv4-address-family/bgp-types:unicast-subsequent-address-family/bgp-inet:ipv4-routes`

Method: `POST`

Content-Type: `application/xml`

Request Body:

```
<ipv4-route xmlns="urn:opendaylight:params:xml:ns:yang:bgp-inet">
    <path-id>0</path-id>
    <prefix>10.0.0.11/32</prefix>
    <attributes>
        <as-path></as-path>
        <origin>
            <value>igp</value>
        </origin>
        <local-pref>
            <pref>100</pref>
        </local-pref>
        <ipv4-next-hop>
            <global>10.11.1.1</global>
        </ipv4-next-hop>
    </attributes>
</ipv4-route>
```

To remove the route added above, following request can be used:

URL: `/restconf/config/bgp-rib:application-rib/10.25.1.9/tables/bgp-types:ipv4-address-family/bgp-types:unicast-subsequent-address-family/bgp-inet:ipv4-routes/ipv4-route/10.0.0.11%2F32/0`

Method: `DELETE`

IPv6 Unicast

This examples show how to originate and remove IPv6 route via programmable RIB:

URL: `/restconf/config/bgp-rib:application-rib/10.25.1.9/tables/bgp-types:ipv6-address-family/bgp-types:unicast-subsequent-address-family/bgp-inet:ipv6-routes`

Method: `POST`

Content-Type: `application/xml`

Request Body:

```
<ipv6-route xmlns="urn:opendaylight:params:xml:ns:yang:bgp-inet">
    <prefix>2001:db8:30::3/128</prefix>
    <path-id>0</path-id>
    <attributes>
        <ipv6-next-hop>
            <global>2001:db8:1::6</global>
        </ipv6-next-hop>
        <as-path/>
        <origin>
            <value>igp</value>
        </origin>
        <local-pref>
            <pref>100</pref>
        </local-pref>
    </attributes>
</ipv6-route>
```

To remove the route added above, following request can be used:

URL: `/restconf/config/bgp-rib:application-rib/10.25.1.9/tables/bgp-types:ipv6-address-family/bgp-types:unicast-subsequent-address-family/bgp-inet:ipv6-routes/ipv6-route/2001:db8:30::3%2F128/0`

Method: `DELETE`

References

- Multiprotocol Extensions for BGP-4

IP Labeled Unicast Family

The BGP Labeled Unicast (BGP-LU) Multiprotocol extension is used to distribute a MPLS label that is mapped to a particular route. It can be used to advertise a MPLS transport path between IGP regions and Autonomous Systems. Also, BGP-LU can help to solve the Inter-domain traffic-engineering problem and can be deployed in large-scale data centers along with MPLS and Spring. In addition, IPv6 Labeled Unicast can be used to interconnect IPv6 islands over IPv4/MPLS networks using 6PE.

Contents

- *Configuration*
  - *BGP Speaker*
  - *BGP Peer*
- *IP Labeled Unicast API*
  - *IPv4 Labeled Unicast Route*
  - *IPv6 Labeled Unicast Route*
- *Usage*
- *Programming*
  - *IPv4 Labeled*
  - *IPv6 Labeled*
- *References*

Configuration

This section shows a way to enable IPv4 and IPv6 Labeled Unicast family in BGP speaker and peer configuration.

BGP Speaker

To enable IPv4 and IPv6 Labeled Unicast support in BGP plugin, first configure BGP speaker instance:

URL: `/restconf/config/openconfig-network-instance:network-instances/network-instance/global-bgp/openconfig-network-instance:protocols`

Method: `POST`

Content-Type: `application/xml`

Request Body:

```
<protocol xmlns="http://openconfig.net/yang/network-instance">
    <name>bgp-example</name>
    <identifier xmlns:x="http://openconfig.net/yang/policy-types">x:BGP</identifier>
    <bgp xmlns="urn:opendaylight:params:xml:ns:yang:bgp:openconfig-extensions">
        <global>
            <config>
                <router-id>192.0.2.2</router-id>
                <as>65000</as>
            </config>
            <afi-safis>
                <afi-safi>
                    <afi-safi-name xmlns:x="http://openconfig.net/yang/bgp-types">
↪x:IPV4-LABELLED-UNICAST</afi-safi-name>
                </afi-safi>
                <afi-safi>
                    <afi-safi-name xmlns:x="http://openconfig.net/yang/bgp-types">
↪x:IPV6-LABELLED-UNICAST</afi-safi-name>
                </afi-safi>
            </afi-safis>
        </global>
```

(continues on next page)

(continued from previous page)

```
    </bgp>
</protocol>
```

BGP Peer

Here is an example for BGP peer configuration with enabled IPv4 and IPv6 Labeled Unicast family.

URL: `/restconf/config/openconfig-network-instance:network-instances/network-instance/global-bgp/openconfig-network-instance:protocols/protocol/openconfig-policy-types:BGP/bgp-example/bgp/neighbors`

Method: `POST`

Content-Type: `application/xml`

Request Body:

```
<neighbor xmlns="urn:opendaylight:params:xml:ns:yang:bgp:openconfig-extensions">
    <neighbor-address>192.0.2.1</neighbor-address>
    <afi-safis>
        <afi-safi>
            <afi-safi-name xmlns:x="http://openconfig.net/yang/bgp-types">x:IPV4-
↪LABELLED-UNICAST</afi-safi-name>
        </afi-safi>
        <afi-safi>
            <afi-safi-name xmlns:x="http://openconfig.net/yang/bgp-types">x:IPV6-
↪LABELLED-UNICAST</afi-safi-name>
        </afi-safi>
    </afi-safis>
</neighbor>
```

IP Labeled Unicast API

Following trees illustrate the BGP IP Labeled Unicast routes structures.

IPv4 Labeled Unicast Route

```
:(labeled-unicast-routes-case)
  +--ro labeled-unicast-routes
     +--ro labeled-unicast-route* [route-key path-id]
        +--ro route-key      string
        +--ro label-stack*
        |  +--ro label-value?   netc:mpls-label
        +--ro prefix?        inet:ip-prefix
        +--ro path-id        path-id
        +--ro attributes
        ...
```

IPv6 Labeled Unicast Route

```
:(labeled-unicast-ipv6-routes-case)
   +--ro labeled-unicast-ipv6-routes
      +--ro labeled-unicast-route* [route-key path-id]
         +--ro route-key      string
         +--ro label-stack*
         |  +--ro label-value?   netc:mpls-label
         +--ro prefix?        inet:ip-prefix
         +--ro path-id        path-id
         +--ro attributes
         ...
```

Usage

The IPv4 Labeled Unicast table in an instance of the speaker's Loc-RIB can be verified via REST:

URL: `/restconf/operational/bgp-rib:bgp-rib/rib/bgp-example/loc-rib/tables/bgp-types:ipv4-address-family/bgp-labeled-unicast:labeled-unicast-subsequent-address-family-bgp-labeled-unicast:labeled-unicast-routes`

Method: `GET`

Response Body:

```
<labeled-unicast-routes xmlns="urn:opendaylight:params:xml:ns:yang:bgp-labeled-unicast
↪">
    <labeled-unicast-route>
        <path-id>0</path-id>
        <route-key>MAA+gRQAAA==</route-key>
        <attributes>
            <local-pref>
                <pref>100</pref>
            </local-pref>
            <ipv4-next-hop>
                <global>200.10.0.101</global>
            </ipv4-next-hop>
            <as-path></as-path>
            <origin>
                <value>igp</value>
            </origin>
        </attributes>
        <label-stack>
            <label-value>1000</label-value>
        </label-stack>
        <prefix>20.0.0.0/24</prefix>
    </labeled-unicast-route>
</labeled-unicast-routes>
```

Programming

IPv4 Labeled

This examples show how to originate and remove IPv4 labeled route via programmable RIB. Make sure the *Application Peer* is configured first.

URL: /restconf/config/bgp-rib:application-rib/10.25.1.9/tables/bgp-types:ipv4-address-family/bgp-labeled-unicast:labeled-unicast-subsequent-address-family/bgp-labeled-unicast:labeled-unicast-routes

Method: POST

Content-Type: application/xml

Request Body:

```
<labeled-unicast-route xmlns="urn:opendaylight:params:xml:ns:yang:bgp-labeled-unicast
↪">
    <route-key>label1</route-key>
    <prefix>1.1.1.1/32</prefix>
    <path-id>0</path-id>
    <label-stack>
        <label-value>800322</label-value>
    </label-stack>
    <attributes>
        <ipv4-next-hop>
            <global>199.20.160.41</global>
        </ipv4-next-hop>
        <origin>
            <value>igp</value>
        </origin>
        <as-path/>
        <local-pref>
            <pref>100</pref>
        </local-pref>
    </attributes>
</labeled-unicast-route>
```

In addition, BGP-LU Spring extension allows to attach BGP Prefix SID attribute to the route, in order to signal the BGP-Prefix-SID, where the SR is applied to MPLS dataplane.

```
<bgp-prefix-sid>
    <bgp-prefix-sid-tlvs>
        <label-index-tlv xmlns="urn:opendaylight:params:xml:ns:yang:bgp-labeled-
↪unicast">322</label-index-tlv>
    </bgp-prefix-sid-tlvs>
    <bgp-prefix-sid-tlvs>
        <srgb-value xmlns="urn:opendaylight:params:xml:ns:yang:bgp-labeled-unicast">
            <base>800000</base>
            <range>4095</range>
        </srgb-value>
    </bgp-prefix-sid-tlvs>
</bgp-prefix-sid>
```

To remove the route added above, following request can be used:

URL: /restconf/config/bgp-rib:application-rib/10.25.1.9/tables/bgp-types:ipv4-address-family/bgp-labeled-unicast:labeled-unicast-subsequent-address-family/bgp-labeled-unicast:labeled-unicast-routes/bgp-labeled-unicast:labeled-unicast-route/label1/0

Method: DELETE

IPv6 Labeled

This examples show how to originate and remove IPv6 labeled route via programmable RIB.

URL: `/restconf/config/bgp-rib:application-rib/10.25.1.9/tables/bgp-types:ipv4-address-family/bgp-labeled-unicast:labeled-unicast-subsequent-address-family/bgp-labeled-unicast:labeled-unicast-ipv6-routes`

Method: `POST`

Content-Type: `application/xml`

Request Body:

```
<labeled-unicast-route xmlns="urn:opendaylight:params:xml:ns:yang:bgp-labeled-unicast
↪">
    <route-key>label1</route-key>
    <prefix>2001:db8:30::3/128</prefix>
    <path-id>0</path-id>
    <label-stack>
        <label-value>123</label-value>
    </label-stack>
    <attributes>
        <ipv6-next-hop>
            <global>2003:4:5:6::7</global>
        </ipv6-next-hop>
        <origin>
            <value>igp</value>
        </origin>
        <as-path/>
        <local-pref>
            <pref>100</pref>
        </local-pref>
    </attributes>
</labeled-unicast-route>
```

To remove the route added above, following request can be used:

URL: `/restconf/config/bgp-rib:application-rib/10.25.1.9/tables/bgp-types:ipv4-address-family/bgp-labeled-unicast:labeled-unicast-subsequent-address-family/bgp-labeled-unicast:labeled-unicast-ipv6-routes/bgp-labeled-unicast:labeled-unicast-route/label1/0`

Method: `DELETE`

References

- Carrying Label Information in BGP-4
- Segment Routing Prefix SID extensions for BGP
- Connecting IPv6 Islands over IPv4 MPLS Using IPv6 Provider Edge Routers (6PE)
- BGP-Prefix Segment in large-scale data centers
- Egress Peer Engineering using BGP-LU

IP L3VPN Family

The BGP/MPLS IP Virtual Private Networks (BGP L3VPN) Multiprotocol extension can be used to exchange particular VPN (customer) routes among the provider's routers attached to that VPN. Also, routes are distributed to specific VPN remote sites.

Contents

- *Configuration*
  - *BGP Speaker*
  - *BGP Peer*
- *IP L3VPN API*
  - *IPv4 L3VPN Route*
  - *IPv6 L3VPN Route*
- *Usage*
  - *IPv4 L3VPN*
  - *IPv6 L3VPN*
- *Programming*
- *References*

Configuration

This section shows a way to enable IPv4 and IPv6 L3VPN family in BGP speaker and peer configuration.

BGP Speaker

To enable IPv4 and IPv6 L3VPN support in BGP plugin, first configure BGP speaker instance:

URL: `/restconf/config/openconfig-network-instance:network-instances/network-instance/global-bgp/openconfig-network-instance:protocols`

Method: `POST`

Content-Type: `application/xml`

Request Body:

```
<protocol xmlns="http://openconfig.net/yang/network-instance">
    <name>bgp-example</name>
    <identifier xmlns:x="http://openconfig.net/yang/policy-types">x:BGP</identifier>
    <bgp xmlns="urn:opendaylight:params:xml:ns:yang:bgp:openconfig-extensions">
        <global>
            <config>
                <router-id>192.0.2.2</router-id>
                <as>65000</as>
            </config>
            <afi-safis>
```

(continues on next page)

(continued from previous page)

```
                <afi-safi>
                    <afi-safi-name xmlns:x="http://openconfig.net/yang/bgp-types">
↪x:L3VPN-IPV4-UNICAST</afi-safi-name>
                </afi-safi>
                <afi-safi>
                    <afi-safi-name xmlns:x="http://openconfig.net/yang/bgp-types">
↪x:L3VPN-IPV6-UNICAST</afi-safi-name>
                </afi-safi>
            </afi-safis>
        </global>
    </bgp>
</protocol>
```

BGP Peer

Here is an example for BGP peer configuration with enabled IPv4 and IPv6 L3VPN family.

URL: `/restconf/config/openconfig-network-instance:network-instances/network-instance/global-bgp/openconfig-network-instance:protocols/protocol/openconfig-policy-types:BGP/bgp-example/bgp/neighbors`

Method: `POST`

Content-Type: `application/xml`

Request Body:

```
<neighbor xmlns="urn:opendaylight:params:xml:ns:yang:bgp:openconfig-extensions">
    <neighbor-address>192.0.2.1</neighbor-address>
    <afi-safis>
        <afi-safi>
            <afi-safi-name xmlns:x="http://openconfig.net/yang/bgp-types">x:L3VPN-
↪IPV4-UNICAST</afi-safi-name>
        </afi-safi>
        <afi-safi>
            <afi-safi-name xmlns:x="http://openconfig.net/yang/bgp-types">x:L3VPN-
↪IPV6-UNICAST</afi-safi-name>
        </afi-safi>
    </afi-safis>
</neighbor>
```

IP L3VPN API

Following trees illustrate the BGP IP L3VPN routes structures.

IPv4 L3VPN Route

```
:(vpn-ipv4-routes-case)
   +--ro vpn-ipv4-routes
      +--ro vpn-route* [route-key]
         +--ro route-key              string
         +--ro label-stack*
```

(continues on next page)

(continued from previous page)

```
         |  +--ro label-value?   netc:mpls-label
         +--ro prefix?               inet:ip-prefix
         +--ro path-id?              path-id
         +--ro route-distinguisher?  bgp-t:route-distinguisher
         +--ro attributes
         ...
```

IPv6 L3VPN Route

```
:(vpn-ipv6-routes-case)
   +--ro vpn-ipv6-routes
      +--ro vpn-route* [route-key]
         +--ro route-key              string
         +--ro label-stack*
         |  +--ro label-value?   netc:mpls-label
         +--ro prefix?               inet:ip-prefix
         +--ro path-id?              path-id
         +--ro route-distinguisher?  bgp-t:route-distinguisher
         +--ro attributes
         ...
```

Usage

IPv4 L3VPN

The IPv4 L3VPN table in an instance of the speaker's Loc-RIB can be verified via REST:

URL: `/restconf/operational/bgp-rib:bgp-rib/rib/bgp-example/loc-rib/tables/bgp-types:ipv4-address-family/bgp-types:mpls-labeled-vpn-subsequent-address-family/bgp-vpn-ipv4:vpn-ipv4-routes`

Method: `GET`

Response Body:

```
<vpn-ipv4-routes xmlns="urn:opendaylight:params:xml:ns:yang:bgp-vpn-ipv4">
    <vpn-route>
        <route-key>cAXdYQABrBAALABlCgIi</route-key>
        <label-stack>
            <label-value>24022</label-value>
        </label-stack>
        <attributes>
            <extended-communities>
                <transitive>true</transitive>
                <route-target-extended-community>
                    <global-administrator>65000</global-administrator>
                    <local-administrator>AAAAZQ==</local-administrator>
                </route-target-extended-community>
            </extended-communities>
            <origin>
                <value>igp</value>
            </origin>
            <as-path></as-path>
```

(continues on next page)

(continued from previous page)

```
            <local-pref>
                <pref>100</pref>
            </local-pref>
            <ipv4-next-hop>
                <global>127.16.0.44</global>
            </ipv4-next-hop>
        </attributes>
        <route-distinguisher>172.16.0.44:101</route-distinguisher>
        <prefix>10.2.34.0/24</prefix>
    </vpn-route>
</vpn-ipv4-routes>
```

IPv6 L3VPN

The IPv6 L3VPN table in an instance of the speaker's Loc-RIB can be verified via REST:

URL: `/restconf/operational/bgp-rib:bgp-rib/rib/bgp-example/loc-rib/tables/bgp-types:ipv6-address-family/bgp-types:mpls-labeled-vpn-subsequent-address-family/bgp-vpn-ipv6:vpn-ipv6-routes`

Method: `GET`

Response Body:

```
<vpn-ipv6-routes xmlns="urn:opendaylight:params:xml:ns:yang:bgp-vpn-ipv6">
    <vpn-route>
        <route-key>mAXdcQABrBAALABlKgILgAAAAAE=</route-key>
        <label-stack>
            <label-value>24023</label-value>
        </label-stack>
        <attributes>
            <local-pref>
                <pref>100</pref>
            </local-pref>
            <extended-communities>
                <route-target-extended-community>
                    <global-administrator>65000</global-administrator>
                    <local-administrator>AAAAZQ==</local-administrator>
                </route-target-extended-community>
                <transitive>true</transitive>
            </extended-communities>
            <ipv6-next-hop>
                <global>2a02:b80:0:2::1</global>
            </ipv6-next-hop>
            <origin>
                <value>igp</value>
            </origin>
            <as-path></as-path>
        </attributes>
        <route-distinguisher>172.16.0.44:101</route-distinguisher>
        <prefix>2a02:b80:0:1::/64</prefix>
    </vpn-route>
</vpn-ipv6-routes>
```

Programming

This examples show how to originate and remove IPv4 L3VPN route via programmable RIB. Make sure the *Application Peer* is configured first.

URL: `/restconf/config/bgp-rib:application-rib/10.25.1.9/tables/bgp-types:ipv4-address-family/bgp-types:mpls-labeled-vpn-subsequent-address-family/bgp-vpn-ipv4:vpn-ipv4-routes`

Method: `POST`

Content-Type: `application/xml`

Request Body:

```
<vpn-route xmlns="urn:opendaylight:params:xml:ns:yang:bgp-vpn-ipv4">
    <route-key>vpn1</route-key>
    <label-stack>
        <label-value>123</label-value>
    </label-stack>
    <route-distinguisher>429496729:1</route-distinguisher>
    <prefix>2.2.2.2/32</prefix>
    <attributes>
        <ipv4-next-hop>
            <global>199.20.166.41</global>
        </ipv4-next-hop>
        <as-path/>
        <origin>
            <value>igp</value>
        </origin>
        <extended-communities>
            <route-target-extended-community>
                <global-administrator>65000</global-administrator>
                <local-administrator>AAAAZQ==</local-administrator>
            </route-target-extended-community>
            <transitive>true</transitive>
        </extended-communities>
    </attributes>
</vpn-route>
```

To remove the route added above, following request can be used:

URL: `/restconf/config/bgp-rib:application-rib/10.25.1.9/tables/bgp-types:ipv4-address-family/bgp-types:mpls-labeled-vpn-subsequent-address-family/bgp-vpn-ipv4:vpn-ipv4-routes/vpn-route/vpn1`

Method: `DELETE`

References

- BGP/MPLS IP Virtual Private Networks (VPNs)
- BGP-MPLS IP Virtual Private Network (VPN) Extension for IPv6 VPN
- BGP/MPLS VPN Virtual PE

Link-State Family

The BGP Link-State (BGP-LS) Multiprotocol extension allows to distribute Link-State and Traffic Engineering (TE) information. This information is typically distributed by IGP routing protocols with in the network, limiting LSDB or TED visibility to the IGP area. The BGP-LS-enabled routers are capable to collect such information from networks (multiple IGP areas, inter-AS) and share with external components (i.e. OpenDaylight BGP). The information is applicable in ALTO servers and PCEs, as both need to gather information about topologies. In addition, link-state information is extended to carry segment information (Spring).

Contents

- *Configuration*
  - *BGP Speaker*
  - *Linkstate path attribute*
  - *BGP Peer*
- *Link-State Route API*
- *Usage*
- *References*

Configuration

This section shows a way to enable IPv4 and IPv6 Labeled Unicast family in BGP speaker and peer configuration.

BGP Speaker

To enable BGP-LS support in BGP plugin, first configure BGP speaker instance:

URL: `/restconf/config/openconfig-network-instance:network-instances/network-instance/global-bgp/openconfig-network-instance:protocols`

Method: `POST`

Content-Type: `application/xml`

Request Body:

```
<protocol xmlns="http://openconfig.net/yang/network-instance">
    <name>bgp-example</name>
    <identifier xmlns:x="http://openconfig.net/yang/policy-types">x:BGP</identifier>
    <bgp xmlns="urn:opendaylight:params:xml:ns:yang:bgp:openconfig-extensions">
        <global>
            <config>
                <router-id>192.0.2.2</router-id>
                <as>65000</as>
            </config>
            <afi-safis>
                <afi-safi>
                    <afi-safi-name>LINKSTATE</afi-safi-name>
                </afi-safi>
            </afi-safis>
```

(continues on next page)

(continued from previous page)

```
        </global>
    </bgp>
</protocol>
```

Linkstate path attribute

IANA allocation for BGP-LS path attribute is TYPE 29. Some older BGP-LS implementations might still require earliest asigned allocation TYPE 99. To use TYPE = 99, you need to set value bellow to false.

URL: `/restconf/config/bgp-linkstate-app-config:bgp-linkstate-app-config`

Method: `PUT`

Content-Type: `application/xml`

Request Body:

```
<bgp-linkstate-app-config xmlns=
↪"urn:opendaylight:params:xml:ns:yang:controller:bgp:linkstate-app-config">
    <iana-linkstate-attribute-type>false</iana-linkstate-attribute-type>
</bgp-linkstate-app-config>
```

BGP Peer

Here is an example for BGP peer configuration with enabled BGP-LS family.

URL: `/restconf/config/openconfig-network-instance:network-instances/network-instance/global-bgp/openconfig-network-instance:protocols/protocol/openconfig-policy-types:BGP/bgp-example/bgp/neighbors`

Method: `POST`

Content-Type: `application/xml`

Request Body:

```
<neighbor xmlns="urn:opendaylight:params:xml:ns:yang:bgp:openconfig-extensions">
    <neighbor-address>192.0.2.1</neighbor-address>
    <afi-safis>
        <afi-safi>
            <afi-safi-name>LINKSTATE</afi-safi-name>
        </afi-safi>
    </afi-safis>
</neighbor>
```

Link-State Route API

Following tree illustrate the BGP Link-State route structure.

```
:(linkstate-routes-case)
   +--ro linkstate-routes
      +--ro linkstate-route* [route-key]
         +--ro route-key                        binary
         +--ro protocol-id                      protocol-id
```

(continues on next page)

(continued from previous page)

```
         +--ro identifier                      identifier
         +--ro (object-type)?
         |  +--:(node-case)
         |  |  +--ro node-descriptors
         |  |     +--ro as-number?         inet:as-number
         |  |     +--ro area-id?           area-identifier
         |  |     +--ro domain-id?         domain-identifier
         |  |     +--ro (c-router-identifier)?
         |  |        +--:(isis-node-case)
         |  |        |  +--ro isis-node
         |  |        |     +--ro iso-system-id    netc:iso-system-identifier
         |  |        +--:(isis-pseudonode-case)
         |  |        |  +--ro isis-pseudonode
         |  |        |     +--ro is-is-router-identifier
         |  |        |     |  +--ro iso-system-id    netc:iso-system-identifier
         |  |        |     +--ro psn                        uint8
         |  |        +--:(ospf-node-case)
         |  |        |  +--ro ospf-node
         |  |        |     +--ro ospf-router-id    uint32
         |  |        +--:(ospf-pseudonode-case)
         |  |           +--ro ospf-pseudonode
         |  |              +--ro ospf-router-id    uint32
         |  |              +--ro lan-interface     ospf-interface-identifier
         |  +--:(link-case)
         |  |  +--ro local-node-descriptors
         |  |  |  +--ro as-number?         inet:as-number
         |  |  |  +--ro area-id?           area-identifier
         |  |  |  +--ro domain-id?         domain-identifier
         |  |  |  +--ro (c-router-identifier)?
         |  |  |  |  +--:(isis-node-case)
         |  |  |  |  |  +--ro isis-node
         |  |  |  |  |     +--ro iso-system-id    netc:iso-system-identifier
         |  |  |  |  +--:(isis-pseudonode-case)
         |  |  |  |  |  +--ro isis-pseudonode
         |  |  |  |  |     +--ro is-is-router-identifier
         |  |  |  |  |     |  +--ro iso-system-id    netc:iso-system-identifier
         |  |  |  |  |     +--ro psn                        uint8
         |  |  |  |  +--:(ospf-node-case)
         |  |  |  |  |  +--ro ospf-node
         |  |  |  |  |     +--ro ospf-router-id    uint32
         |  |  |  |  +--:(ospf-pseudonode-case)
         |  |  |  |     +--ro ospf-pseudonode
         |  |  |  |        +--ro ospf-router-id    uint32
         |  |  |  |        +--ro lan-interface     ospf-interface-identifier
         |  |  |  +--ro bgp-router-id?     inet:ipv4-address
         |  |  |  +--ro member-asn?        inet:as-number
         |  |  +--ro remote-node-descriptors
         |  |  |  +--ro as-number?         inet:as-number
         |  |  |  +--ro area-id?           area-identifier
         |  |  |  +--ro domain-id?         domain-identifier
         |  |  |  +--ro (c-router-identifier)?
         |  |  |  |  +--:(isis-node-case)
         |  |  |  |  |  +--ro isis-node
         |  |  |  |  |     +--ro iso-system-id    netc:iso-system-identifier
         |  |  |  |  +--:(isis-pseudonode-case)
         |  |  |  |  |  +--ro isis-pseudonode
         |  |  |  |  |     +--ro is-is-router-identifier
```

(continues on next page)

(continued from previous page)

```
|  |  |  |  |     |  +--ro iso-system-id    netc:iso-system-identifier
|  |  |  |  |     +--ro psn                        uint8
|  |  |  |  +--:(ospf-node-case)
|  |  |  |  |  +--ro ospf-node
|  |  |  |  |     +--ro ospf-router-id    uint32
|  |  |  |  +--:(ospf-pseudonode-case)
|  |  |  |     +--ro ospf-pseudonode
|  |  |  |        +--ro ospf-router-id    uint32
|  |  |  |        +--ro lan-interface     ospf-interface-identifier
|  |  |  +--ro bgp-router-id?     inet:ipv4-address
|  |  |  +--ro member-asn?        inet:as-number
|  |  +--ro link-descriptors
|  |     +--ro link-local-identifier?    uint32
|  |     +--ro link-remote-identifier?   uint32
|  |     +--ro ipv4-interface-address?   ipv4-interface-identifier
|  |     +--ro ipv6-interface-address?   ipv6-interface-identifier
|  |     +--ro ipv4-neighbor-address?    ipv4-interface-identifier
|  |     +--ro ipv6-neighbor-address?    ipv6-interface-identifier
|  |     +--ro multi-topology-id?        topology-identifier
|  +--:(prefix-case)
|  |  +--ro advertising-node-descriptors
|  |  |  +--ro as-number?         inet:as-number
|  |  |  +--ro area-id?           area-identifier
|  |  |  +--ro domain-id?         domain-identifier
|  |  |  +--ro (c-router-identifier)?
|  |  |     +--:(isis-node-case)
|  |  |     |  +--ro isis-node
|  |  |     |     +--ro iso-system-id    netc:iso-system-identifier
|  |  |     +--:(isis-pseudonode-case)
|  |  |     |  +--ro isis-pseudonode
|  |  |     |     +--ro is-is-router-identifier
|  |  |     |     |  +--ro iso-system-id    netc:iso-system-identifier
|  |  |     |     +--ro psn                        uint8
|  |  |     +--:(ospf-node-case)
|  |  |     |  +--ro ospf-node
|  |  |     |     +--ro ospf-router-id    uint32
|  |  |     +--:(ospf-pseudonode-case)
|  |  |        +--ro ospf-pseudonode
|  |  |           +--ro ospf-router-id    uint32
|  |  |           +--ro lan-interface     ospf-interface-identifier
|  |  +--ro prefix-descriptors
|  |     +--ro multi-topology-id?             topology-identifier
|  |     +--ro ospf-route-type?               ospf-route-type
|  |     +--ro ip-reachability-information?   inet:ip-prefix
|  +--:(te-lsp-case)
|     +--ro (address-family)?
|     |  +--:(ipv4-case)
|     |  |  +--ro ipv4-tunnel-sender-address      inet:ipv4-address
|     |  |  +--ro ipv4-tunnel-endpoint-address    inet:ipv4-address
|     |  +--:(ipv6-case)
|     |     +--ro ipv6-tunnel-sender-address      inet:ipv6-address
|     |     +--ro ipv6-tunnel-endpoint-address    inet:ipv6-address
|     +--ro tunnel-id?                       rsvp:tunnel-id
|     +--ro lsp-id?                          rsvp:lsp-id
+--ro attributes
   +--ro (link-state-attribute)?
      +--:(node-attributes-case)
```

(continues on next page)

(continued from previous page)

```
|  +--ro node-attributes
|     +--ro topology-identifier*   topology-identifier
|     +--ro node-flags?            node-flag-bits
|     +--ro isis-area-id*          isis-area-identifier
|     +--ro dynamic-hostname?      string
|     +--ro ipv4-router-id?        ipv4-router-identifier
|     +--ro ipv6-router-id?        ipv6-router-identifier
|     +--ro sr-capabilities
|     |  +--ro mpls-ipv4?      boolean
|     |  +--ro mpls-ipv6?      boolean
|     |  +--ro sr-ipv6?        boolean
|     |  +--ro range-size?     uint32
|     |  +--ro (sid-label-index)?
|     |     +--:(local-label-case)
|     |     |  +--ro local-label?    netc:mpls-label
|     |     +--:(ipv6-address-case)
|     |     |  +--ro ipv6-address?   inet:ipv6-address
|     |     +--:(sid-case)
|     |        +--ro sid?            uint32
|     +--ro sr-algorithm
|        +--ro algorithms*   algorithm
+--:(link-attributes-case)
|  +--ro link-attributes
|     +--ro local-ipv4-router-id?       ipv4-router-identifier
|     +--ro local-ipv6-router-id?       ipv6-router-identifier
|     +--ro remote-ipv4-router-id?      ipv4-router-identifier
|     +--ro remote-ipv6-router-id?      ipv6-router-identifier
|     +--ro mpls-protocol?              mpls-protocol-mask
|     +--ro te-metric?                  netc:te-metric
|     +--ro metric?                     netc:metric
|     +--ro shared-risk-link-groups*    rsvp:srlg-id
|     +--ro link-name?                  string
|     +--ro max-link-bandwidth?         netc:bandwidth
|     +--ro max-reservable-bandwidth?   netc:bandwidth
|     +--ro unreserved-bandwidth* [priority]
|     |  +--ro priority     uint8
|     |  +--ro bandwidth?   netc:bandwidth
|     +--ro link-protection?            link-protection-type
|     +--ro admin-group?                administrative-group
|     +--ro sr-adj-ids*
|     |  +--ro (flags)?
|     |  |  +--:(ospf-adj-flags-case)
|     |  |  |  +--ro backup?           boolean
|     |  |  |  +--ro set?              boolean
|     |  |  +--:(isis-adj-flags-case)
|     |  |     +--ro backup?           boolean
|     |  |     +--ro set?              boolean
|     |  |     +--ro address-family?   boolean
|     |  +--ro weight?           weight
|     |  +--ro (sid-label-index)?
|     |     +--:(local-label-case)
|     |     |  +--ro local-label?      netc:mpls-label
|     |     +--:(ipv6-address-case)
|     |     |  +--ro ipv6-address?     inet:ipv6-address
|     |     +--:(sid-case)
|     |        +--ro sid?              uint32
|     +--ro sr-lan-adj-ids*
```

(continues on next page)

(continued from previous page)

```
|      |  +--ro (flags)?
|      |  |  +--:(ospf-adj-flags-case)
|      |  |  |  +--ro backup?           boolean
|      |  |  |  +--ro set?              boolean
|      |  |  +--:(isis-adj-flags-case)
|      |  |     +--ro backup?           boolean
|      |  |     +--ro set?              boolean
|      |  |     +--ro address-family?   boolean
|      |  +--ro weight?           weight
|      |  +--ro iso-system-id?    netc:iso-system-identifier
|      |  +--ro neighbor-id?      inet:ipv4-address
|      |  +--ro (sid-label-index)?
|      |     +--:(local-label-case)
|      |     |  +--ro local-label?      netc:mpls-label
|      |     +--:(ipv6-address-case)
|      |     |  +--ro ipv6-address?     inet:ipv6-address
|      |     +--:(sid-case)
|      |        +--ro sid?              uint32
|      +--ro peer-node-sid
|      |  +--ro weight?         weight
|      |  +--ro (sid-label-index)?
|      |     +--:(local-label-case)
|      |     |  +--ro local-label?    netc:mpls-label
|      |     +--:(ipv6-address-case)
|      |     |  +--ro ipv6-address?   inet:ipv6-address
|      |     +--:(sid-case)
|      |        +--ro sid?            uint32
|      +--ro peer-adj-sid
|      |  +--ro weight?         weight
|      |  +--ro (sid-label-index)?
|      |     +--:(local-label-case)
|      |     |  +--ro local-label?    netc:mpls-label
|      |     +--:(ipv6-address-case)
|      |     |  +--ro ipv6-address?   inet:ipv6-address
|      |     +--:(sid-case)
|      |        +--ro sid?            uint32
|      +--ro peer-set-sids*
|         +--ro weight?         weight
|         +--ro (sid-label-index)?
|            +--:(local-label-case)
|            |  +--ro local-label?    netc:mpls-label
|            +--:(ipv6-address-case)
|            |  +--ro ipv6-address?   inet:ipv6-address
|            +--:(sid-case)
|               +--ro sid?            uint32
+--:(prefix-attributes-case)
|  +--ro prefix-attributes
|     +--ro igp-bits
|     |  x--ro up-down?               bits
|     |  +--ro is-is-up-down?         boolean
|     |  +--ro ospf-no-unicast?       boolean
|     |  +--ro ospf-local-address?    boolean
|     |  +--ro ospf-propagate-nssa?   boolean
|     +--ro route-tags*                route-tag
|     +--ro extended-tags*             extended-route-tag
|     +--ro prefix-metric?             netc:igp-metric
|     +--ro ospf-forwarding-address?   inet:ip-address
```

(continues on next page)

(continued from previous page)

```
                |      +--ro sr-prefix
                |      |  +--ro (flags)?
                |      |  |  +--:(isis-prefix-flags-case)
                |      |  |  |  +--ro no-php?            boolean
                |      |  |  |  +--ro explicit-null?     boolean
                |      |  |  |  +--ro readvertisement?   boolean
                |      |  |  |  +--ro node-sid?          boolean
                |      |  |  +--:(ospf-prefix-flags-case)
                |      |  |     +--ro no-php?            boolean
                |      |  |     +--ro explicit-null?     boolean
                |      |  |     +--ro mapping-server?    boolean
                |      |  +--ro algorithm?         algorithm
                |      |  +--ro (sid-label-index)?
                |      |     +--:(local-label-case)
                |      |     |  +--ro local-label?       netc:mpls-label
                |      |     +--:(ipv6-address-case)
                |      |     |  +--ro ipv6-address?      inet:ipv6-address
                |      |     +--:(sid-case)
                |      |        +--ro sid?               uint32
                |      +--ro ipv6-sr-prefix
                |      |  +--ro algorithm?   algorithm
                |      +--ro sr-range
                |      |  +--ro inter-area?   boolean
                |      |  +--ro range-size?   uint16
                |      |  +--ro sub-tlvs*
                |      |     +--ro (range-sub-tlv)?
                |      |        +--:(binding-sid-tlv-case)
                |      |        |  +--ro weight?                weight
                |      |        |  +--ro (flags)?
                |      |        |  |  +--:(isis-binding-flags-case)
                |      |        |  |  |  +--ro address-family?        boolean
                |      |        |  |  |  +--ro mirror-context?        boolean
                |      |        |  |  |  +--ro spread-tlv?            boolean
                |      |        |  |  |  +--ro leaked-from-level-2?   boolean
                |      |        |  |  |  +--ro attached-flag?         boolean
                |      |        |  |  +--:(ospf-binding-flags-case)
                |      |        |  |     +--ro mirroring?             boolean
                |      |        |  +--ro binding-sub-tlvs*
                |      |        |     +--ro (binding-sub-tlv)?
                |      |        |        +--:(prefix-sid-case)
                |      |        |        |  +--ro (flags)?
                |      |        |        |  |  +--:(isis-prefix-flags-case)
                |      |        |        |  |  |  +--ro no-php?            boolean
                |      |        |        |  |  |  +--ro explicit-null?     boolean
                |      |        |        |  |  |  +--ro readvertisement?   boolean
                |      |        |        |  |  |  +--ro node-sid?          boolean
                |      |        |        |  |  +--:(ospf-prefix-flags-case)
                |      |        |        |  |     +--ro no-php?            boolean
                |      |        |        |  |     +--ro explicit-null?     boolean
                |      |        |        |  |     +--ro mapping-server?    boolean
                |      |        |        |  +--ro algorithm?         algorithm
                |      |        |        |  +--ro (sid-label-index)?
                |      |        |        |     +--:(local-label-case)
                |      |        |        |     |  +--ro local-label?       netc:mpls-
↪label
                |      |        |        |     +--:(ipv6-address-case)
                |      |        |        |     |  +--ro ipv6-address?      inet:ipv6-
↪address
```

(continues on next page)

(continued from previous page)

```
                |      |         |         |     +--:(sid-case)
                |      |         |         |        +--ro sid?               uint32
                |      |         |         +--:(ipv6-prefix-sid-case)
                |      |         |         |  +--ro algorithm?         algorithm
                |      |         |         +--:(sid-label-case)
                |      |         |         |  +--ro (sid-label-index)?
                |      |         |         |     +--:(local-label-case)
                |      |         |         |     |  +--ro local-label?       netc:mpls-
↪label
                |      |         |         |     +--:(ipv6-address-case)
                |      |         |         |     |  +--ro ipv6-address?      inet:ipv6-
↪address
                |      |         |         |     +--:(sid-case)
                |      |         |         |        +--ro sid?               uint32
                |      |         |         +--:(ero-metric-case)
                |      |         |         |  +--ro ero-metric?        netc:te-metric
                |      |         |         +--:(ipv4-ero-case)
                |      |         |         |  +--ro loose?             boolean
                |      |         |         |  +--ro address            inet:ipv4-address
                |      |         |         +--:(ipv6-ero-case)
                |      |         |         |  +--ro loose?             boolean
                |      |         |         |  +--ro address            inet:ipv6-address
                |      |         |         +--:(unnumbered-interface-id-ero-case)
                |      |         |         |  +--ro loose?             boolean
                |      |         |         |  +--ro router-id?         uint32
                |      |         |         |  +--ro interface-id?      uint32
                |      |         |         +--:(ipv4-ero-backup-case)
                |      |         |         |  +--ro loose?             boolean
                |      |         |         |  +--ro address            inet:ipv4-address
                |      |         |         +--:(ipv6-ero-backup-case)
                |      |         |         |  +--ro loose?             boolean
                |      |         |         |  +--ro address            inet:ipv6-address
                |      |         |         +--:(unnumbered-interface-id-backup-ero-case)
                |      |         |            +--ro loose?             boolean
                |      |         |            +--ro router-id?         uint32
                |      |         |            +--ro interface-id?      uint32
                |      |         +--:(prefix-sid-tlv-case)
                |      |         |  +--ro (flags)?
                |      |         |  |  +--:(isis-prefix-flags-case)
                |      |         |  |  |  +--ro no-php?                boolean
                |      |         |  |  |  +--ro explicit-null?         boolean
                |      |         |  |  |  +--ro readvertisement?       boolean
                |      |         |  |  |  +--ro node-sid?              boolean
                |      |         |  |  +--:(ospf-prefix-flags-case)
                |      |         |  |     +--ro no-php?                boolean
                |      |         |  |     +--ro explicit-null?         boolean
                |      |         |  |     +--ro mapping-server?        boolean
                |      |         |  +--ro algorithm?             algorithm
                |      |         |  +--ro (sid-label-index)?
                |      |         |     +--:(local-label-case)
                |      |         |     |  +--ro local-label?           netc:mpls-label
                |      |         |     +--:(ipv6-address-case)
                |      |         |     |  +--ro ipv6-address?          inet:ipv6-address
                |      |         |     +--:(sid-case)
                |      |         |        +--ro sid?                   uint32
                |      |         +--:(ipv6-prefix-sid-tlv-case)
                |      |         |  +--ro algorithm?             algorithm
```

(continues on next page)

(continued from previous page)

```
|     |        +--:(sid-label-tlv-case)
|     |           +--ro (sid-label-index)?
|     |              +--:(local-label-case)
|     |              |  +--ro local-label?           netc:mpls-label
|     |              +--:(ipv6-address-case)
|     |              |  +--ro ipv6-address?          inet:ipv6-address
|     |              +--:(sid-case)
|     |                 +--ro sid?                    uint32
|     +--ro sr-binding-sid-labels*
|        +--ro weight?               weight
|        +--ro (flags)?
|        |  +--:(isis-binding-flags-case)
|        |  |  +--ro address-family?        boolean
|        |  |  +--ro mirror-context?        boolean
|        |  |  +--ro spread-tlv?            boolean
|        |  |  +--ro leaked-from-level-2?   boolean
|        |  |  +--ro attached-flag?         boolean
|        |  +--:(ospf-binding-flags-case)
|        |     +--ro mirroring?             boolean
|        +--ro binding-sub-tlvs*
|           +--ro (binding-sub-tlv)?
|              +--:(prefix-sid-case)
|              |  +--ro (flags)?
|              |  |  +--:(isis-prefix-flags-case)
|              |  |  |  +--ro no-php?            boolean
|              |  |  |  +--ro explicit-null?     boolean
|              |  |  |  +--ro readvertisement?   boolean
|              |  |  |  +--ro node-sid?          boolean
|              |  |  +--:(ospf-prefix-flags-case)
|              |  |     +--ro no-php?            boolean
|              |  |     +--ro explicit-null?     boolean
|              |  |     +--ro mapping-server?    boolean
|              |  +--ro algorithm?         algorithm
|              |  +--ro (sid-label-index)?
|              |     +--:(local-label-case)
|              |     |  +--ro local-label?       netc:mpls-label
|              |     +--:(ipv6-address-case)
|              |     |  +--ro ipv6-address?      inet:ipv6-address
|              |     +--:(sid-case)
|              |        +--ro sid?               uint32
|              +--:(ipv6-prefix-sid-case)
|              |  +--ro algorithm?         algorithm
|              +--:(sid-label-case)
|              |  +--ro (sid-label-index)?
|              |     +--:(local-label-case)
|              |     |  +--ro local-label?       netc:mpls-label
|              |     +--:(ipv6-address-case)
|              |     |  +--ro ipv6-address?      inet:ipv6-address
|              |     +--:(sid-case)
|              |        +--ro sid?               uint32
|              +--:(ero-metric-case)
|              |  +--ro ero-metric?        netc:te-metric
|              +--:(ipv4-ero-case)
|              |  +--ro loose?             boolean
|              |  +--ro address            inet:ipv4-address
|              +--:(ipv6-ero-case)
|              |  +--ro loose?             boolean
```

(continues on next page)

(continued from previous page)

```
|              |  +--ro address            inet:ipv6-address
|              +--:(unnumbered-interface-id-ero-case)
|              |  +--ro loose?             boolean
|              |  +--ro router-id?         uint32
|              |  +--ro interface-id?      uint32
|              +--:(ipv4-ero-backup-case)
|              |  +--ro loose?             boolean
|              |  +--ro address            inet:ipv4-address
|              +--:(ipv6-ero-backup-case)
|              |  +--ro loose?             boolean
|              |  +--ro address            inet:ipv6-address
|              +--:(unnumbered-interface-id-backup-ero-case)
|                 +--ro loose?             boolean
|                 +--ro router-id?         uint32
|                 +--ro interface-id?      uint32
x--:(te-lsp-attributes-case)
   +--ro te-lsp-attributes
```

Usage

The Link-State table in a instance of the speaker's Loc-RIB can be verified via REST:

URL: `/restconf/operational/bgp-rib:bgp-rib/rib/bgp-example/loc-rib/tables/bgp-linkstate:linkstate-address-family/bgp-linkstate:linkstate-subsequent-address-family/linkstate-routes`

Method: `GET`

Response Body:

```
<linkstate-routes xmlns="urn:opendaylight:params:xml:ns:yang:bgp-linkstate">
   ...
</linkstate-routes>
```

Note: Link-State routes mapping to topology links/nodes/prefixes is supported by BGP Topology Provider.

References

- North-Bound Distribution of Link-State and Traffic Engineering (TE) Information Using BGP
- BGP Link-State extensions for Segment Routing
- Segment Routing BGP Egress Peer Engineering BGP-LS Extensions
- BGP Link-State Information Distribution Implementation Report

Flow Specification Family

The BGP Flow Specification (BGP-FS) Multiprotocol extension can be used to distribute traffic flow specifications. For example, the BGP-FS can be used in a case of (distributed) denial-of-service (DDoS) attack mitigation procedures and traffic filtering (BGP/MPLS VPN service, DC).

Contents

- *Configuration*
  - *BGP Speaker*
  - *BGP Peer*
- *Flow Specification API*
  - *IPv4 Flow Specification Route*
  - *IPv6 Flow Specification Route*
- *Usage*
  - *IPv4 Flow Specification*
  - *IPv6 Flows Specification*
  - *IPv4 L3VPN Flows Specification*
- *Programming*
  - *IPv4 Flow Specification*
  - *IPv4 L3VPN Flow Specification*
  - *IPv6 Flow Specification*
- *References*

Configuration

This section shows a way to enable BGP-FS family in BGP speaker and peer configuration.

BGP Speaker

To enable BGP-FS support in BGP plugin, first configure BGP speaker instance:

URL: `/restconf/config/openconfig-network-instance:network-instances/network-instance/global-bgp/openconfig-network-instance:protocols`

Method: `POST`

Content-Type: `application/xml`

Request Body:

```
<protocol xmlns="http://openconfig.net/yang/network-instance">
    <name>bgp-example</name>
    <identifier xmlns:x="http://openconfig.net/yang/policy-types">x:BGP</identifier>
    <bgp xmlns="urn:opendaylight:params:xml:ns:yang:bgp:openconfig-extensions">
        <global>
            <config>
                <router-id>192.0.2.2</router-id>
                <as>65000</as>
            </config>
            <afi-safis>
```

(continues on next page)

(continued from previous page)

```
                <afi-safi>
                    <afi-safi-name>IPV4-FLOW</afi-safi-name>
                </afi-safi>
                <afi-safi>
                    <afi-safi-name>IPV6-FLOW</afi-safi-name>
                </afi-safi>
                <afi-safi>
                    <afi-safi-name>IPV4-L3VPN-FLOW</afi-safi-name>
                </afi-safi>
                <afi-safi>
                    <afi-safi-name>IPV6-L3VPN-FLOW</afi-safi-name>
                </afi-safi>
            </afi-safis>
        </global>
    </bgp>
</protocol>
```

BGP Peer

Here is an example for BGP peer configuration with enabled BGP-FS family.

URL: `/restconf/config/openconfig-network-instance:network-instances/network-instance/global-bgp/openconfig-network-instance:protocols/protocol/openconfig-policy-types:BGP/bgp-example/bgp/neighbors`

Method: `POST`

Content-Type: `application/xml`

Request Body:

```
<neighbor xmlns="urn:opendaylight:params:xml:ns:yang:bgp:openconfig-extensions">
    <neighbor-address>192.0.2.1</neighbor-address>
    <afi-safis>
        <afi-safi>
            <afi-safi-name>IPV4-FLOW</afi-safi-name>
        </afi-safi>
        <afi-safi>
            <afi-safi-name>IPV6-FLOW</afi-safi-name>
        </afi-safi>
        <afi-safi>
            <afi-safi-name>IPV4-L3VPN-FLOW</afi-safi-name>
        </afi-safi>
        <afi-safi>
            <afi-safi-name>IPV6-L3VPN-FLOW</afi-safi-name>
        </afi-safi>
    </afi-safis>
</neighbor>
```

Flow Specification API

Following trees illustrate the BGP Flow Specification routes structure.

IPv4 Flow Specification Route

```
:(flowspec-routes-case)
  +--ro flowspec-routes
     +--ro flowspec-route* [route-key path-id]
        +--ro route-key     string
        +--ro flowspec*
        |  +--ro (flowspec-type)?
        |     +--:(port-case)
        |     |  +--ro ports*
        |     |     +--ro op?      numeric-operand
        |     |     +--ro value?   uint16
        |     +--:(destination-port-case)
        |     |  +--ro destination-ports*
        |     |     +--ro op?      numeric-operand
        |     |     +--ro value?   uint16
        |     +--:(source-port-case)
        |     |  +--ro source-ports*
        |     |     +--ro op?      numeric-operand
        |     |     +--ro value?   uint16
        |     +--:(icmp-type-case)
        |     |  +--ro types*
        |     |     +--ro op?      numeric-operand
        |     |     +--ro value?   uint8
        |     +--:(icmp-code-case)
        |     |  +--ro codes*
        |     |     +--ro op?      numeric-operand
        |     |     +--ro value?   uint8
        |     +--:(tcp-flags-case)
        |     |  +--ro tcp-flags*
        |     |     +--ro op?      bitmask-operand
        |     |     +--ro value?   uint16
        |     +--:(packet-length-case)
        |     |  +--ro packet-lengths*
        |     |     +--ro op?      numeric-operand
        |     |     +--ro value?   uint16
        |     +--:(dscp-case)
        |     |  +--ro dscps*
        |     |     +--ro op?      numeric-operand
        |     |     +--ro value?   dscp
        |     +--:(fragment-case)
        |     |  +--ro fragments*
        |     |     +--ro op?      bitmask-operand
        |     |     +--ro value?   fragment
        |     +--:(destination-prefix-case)
        |     |  +--ro destination-prefix?   inet:ipv4-prefix
        |     +--:(source-prefix-case)
        |     |  +--ro source-prefix?        inet:ipv4-prefix
        |     +--:(protocol-ip-case)
        |        +--ro protocol-ips*
        |           +--ro op?      numeric-operand
        |           +--ro value?   uint8
        +--ro path-id       path-id
        +--ro attributes
           +--ro extended-communities*
              +--ro transitive?                          boolean
              +--ro (extended-community)?
```

(continues on next page)

(continued from previous page)

```
                +--:(traffic-rate-extended-community-case)
                |  +--ro traffic-rate-extended-community
                |     +--ro informative-as?        bgp-t:short-as-number
                |     +--ro local-administrator?   netc:bandwidth
                +--:(traffic-action-extended-community-case)
                |  +--ro traffic-action-extended-community
                |     +--ro sample?            boolean
                |     +--ro terminal-action?   boolean
                +--:(redirect-extended-community-case)
                |  +--ro redirect-extended-community
                |     +--ro global-administrator?   bgp-t:short-as-number
                |     +--ro local-administrator?    binary
                +--:(traffic-marking-extended-community-case)
                |  +--ro traffic-marking-extended-community
                |     +--ro global-administrator?   dscp
                +--:(redirect-ipv4-extended-community-case)
                |  +--ro redirect-ipv4
                |     +--ro global-administrator?   inet:ipv4-address
                |     +--ro local-administrator?    uint16
                +--:(redirect-as4-extended-community-case)
                |  +--ro redirect-as4
                |     +--ro global-administrator?   inet:as-number
                |     +--ro local-administrator?    uint16
                +--:(redirect-ip-nh-extended-community-case)
                  +--ro redirect-ip-nh-extended-community
                     +--ro next-hop-address?   inet:ip-address
                     +--ro copy?               boolean
```

IPv6 Flow Specification Route

```
:(flowspec-ipv6-routes-case)
  +--ro flowspec-ipv6-routes
     +--ro flowspec-route* [route-key path-id]
        +--ro flowspec*
        |  +--ro (flowspec-type)?
        |     +--:(port-case)
        |     |  +--ro ports*
        |     |     +--ro op?      numeric-operand
        |     |     +--ro value?   uint16
        |     +--:(destination-port-case)
        |     |  +--ro destination-ports*
        |     |     +--ro op?      numeric-operand
        |     |     +--ro value?   uint16
        |     +--:(source-port-case)
        |     |  +--ro source-ports*
        |     |     +--ro op?      numeric-operand
        |     |     +--ro value?   uint16
        |     +--:(icmp-type-case)
        |     |  +--ro types*
        |     |     +--ro op?      numeric-operand
        |     |     +--ro value?   uint8
        |     +--:(icmp-code-case)
        |     |  +--ro codes*
        |     |     +--ro op?      numeric-operand
```

(continues on next page)

(continued from previous page)

```
|      |     +--ro value?   uint8
|      +--:(tcp-flags-case)
|      |  +--ro tcp-flags*
|      |     +--ro op?      bitmask-operand
|      |     +--ro value?   uint16
|      +--:(packet-length-case)
|      |  +--ro packet-lengths*
|      |     +--ro op?      numeric-operand
|      |     +--ro value?   uint16
|      +--:(dscp-case)
|      |  +--ro dscps*
|      |     +--ro op?      numeric-operand
|      |     +--ro value?   dscp
|      +--:(fragment-case)
|      |  +--ro fragments*
|      |     +--ro op?      bitmask-operand
|      |     +--ro value?   fragment
|      +--:(destination-ipv6-prefix-case)
|      |  +--ro destination-prefix?   inet:ipv6-prefix
|      +--:(source-ipv6-prefix-case)
|      |  +--ro source-prefix?        inet:ipv6-prefix
|      +--:(next-header-case)
|      |  +--ro next-headers*
|      |     +--ro op?      numeric-operand
|      |     +--ro value?   uint8
|      +--:(flow-label-case)
|         +--ro flow-label*
|            +--ro op?      numeric-operand
|            +--ro value?   uint32
+--ro path-id       path-id
+--ro attributes
   +--ro extended-communities*
      +--ro transitive?                              boolean
      +--ro (extended-community)?
         +--:(traffic-rate-extended-community-case)
         |  +--ro traffic-rate-extended-community
         |     +--ro informative-as?        bgp-t:short-as-number
         |     +--ro local-administrator?   netc:bandwidth
         +--:(traffic-action-extended-community-case)
         |  +--ro traffic-action-extended-community
         |     +--ro sample?            boolean
         |     +--ro terminal-action?   boolean
         +--:(redirect-extended-community-case)
         |  +--ro redirect-extended-community
         |     +--ro global-administrator?   bgp-t:short-as-number
         |     +--ro local-administrator?    binary
         +--:(traffic-marking-extended-community-case)
         |  +--ro traffic-marking-extended-community
         |     +--ro global-administrator?   dscp
         +--:(redirect-ipv6-extended-community-case)
         |  +--ro redirect-ipv6
         |     +--ro global-administrator?   inet:ipv6-address
         |     +--ro local-administrator?    uint16
         +--:(redirect-as4-extended-community-case)
         |  +--ro redirect-as4
         |     +--ro global-administrator?   inet:as-number
         |     +--ro local-administrator?    uint16
```

(continues on next page)

(continued from previous page)

```
                +--:(redirect-ip-nh-extended-community-case)
                   +--ro redirect-ip-nh-extended-community
                      +--ro next-hop-address?   inet:ip-address
                      +--ro copy?               boolean
```

Usage

The flowspec route represents rules and an action, defined as an extended community.

IPv4 Flow Specification

The IPv4 Flowspec table in an instance of the speaker's Loc-RIB can be verified via REST:

URL: `/restconf/operational/bgp-rib:bgp-rib/rib/bgp-example/loc-rib/tables/bgp-types:ipv4-address-family/bgp-flowspec:flowspec-subsequent-address-family/bgp-flowspec:flowspec-routes`

Method: `GET`

Response Body:

```
<flowspec-routes xmlns="urn:opendaylight:params:xml:ns:yang:bgp-flowspec">
    <flowspec-route>
        <path-id>0</path-id>
        <route-key>all packets to 192.168.0.1/32 AND from 10.0.0.2/32 AND where IP␣
↪protocol equals to 17 or equals to 6 AND where port equals to 80 or equals to 8080␣
↪AND where destination port is greater than 8080 and is less than 8088 or equals to␣
↪3128 AND where source port is greater than 1024 </route-key>
        <attributes>
            <local-pref>
                <pref>100</pref>
            </local-pref>
            <origin>
                <value>igp</value>
            </origin>
            <as-path></as-path>
            <extended-communities>
                <transitive>true</transitive>
                <redirect-extended-community>
                    <local-administrator>AgMWLg==</local-administrator>
                    <global-administrator>258</global-administrator>
                </redirect-extended-community>
            </extended-communities>
        </attributes>
        <flowspec>
            <destination-prefix>192.168.0.1/32</destination-prefix>
        </flowspec>
        <flowspec>
            <source-prefix>10.0.0.2/32</source-prefix>
        </flowspec>
        <flowspec>
            <protocol-ips>
                <op>equals</op>
                <value>17</value>
```

(continues on next page)

(continued from previous page)

```
            </protocol-ips>
            <protocol-ips>
                <op>equals end-of-list</op>
                <value>6</value>
            </protocol-ips>
        </flowspec>
        <flowspec>
            <ports>
                <op>equals</op>
                <value>80</value>
            </ports>
            <ports>
                <op>equals end-of-list</op>
                <value>8080</value>
            </ports>
        </flowspec>
        <flowspec>
            <destination-ports>
                <op>greater-than</op>
                <value>8080</value>
            </destination-ports>
            <destination-ports>
                <op>less-than and-bit</op>
                <value>8088</value>
            </destination-ports>
            <destination-ports>
                <op>equals end-of-list</op>
                <value>3128</value>
            </destination-ports>
        </flowspec>
        <flowspec>
            <source-ports>
                <op>end-of-list greater-than</op>
                <value>1024</value>
            </source-ports>
        </flowspec>
    </flowspec-route>
</flowspec-routes>
```

IPv6 Flows Specification

The IPv6 Flowspec table in an instance of the speaker's Loc-RIB can be verified via REST:

URL: `/restconf/operational/bgp-rib:bgp-rib/rib/bgp-example/loc-rib/tables/bgp-types:ipv6-address-family/bgp-flowspec:flowspec-subsequent-address-family/bgp-flowspec:flowspec-ipv6-routes`

Method: `GET`

Response Body:

```
<flowspec-ipv6-routes xmlns="urn:opendaylight:params:xml:ns:yang:bgp-flowspec">
    <flowspec-route>
        <path-id>0</path-id>
        <route-key>all packets to 2001:db8:31::/64 AND from 2001:db8:30::/64 AND␣
↪where next header equals to 17 AND where DSCP equals to 50 AND where flow label␣
↪equals to 2013 </route-key>
```

(continues on next page)

(continued from previous page)

```
        <attributes>
            <local-pref>
                <pref>100</pref>
            </local-pref>
            <origin>
                <value>igp</value>
            </origin>
            <as-path></as-path>
            <extended-communities>
                <transitive>true</transitive>
                <traffic-rate-extended-community>
                    <informative-as>0</informative-as>
                    <local-administrator>AAAAAA==</local-administrator>
                </traffic-rate-extended-community>
            </extended-communities>
        </attributes>
        <flowspec>
            <destination-prefix>2001:db8:31::/64</destination-prefix>
        </flowspec>
        <flowspec>
            <source-prefix>2001:db8:30::/64</source-prefix>
        </flowspec>
        <flowspec>
            <next-headers>
                <op>equals end-of-list</op>
                <value>17</value>
            </next-headers>
        </flowspec>
        <flowspec>
            <dscps>
                <op>equals end-of-list</op>
                <value>50</value>
            </dscps>
        </flowspec>
        <flowspec>
            <flow-label>
                <op>equals end-of-list</op>
                <value>2013</value>
            </flow-label>
        </flowspec>
    </flowspec-route>
</flowspec-ipv6-routes>
```

IPv4 L3VPN Flows Specification

The IPv4 L3VPN Flowspec table in an instance of the speaker's Loc-RIB can be verified via REST:

URL: `/restconf/operational/bgp-rib:bgp-rib/rib/bgp-example/loc-rib/tables/bgp-types:ipv4-address-family/bgp-flowspec:flowspec-l3vpn-subsequent-address-family/bgp-flowspec:flowspec-l3vpn-ipv4-routes`

Method: `GET`

Response Body:

```
<flowspec-l3vpn-ipv4-routes xmlns="urn:opendaylight:params:xml:ns:yang:bgp-flowspec">
    <flowspec-l3vpn-route>
        <path-id>0</path-id>
        <route-key>[l3vpn with route-distinguisher 172.16.0.44:101] all packets from␣
↪10.0.0.3/32</route-key>
        <attributes>
            <local-pref>
                <pref>100</pref>
            </local-pref>
            <ipv4-next-hop>
                <global>5.6.7.8</global>
            </ipv4-next-hop>
            <origin>
                <value>igp</value>
            </origin>
            <as-path></as-path>
            <extended-communities>
                <transitive>true</transitive>
                <redirect-ip-nh-extended-community>
                    <copy>false</copy>
                    <next-hop-address>0.0.0.0</next-hop-address>
                </redirect-ip-nh-extended-community>
            </extended-communities>
        </attributes>
        <route-distinguisher>172.16.0.44:101</route-distinguisher>
        <flowspec>
            <source-prefix>10.0.0.3/32</source-prefix>
        </flowspec>
    </flowspec-l3vpn-route>
</flowspec-l3vpn-ipv4-routes>
```

Programming

IPv4 Flow Specification

This examples show how to originate and remove IPv4 fowspec route via programmable RIB. Make sure the *Application Peer* is configured first.

URL: `/restconf/config/bgp-rib:application-rib/10.25.1.9/tables/bgp-types:ipv4-address-family/bgp-flowspec:flowspec-subsequent-address-family/bgp-flowspec:flowspec-routes`

Method: `POST`

Content-Type: `application/xml`

Request Body:

```
<flowspec-route xmlns="urn:opendaylight:params:xml:ns:yang:bgp-flowspec">
    <route-key>flow1</route-key>
    <path-id>0</path-id>
    <flowspec>
        <destination-prefix>192.168.0.1/32</destination-prefix>
    </flowspec>
    <flowspec>
        <source-prefix>10.0.0.1/32</source-prefix>
```

(continues on next page)

(continued from previous page)

```
    </flowspec>
    <flowspec>
        <protocol-ips>
            <op>equals end-of-list</op>
            <value>6</value>
        </protocol-ips>
    </flowspec>
    <flowspec>
        <ports>
            <op>equals end-of-list</op>
            <value>80</value>
        </ports>
    </flowspec>
    <flowspec>
        <destination-ports>
            <op>greater-than</op>
            <value>8080</value>
        </destination-ports>
        <destination-ports>
            <op>and-bit less-than end-of-list</op>
            <value>8088</value>
        </destination-ports>
    </flowspec>
    <flowspec>
        <source-ports>
            <op>greater-than end-of-list</op>
            <value>1024</value>
        </source-ports>
    </flowspec>
    <flowspec>
        <types>
            <op>equals end-of-list</op>
            <value>0</value>
        </types>
    </flowspec>
    <flowspec>
        <codes>
            <op>equals end-of-list</op>
            <value>0</value>
        </codes>
    </flowspec>
    <flowspec>
        <tcp-flags>
            <op>match end-of-list</op>
            <value>32</value>
        </tcp-flags>
    </flowspec>
    <flowspec>
        <packet-lengths>
            <op>greater-than</op>
            <value>400</value>
        </packet-lengths>
        <packet-lengths>
            <op>and-bit less-than end-of-list</op>
            <value>500</value>
        </packet-lengths>
    </flowspec>
```

(continues on next page)

(continued from previous page)

```
    <flowspec>
        <dscps>
            <op>equals end-of-list</op>
            <value>20</value>
        </dscps>
    </flowspec>
    <flowspec>
        <fragments>
            <op>match end-of-list</op>
            <value>first</value>
        </fragments>
    </flowspec>
    <attributes>
        <origin>
            <value>igp</value>
        </origin>
        <as-path/>
        <local-pref>
            <pref>100</pref>
        </local-pref>
        <extended-communities>
            ....
        </extended-communities>
    </attributes>
</flowspec-route>
```

Extended Communities

- **Traffic Rate**

```
1 <extended-communities>
2     <transitive>true</transitive>
3     <traffic-rate-extended-community>
4         <informative-as>123</informative-as>
5         <local-administrator>AAAAAA==</local-administrator>
6     </traffic-rate-extended-community>
7 </extended-communities>
```

@line 5: A rate in bytes per second, *AAAAAA*== (0) means traffic discard.

- **Traffic Action**

```
<extended-communities>
    <transitive>true</transitive>
    <traffic-action-extended-community>
        <sample>true</sample>
        <terminal-action>false</terminal-action>
    </traffic-action-extended-community>
</extended-communities>
```

- **Redirect to VRF AS 2byte format**

```
<extended-communities>
    <transitive>true</transitive>
    <redirect-extended-community>
        <global-administrator>123</global-administrator>
```

(continues on next page)

(continued from previous page)

```
        <local-administrator>AAAAew==</local-administrator>
    </redirect-extended-community>
</extended-communities>
```

- **Redirect to VRF IPv4 format**

```
<extended-communities>
    <transitive>true</transitive>
    <redirect-ipv4>
        <global-administrator>192.168.0.1</global-administrator>
        <local-administrator>12345</local-administrator>
    </redirect-ipv4>
</extended-communities>
```

- **Redirect to VRF AS 4byte format**

```
<extended-communities>
    <transitive>true</transitive>
    <redirect-as4>
        <global-administrator>64495</global-administrator>
        <local-administrator>12345</local-administrator>
    </redirect-as4>
</extended-communities>
```

- **Redirect to IP**

```
<extended-communities>
    <transitive>true</transitive>
    <redirect-ip-nh-extended-community>
        <copy>false</false>
    </redirect-ip-nh-extended-community>
</extended-communities>
```

- **Traffic Marking**

```
<extended-communities>
    <transitive>true</transitive>
    <traffic-marking-extended-community>
        <global-administrator>20</global-administrator>
    </traffic-marking-extended-community>
</extended-communities>
```

To remove the route added above, following request can be used:

URL: `/restconf/config/bgp-rib:application-rib/10.25.1.9/tables/bgp-types:ipv4-address-family/bgp-flowspec:flowspec-subsequent-address-family/bgp-flowspec:flowspec-routes/bgp-flowspec:flowspec-route/flow1/0`

Method: `DELETE`

IPv4 L3VPN Flow Specification

This examples show how to originate and remove IPv4 L3VPN fowspec route via programmable RIB.

URL: /restconf/config/bgp-rib:application-rib/10.25.1.9/tables/bgp-types:ipv4-address-family/bgp-flowspec:flowspec-l3vpn-subsequent-address-family/bgp-flowspec:flowspec-l3vpn-ipv4-routes

Method: POST

Content-Type: application/xml

Request Body:

```
<flowspec-l3vpn-route xmlns="urn:opendaylight:params:xml:ns:yang:bgp-flowspec">
    <path-id>0</path-id>
    <route-key>flow-l3vpn</route-key>
    <route-distinguisher>172.16.0.44:101</route-distinguisher>
    <flowspec>
        <source-prefix>10.0.0.3/32</source-prefix>
    </flowspec>
    <attributes>
        <local-pref>
            <pref>100</pref>
        </local-pref>
        <origin>
            <value>igp</value>
        </origin>
        <as-path></as-path>
           <extended-communities>
               <transitive>true</transitive>
               <redirect-ipv4>
                   <global-administrator>172.16.0.44</global-administrator>
                   <local-administrator>102</local-administrator>
               </redirect-ipv4>
           </extended-communities>
    </attributes>
</flowspec-l3vpn-route>
```

To remove the route added above, following request can be used:

URL: /restconf/config/bgp-rib:application-rib/10.25.1.9/tables/bgp-types:ipv4-address-family/bgp-flowspec:flowspec-l3vpn-subsequent-address-family/bgp-flowspec:flowspec-l3vpn-ipv4-routes/flowspec-l3vpn-route/flow-l3vpn/0

Method: DELETE

IPv6 Flow Specification

This examples show how to originate and remove IPv6 fowspec route via programmable RIB.

URL: /restconf/config/bgp-rib:application-rib/10.25.1.9/tables/bgp-types:ipv6-address-family/bgp-flowspec:flowspec-subsequent-address-family/bgp-flowspec:flowspec-ipv6-routes

Method: POST

Content-Type: application/xml

Request Body:

```
<flowspec-route xmlns="urn:opendaylight:params:xml:ns:yang:bgp-flowspec">
    <route-key>flow-v6</route-key>
    <path-id>0</path-id>
    <flowspec>
        <destination-prefix>2001:db8:30::3/128</destination-prefix>
    </flowspec>
    <flowspec>
        <source-prefix>2001:db8:31::3/128</source-prefix>
     </flowspec>
    <flowspec>
        <flow-label>
            <op>equals end-of-list</op>
            <value>1</value>
        </flow-label>
    </flowspec>
    <attributes>
        <extended-communities>
            <transitive>true</transitive>
            <redirect-ipv6>
                <global-administrator>2001:db8:1::6</global-administrator>
                <local-administrator>12345</local-administrator>
            </redirect-ipv6>
        </extended-communities>
        <origin>
            <value>igp</value>
        </origin>
        <as-path/>
        <local-pref>
            <pref>100</pref>
        </local-pref>
    </attributes>
</flowspec-route>
```

To remove the route added above, following request can be used:

URL: `/restconf/config/bgp-rib:application-rib/10.25.1.9/tables/bgp-types:ipv6-address-family/bgp-flowspec:flowspec-subsequent-address-family/bgp-flowspec:flowspec-ipv6-routes/bgp-flowspec:flowspec-route/flow-v6/0`

Method: `DELETE`

References

- Dissemination of Flow Specification Rules
- Dissemination of Flow Specification Rules for IPv6
- BGP Flow-Spec Extended Community for Traffic Redirect to IP Next Hop
- Clarification of the Flowspec Redirect Extended Community
- Revised Validation Procedure for BGP Flow Specifications

EVPN Family

The BGP MPLS-Based Ethernet VPN (BGP EVPN) Multiprotocol extension can be used to distribute Ethernet L2VPN service related routes in order to support a concept of MAC routing. A major use-case for BGP EVPN is data-center interconnection (DCI), where advantage of BGP EVPN are MAC/IP address advertising across MPLS network, Multihoming functionality including Fast Convergence, Split Horizon and Aliasing support, VM (MAC) Mobility, support Multicast and Broadcast traffic. In addition to MPLS, IP tunnelling encapsulation techniques like VXLAN, NVGRE, MPLSoGRE and others can be used for packet transportation. Also, Provider Backbone Bridging (PBB) can be combined with EVPN in order to reduce a number of MAC Advertisement routes.

Contents

- *Configuration*
  - *BGP Speaker*
  - *BGP Peer*
- *EVPN Route API*
- *Usage*
- *Programming*
- *References*

Configuration

This section shows a way to enable EVPN family in BGP speaker and peer configuration.

BGP Speaker

To enable EVPN support in BGP plugin, first configure BGP speaker instance:

URL: `/restconf/config/openconfig-network-instance:network-instances/network-instance/global-bgp/openconfig-network-instance:protocols`

Method: `POST`

Content-Type: `application/xml`

Request Body:

```
<protocol xmlns="http://openconfig.net/yang/network-instance">
    <name>bgp-example</name>
    <identifier xmlns:x="http://openconfig.net/yang/policy-types">x:BGP</identifier>
    <bgp xmlns="urn:opendaylight:params:xml:ns:yang:bgp:openconfig-extensions">
        <global>
            <config>
                <router-id>192.0.2.2</router-id>
                <as>65000</as>
            </config>
            <afi-safis>
                <afi-safi>
                    <afi-safi-name xmlns:x="http://openconfig.net/yang/bgp-types">
↪x:L2VPN-EVPN</afi-safi-name>
```

(continues on next page)

(continued from previous page)

```
                </afi-safi>
            </afi-safis>
        </global>
    </bgp>
</protocol>
```

BGP Peer

Here is an example for BGP peer configuration with enabled EVPN family.

URL: `/restconf/config/openconfig-network-instance:network-instances/network-instance/global-bgp/openconfig-network-instance:protocols/protocol/openconfig-policy-types:BGP/bgp-example/bgp/neighbors`

Method: `POST`

Content-Type: `application/xml`

Request Body:

```
<neighbor xmlns="urn:opendaylight:params:xml:ns:yang:bgp:openconfig-extensions">
    <neighbor-address>192.0.2.1</neighbor-address>
    <afi-safis>
        <afi-safi>
            <afi-safi-name xmlns:x="http://openconfig.net/yang/bgp-types">x:L2VPN-EVPN
↪</afi-safi-name>
        </afi-safi>
    </afi-safis>
</neighbor>
```

EVPN Route API

Following tree illustrate the BGP EVPN route structure.

```
:(evpn-routes-case)
   +--ro evpn-routes
      +--ro evpn-route* [route-key]
         +--ro route-key                    string
         +--ro (evpn-choice)
         |  +--:(ethernet-a-d-route-case)
         |  |  +--ro ethernet-a-d-route
         |  |     +--ro (esi)
         |  |     |  +--:(arbitrary-case)
         |  |     |  |  +--ro arbitrary
         |  |     |  |     +--ro arbitrary    binary
         |  |     |  +--:(lacp-auto-generated-case)
         |  |     |  |  +--ro lacp-auto-generated
         |  |     |  |     +--ro ce-lacp-mac-address    yang:mac-address
         |  |     |  |     +--ro ce-lacp-port-key       uint16
         |  |     |  +--:(lan-auto-generated-case)
         |  |     |  |  +--ro lan-auto-generated
         |  |     |  |     +--ro root-bridge-mac-address    yang:mac-address
         |  |     |  |     +--ro root-bridge-priority       uint16
         |  |     |  +--:(mac-auto-generated-case)
```

(continues on next page)

(continued from previous page)

```
|  |     |  |  +--ro mac-auto-generated
|  |     |  |     +--ro system-mac-address     yang:mac-address
|  |     |  |     +--ro local-discriminator    uint24
|  |     |  +--:(router-id-generated-case)
|  |     |  |  +--ro router-id-generated
|  |     |  |     +--ro router-id              inet:ipv4-address
|  |     |  |     +--ro local-discriminator    uint32
|  |     |  +--:(as-generated-case)
|  |     |     +--ro as-generated
|  |     |        +--ro as                     inet:as-number
|  |     |        +--ro local-discriminator    uint32
|  |     +--ro ethernet-tag-id
|  |     |  +--ro vlan-id    uint32
|  |     +--ro mpls-label             netc:mpls-label
|  +--:(mac-ip-adv-route-case)
|  |  +--ro mac-ip-adv-route
|  |     +--ro (esi)
|  |     |  +--:(arbitrary-case)
|  |     |  |  +--ro arbitrary
|  |     |  |     +--ro arbitrary    binary
|  |     |  +--:(lacp-auto-generated-case)
|  |     |  |  +--ro lacp-auto-generated
|  |     |  |     +--ro ce-lacp-mac-address    yang:mac-address
|  |     |  |     +--ro ce-lacp-port-key       uint16
|  |     |  +--:(lan-auto-generated-case)
|  |     |  |  +--ro lan-auto-generated
|  |     |  |     +--ro root-bridge-mac-address    yang:mac-address
|  |     |  |     +--ro root-bridge-priority       uint16
|  |     |  +--:(mac-auto-generated-case)
|  |     |  |  +--ro mac-auto-generated
|  |     |  |     +--ro system-mac-address     yang:mac-address
|  |     |  |     +--ro local-discriminator    uint24
|  |     |  +--:(router-id-generated-case)
|  |     |  |  +--ro router-id-generated
|  |     |  |     +--ro router-id              inet:ipv4-address
|  |     |  |     +--ro local-discriminator    uint32
|  |     |  +--:(as-generated-case)
|  |     |     +--ro as-generated
|  |     |        +--ro as                     inet:as-number
|  |     |        +--ro local-discriminator    uint32
|  |     +--ro ethernet-tag-id
|  |     |  +--ro vlan-id    uint32
|  |     +--ro mac-address            yang:mac-address
|  |     +--ro ip-address?            inet:ip-address
|  |     +--ro mpls-label1            netc:mpls-label
|  |     +--ro mpls-label2?           netc:mpls-label
|  +--:(inc-multi-ethernet-tag-res-case)
|  |  +--ro inc-multi-ethernet-tag-res
|  |     +--ro ethernet-tag-id
|  |     |  +--ro vlan-id    uint32
|  |     +--ro orig-route-ip?     inet:ip-address
|  +--:(es-route-case)
|     +--ro es-route
|        +--ro (esi)
|        |  +--:(arbitrary-case)
|        |  |  +--ro arbitrary
|        |  |     +--ro arbitrary    binary
```

(continues on next page)

(continued from previous page)

```
|        |  +--:(lacp-auto-generated-case)
|        |  |  +--ro lacp-auto-generated
|        |  |     +--ro ce-lacp-mac-address    yang:mac-address
|        |  |     +--ro ce-lacp-port-key       uint16
|        |  +--:(lan-auto-generated-case)
|        |  |  +--ro lan-auto-generated
|        |  |     +--ro root-bridge-mac-address    yang:mac-address
|        |  |     +--ro root-bridge-priority       uint16
|        |  +--:(mac-auto-generated-case)
|        |  |  +--ro mac-auto-generated
|        |  |     +--ro system-mac-address     yang:mac-address
|        |  |     +--ro local-discriminator    uint24
|        |  +--:(router-id-generated-case)
|        |  |  +--ro router-id-generated
|        |  |     +--ro router-id              inet:ipv4-address
|        |  |     +--ro local-discriminator    uint32
|        |  +--:(as-generated-case)
|        |     +--ro as-generated
|        |        +--ro as                     inet:as-number
|        |        +--ro local-discriminator    uint32
|        +--ro orig-route-ip          inet:ip-address
+--ro route-distinguisher           bgp-t:route-distinguisher
+--ro attributes
   +--ro extended-communities*
   |  +--ro transitive?                            boolean
   |  +--ro (extended-community)?
   |     +--:(encapsulation-case)
   |     |  +--ro encapsulation-extended-community
   |     |     +--ro tunnel-type    encapsulation-tunnel-type
   |     +--:(esi-label-extended-community-case)
   |     |  +--ro esi-label-extended-community
   |     |     +--ro single-active-mode?   boolean
   |     |     +--ro esi-label             netc:mpls-label
   |     +--:(es-import-route-extended-community-case)
   |     |  +--ro es-import-route-extended-community
   |     |     +--ro es-import    yang:mac-address
   |     +--:(mac-mobility-extended-community-case)
   |     |  +--ro mac-mobility-extended-community
   |     |     +--ro static?       boolean
   |     |     +--ro seq-number    uint32
   |     +--:(default-gateway-extended-community-case)
   |     |  +--ro default-gateway-extended-community!
   |     +--:(layer-2-attributes-extended-community-case)
   |        +--ro layer-2-attributes-extended-community
   |           +--ro primary-pe?     boolean
   |           +--ro backup-pe?      boolean
   |           +--ro control-word?   boolean
   |           +--ro l2-mtu          uint16
   +--ro pmsi-tunnel!
      +--ro leaf-information-required    boolean
      +--ro mpls-label?                  netc:mpls-label
      +--ro (tunnel-identifier)?
         +--:(rsvp-te-p2mp-lsp)
         |  +--ro rsvp-te-p2mp-lps
         |     +--ro p2mp-id               uint32
         |     +--ro tunnel-id             uint16
         |     +--ro extended-tunnel-id    inet:ip-address
```

(continues on next page)

(continued from previous page)

```
                    +--:(mldp-p2mp-lsp)
                    |  +--ro mldp-p2mp-lsp
                    |     +--ro address-family       identityref
                    |     +--ro root-node-address    inet:ip-address
                    |     +--ro opaque-value*
                    |        +--ro opaque-type             uint8
                    |        +--ro opaque-extended-type?   uint16
                    |        +--ro opaque                  yang:hex-string
                    +--:(pim-ssm-tree)
                    |  +--ro pim-ssm-tree
                    |     +--ro p-address            inet:ip-address
                    |     +--ro p-multicast-group    inet:ip-address
                    +--:(pim-sm-tree)
                    |  +--ro pim-sm-tree
                    |     +--ro p-address            inet:ip-address
                    |     +--ro p-multicast-group    inet:ip-address
                    +--:(bidir-pim-tree)
                    |  +--ro bidir-pim-tree
                    |     +--ro p-address            inet:ip-address
                    |     +--ro p-multicast-group    inet:ip-address
                    +--:(ingress-replication)
                    |  +--ro ingress-replication
                    |     +--ro receiving-endpoint-address?   inet:ip-address
                    +--:(mldp-mp2mp-lsp)
                       +--ro mldp-mp2mp-lsp
                          +--ro opaque-type             uint8
                          +--ro opaque-extended-type?   uint16
                          +--ro opaque
                    ...
```

Usage

The L2VPN EVPN table in an instance of the speaker's Loc-RIB can be verified via REST:

URL: `/restconf/operational/bgp-rib:bgp-rib/rib/bgp-example/loc-rib/tables/odl-bgp-evpn:l2vpn-address-family/odl-bgp-evpn:evpn-subsequent-address-family/evpn-routes`

Method: `GET`

Response Body:

```
<evpn-routes xmlns="urn:opendaylight:params:xml:ns:yang:bgp-evpn">
   <evpn-route>
      <route-key>AxEAAcCoZAED6AAAAQAgwKhkAQ==</route-key>
      <route-distinguisher>192.168.100.1:1000</route-distinguisher>
      <inc-multi-ethernet-tag-res>
         <ethernet-tag-id>
            <vlan-id>256</vlan-id>
         </ethernet-tag-id>
         <orig-route-ip>192.168.100.1</orig-route-ip>
      </inc-multi-ethernet-tag-res>
      <attributes>
         <ipv4-next-hop>
            <global>172.23.29.104</global>
         </ipv4-next-hop>
```

(continues on next page)

(continued from previous page)

```
         <as-path/>
         <origin>
            <value>igp</value>
         </origin>
         <extended-communities>
            <extended-communities>
                <transitive>true</transitive>
                <route-target-extended-community>
                    <global-administrator>65504</global-administrator>
                    <local-administrator>AAAD6A==</local-administrator>
                </route-target-extended-community>
            </extended-communities>
         </extended-communities>
         <pmsi-tunnel>
             <leaf-information-required>true</leaf-information-required>
             <mpls-label>20024</mpls-label>
             <ingress-replication>
                 <receiving-endpoint-address>192.168.100.1</receiving-endpoint-
↪address>
             </ingress-replication>
         </pmsi-tunnel>
      </attributes>
   </evpn-route>
</evpn-routes>
```

Programming

This examples show how to originate and remove EVPN routes via programmable RIB. There are four different types of EVPN routes, and several extended communities. Routes can be used for variety of use-cases supported by BGP/MPLS EVPN, PBB EVPN and NVO EVPN. Make sure the *Application Peer* is configured first.

URL: `/restconf/config/bgp-rib:application-rib/10.25.1.9/tables/odl-bgp-evpn:l2vpn-address-family/odl-bgp-evpn:evpn-subsequent-address-family/odl-bgp-evpn:evpn-routes`

Method: `POST`

Content-Type: `application/xml`

Request Body:

```
<evpn-route xmlns="urn:opendaylight:params:xml:ns:yang:bgp-evpn">
    <route-key>evpn</route-key>
    <route-distinguisher>172.12.123.3:200</route-distinguisher>
    ....
    <attributes>
        <ipv4-next-hop>
            <global>199.20.166.41</global>
        </ipv4-next-hop>
        <as-path/>
        <origin>
            <value>igp</value>
        </origin>
        <extended-communities>
        ....
        </extended-communities>
```

(continues on next page)

(continued from previous page)

```
16      </attributes>
17  </evpn-route>
```

@line 3: Route Distinguisher (RD) - set to RD of the MAC-VRF advertising the NLRI, recommended format *<IP>:<VLAN_ID>*

@line 4: One of the EVPN route must be set here.

@line 14: In some cases, specific extended community presence is required. The route may carry one or more Route Target attributes.

EVPN Routes:

- **Ethernet AD per ESI**

```
<ethernet-a-d-route>
    <mpls-label>0</mpls-label>
    <ethernet-tag-id>
        <vlan-id>4294967295</vlan-id>
    </ethernet-tag-id>
    <arbitrary>
        <arbitrary>AAAAAAAAAAAA</arbitrary>
    </arbitrary>
</ethernet-a-d-route>
```

- **Ethernet AD per EVI**

```
<ethernet-a-d-route>
    <mpls-label>24001</mpls-label>
    <ethernet-tag-id>
        <vlan-id>2200</vlan-id>
    </ethernet-tag-id>
    <arbitrary>
        <arbitrary>AAAAAAAAAAAA</arbitrary>
    </arbitrary>
</ethernet-a-d-route>
```

- **MAC/IP Advertisement**

```
<mac-ip-adv-route>
    <arbitrary>
        <arbitrary>AAAAAAAAAAAA</arbitrary>
    </arbitrary>
    <ethernet-tag-id>
        <vlan-id>2100</vlan-id>
    </ethernet-tag-id>
    <mac-address>f2:0c:dd:80:9f:f7</mac-address>
    <ip-address>10.0.1.12</ip-address>
    <mpls-label1>299776</mpls-label1>
</mac-ip-adv-route>
```

- **Inclusive Multicast Ethernet Tag**

```
<inc-multi-ethernet-tag-res>
    <ethernet-tag-id>
        <vlan-id>2100</vlan-id>
```

(continues on next page)

(continued from previous page)

```
    </ethernet-tag-id>
    <orig-route-ip>43.43.43.43</orig-route-ip>
</inc-multi-ethernet-tag-res>
```

- **Ethernet Segment**

```
<es-route>
    <orig-route-ip>43.43.43.43</orig-route-ip>
    <arbitrary>
        <arbitrary>AAAAAAAAAAAA</arbitrary>
    </arbitrary>
</es-route>
```

EVPN Ethernet Segment Identifier (ESI):

- **Type 0** Indicates an arbitrary 9-octet ESI.

```
<arbitrary>
    <arbitrary>AAAAAAAAAAAA</arbitrary>
</arbitrary>
```

- **Type 1** IEEE 802.1AX LACP is used.

```
<lacp-auto-generated>
    <ce-lacp-mac-address>f2:0c:dd:80:9f:f7</ce-lacp-mac-address>
    <ce-lacp-port-key>22</ce-lacp-port-key>
</lacp-auto-generated>
```

- **Type 2** Indirectly connected hosts via a bridged LAN.

```
<lan-auto-generated>
    <root-bridge-mac-address>f2:0c:dd:80:9f:f7</root-bridge-mac-address>
    <root-bridge-priority>20</root-bridge-priority>
</lan-auto-generated>
```

- **Type 3** MAC-based ESI.

```
<mac-auto-generated>
    <system-mac-address>f2:0c:dd:80:9f:f7</system-mac-address>
    <local-discriminator>2000</local-discriminator>
</mac-auto-generated>
```

- **Type 4** Router-ID ESI

```
<router-id-generated>
    <router-id>43.43.43.43</router-id>
    <local-discriminator>2000</local-discriminator>
</router-id-generated>
```

- **Type 5** AS-based ESI

```
<as-generated>
    <as>16843009</as>
    <local-discriminator>2000</local-discriminator>
</as-generated>
```

Extended Communities:

- **ESI Label Extended Community**

```
<extended-communities>
    <transitive>true</transitive>
    <esi-label-extended-community>
        <single-active-mode>false</single-active-mode>
        <esi-label>24001</esi-label>
    </esi-label-extended-community >
</extended-communities>
```

- **ES-Import Route Target**

```
<extended-communities>
    <transitive>true</transitive>
    <es-import-route-extended-community>
        <es-import>f2:0c:dd:80:9f:f7</es-import>
    </es-import-route-extended-community>
</extended-communities>
```

- **MAC Mobility Extended Community**

```
<extended-communities>
    <transitive>true</transitive>
    <mac-mobility-extended-community>
        <static>true</static>
        <seq-number>200</seq-number>
    </mac-mobility-extended-community>
</extended-communities>
```

- **Default Gateway Extended Community**

```
<extended-communities>
    <transitive>true</transitive>
    <default-gateway-extended-community>
    </default-gateway-extended-community>
</extended-communities>
```

- **EVPN Layer 2 attributes extended community**

```
<extended-communities>
    <transitive>false</transitive>
    <layer-2-attributes-extended-community>
        <primary-pe>true</primary-pe>
        <backup-pe>true</backup-pe>
        <control-word >true</control-word>
        <l2-mtu>200</l2-mtu>
    </layer-2-attributes-extended-community>
</extended-communities>
```

- **BGP Encapsulation extended community**

```
1 <extended-communities>
2     <transitive>false</transitive>
3     <encapsulation-extended-community>
4         <tunnel-type>vxlan</tunnel-type>
5     </encapsulation-extended-community>
6 </extended-communities>
```

@line 4: full list of tunnel types

- **P-Multicast Service Interface Tunnel (PMSI) attribute**

```
<pmsi-tunnel>
    <leaf-information-required>true</leaf-information-required>
    <mpls-label>20024</mpls-label>
    <ingress-replication>
        <receiving-endpoint-address>172.12.123.3</receiving-endpoint-address>
    </ingress-replication>
</pmsi-tunnel>
```

To remove the route added above, following request can be used:

URL: `/restconf/config/bgp-rib:application-rib/10.25.1.9/tables/bgp-types:ipv4-address-family/odl-bgp-evpn:l2vpn-address-family/odl-bgp-evpn:evpn-subsequent-address-family/odl-bgp-evpn:evpn-routes/evpn-route/evpn`

Method: `DELETE`

Table 4: EVPN Routes Usage.

EVN Route Type	Extended Communities	Usage
Ethernet Auto-discovery	ESI Label, BGP EncapsulationEVPN Layer 2 attributes	Fast Convergence, Split Horizon, Aliasing
MAC/IP Advertisement	BGP Encapsulation, MAC Mobility, Default Gateway	MAC address reachability
Inclusive Multicast Ethernet Tag	PMSI Tunnel, BGP Encapsulation	Handling of Multi-destination traffic
Ethernet Segment	BGP Encapsulation, ES-Import Route Target	Designated Forwarder Election

References

- BGP MPLS-Based Ethernet VPN
- Provider Backbone Bridging Combined with Ethernet VPN
- VPWS support in EVPN
- A Network Virtualization Overlay Solution using EVPN
- Interconnect Solution for EVPN Overlay networks
- Usage and applicability of BGP MPLS based Ethernet VPN

Additional Path

The ADD-PATH capability allows to advertise multiple paths for the same address prefix. It can help with optimal routing and routing convergence in a network by providing potential alternate or backup paths.

Contents

- *Configuration*

        - *BGP Speaker*
        - *BGP Peer*
    - *Usage*
    - *References*

Configuration

This section shows a way to enable ADD-PATH capability in BGP speaker and peer configuration.

Note: The capability is applicable for IP Unicast, IP Labeled Unicast and Flow Specification address families.

BGP Speaker

To enable ADD-PATH capability in BGP plugin, first configure BGP speaker instance:

URL: `/restconf/config/openconfig-network-instance:network-instances/network-instance/global-bgp/openconfig-network-instance:protocols`

Method: `POST`

Content-Type: `application/xml`

Request Body:

```
1  <protocol xmlns="http://openconfig.net/yang/network-instance">
2      <name>bgp-example</name>
3      <identifier xmlns:x="http://openconfig.net/yang/policy-types">x:BGP</identifier>
4      <bgp xmlns="urn:opendaylight:params:xml:ns:yang:bgp:openconfig-extensions">
5          <global>
6              <config>
7                  <router-id>192.0.2.2</router-id>
8                  <as>65000</as>
9              </config>
10             <afi-safis>
11                 <afi-safi>
12                     <afi-safi-name xmlns:x="http://openconfig.net/yang/bgp-types">
   ↪x:IPV4-UNICAST</afi-safi-name>
13                     <receive>true</receive>
14                     <send-max>2</send-max>
15                 </afi-safi>
16             </afi-safis>
17         </global>
18     </bgp>
19 </protocol>
```

@line 14: Defines path selection strategy: *send-max* > 1 -> Advertise N Paths or *send-max* = 0 -> Advertise All Paths

Here is an example for update a specific family with enable ADD-PATH capability

URL: `/restconf/config/openconfig-network-instance:network-instances/network-instance/global-bgp/openconfig-network-instance:protocols/protocol/openconfig-policy-types:BGP/bgp-example/bgp/global/afi-safis/afi-safi/openconfig-bgp-types:IPV4%2DUNICAST`

Method: `PUT`

Content-Type: `application/xml`

Request Body:

```
<afi-safi xmlns="urn:opendaylight:params:xml:ns:yang:bgp:openconfig-extensions">
   <afi-safi-name xmlns:x="http://openconfig.net/yang/bgp-types">x:IPV4-UNICAST</afi-
↪safi-name>
   <receive>true</receive>
   <send-max>0</send-max>
</afi-safi>
```

BGP Peer

Here is an example for BGP peer configuration with enabled ADD-PATH capability.

URL: `/restconf/config/openconfig-network-instance:network-instances/network-instance/global-bgp/openconfig-network-instance:protocols/protocol/openconfig-policy-types:BGP/bgp-example/bgp/neighbors`

Method: `POST`

Content-Type: `application/xml`

Request Body:

```
<neighbor xmlns="urn:opendaylight:params:xml:ns:yang:bgp:openconfig-extensions">
    <neighbor-address>192.0.2.1</neighbor-address>
    <afi-safis>
        <afi-safi>
            <afi-safi-name xmlns:x="http://openconfig.net/yang/bgp-types">x:IPV4-
↪LABELLED-UNICAST</afi-safi-name>
        </afi-safi>
        <afi-safi>
            <afi-safi-name xmlns:x="http://openconfig.net/yang/bgp-types">x:IPV4-
↪UNICAST</afi-safi-name>
            <receive>true</receive>
            <send-max>0</send-max>
        </afi-safi>
    </afi-safis>
</neighbor>
```

Note: The path selection strategy is not configurable on per peer basis. The send-max presence indicates a willingness to send ADD-PATH NLRIs to the neighbor.

Here is an example for update specific family BGP peer configuration with enabled ADD-PATH capability.

URL: `/restconf/config/openconfig-network-instance:network-instances/network-instance/global-bgp/openconfig-network-instance:protocols/protocol/openconfig-policy-types:BGP/bgp-example/bgp/neighbors/neighbor/192.0.2.1/afi-safis/afi-safi/openconfig-bgp-types:IPV4%2DUNICAST`

Method: `PUT`

Content-Type: `application/xml`

Request Body:

```
<afi-safi xmlns="urn:opendaylight:params:xml:ns:yang:bgp:openconfig-extensions">
   <afi-safi-name xmlns:x="http://openconfig.net/yang/bgp-types">x:IPV4-UNICAST</afi-
↪safi-name>
   <receive>true</receive>
   <send-max>0</send-max>
</afi-safi>
```

Usage

The IPv4 Unicast table with enabled ADD-PATH capability in an instance of the speaker's Loc-RIB can be verified via REST:

URL: `/restconf/operational/bgp-rib:bgp-rib/rib/bgp-example/loc-rib/tables/bgp-types:ipv4-address-family/bgp-types:unicast-subsequent-address-family/ipv4-routes`

Method: `GET`

Response Body:

```
1  <ipv4-routes xmlns="urn:opendaylight:params:xml:ns:yang:bgp-inet">
2      <ipv4-route>
3          <path-id>1</path-id>
4          <prefix>193.0.2.1/32</prefix>
5          <attributes>
6              <as-path></as-path>
7              <origin>
8                  <value>igp</value>
9              </origin>
10             <local-pref>
11                 <pref>100</pref>
12             </local-pref>
13             <ipv4-next-hop>
14                 <global>10.0.0.1</global>
15             </ipv4-next-hop>
16         </attributes>
17     </ipv4-route>
18     <ipv4-route>
19         <path-id>2</path-id>
20         <prefix>193.0.2.1/32</prefix>
21         <attributes>
22             <as-path></as-path>
23             <origin>
24                 <value>igp</value>
25             </origin>
26             <local-pref>
27                 <pref>100</pref>
28             </local-pref>
29             <ipv4-next-hop>
30                 <global>10.0.0.2</global>
31             </ipv4-next-hop>
32         </attributes>
33     </ipv4-route>
34 </ipv4-routes>
```

@line 3: The routes with the same destination are distinguished by *path-id* attribute.

References

- Advertisement of Multiple Paths in BGP
- Best Practices for Advertisement of Multiple Paths in IBGP

Route Refresh

The Route Refresh Capability allows to dynamically request a re-advertisement of the Adj-RIB-Out from a BGP peer. This is useful when the inbound routing policy for a peer changes and all prefixes from a peer must be reexamined against a new policy.

Contents

- *Configuration*
- *Usage*
- *References*

Configuration

The capability is enabled by default, no additional configuration is required.

Usage

To send a Route Refresh request from OpenDaylight BGP speaker instance to its neighbor, invoke RPC:

URL: `/restconf/operations/bgp-peer-rpc:route-refresh-request`

Method: `POST`

Content-Type: `application/xml`

Request Body:

```
<input xmlns="urn:opendaylight:params:xml:ns:yang:bgp-peer-rpc">
    <afi xmlns:types="urn:opendaylight:params:xml:ns:yang:bgp-types">types:ipv4-
↪address-family</afi>
    <safi xmlns:types="urn:opendaylight:params:xml:ns:yang:bgp-types">types:unicast-
↪subsequent-address-family</safi>
    <peer-ref xmlns:rib="urn:opendaylight:params:xml:ns:yang:bgp-rib">/rib:bgp-rib/
↪rib:rib[rib:id="bgp-example"]/rib:peer[rib:peer-id="bgp://10.25.1.9"]</peer-ref>
</input>
```

References

- Route Refresh Capability for BGP-4

High Availability

Running OpenDaylight BGP in clustered environment brings an advantage of the plugin's high availability (HA). This section illustrates a basic scenario for HA, also presents a configuration for clustered OpenDaylight BGP.

Contents

- *Configuration*
- *Failover scenario*

Configuration

Following example shows a configuration for running BGP in clustered environment.

1. As the first step, configure (replicated *deafult* shard and *topology* shard if needed) and run OpenDaylight in clustered environment, install BGP and RESTCONF.
2. On one node (OpenDaylight instance), configure BGP speaker instance and neighbor. In addition, configure BGP topology exporter if required. The configuration is shared across all interconnected cluster nodes, however BGP become active only on one node. Other nodes with configured BGP serves as stand-by backups.
3. Remote peer should be configured to accept/initiate connection from/to all OpenDaylight cluster nodes with configured BGP plugin.
4. Connect remote peer, let it advertise some routes. Verify routes presence in Loc-RIB and/or BGP topology exporter instance on all nodes of the OpenDaylight cluster.

Warning: Replicating RIBs across the cluster nodes is causing severe scalability issue and overall performance degradation. To avoid this problems, configure BGP RIB module as a separate shard without enabled replication. Change configuration on all nodes as following (*configuration/initial*):

- **In `modules.conf` add a new module:**

```
{
    name = "bgp_rib"
    namespace = "urn:opendaylight:params:xml:ns:yang:bgp-rib"
    shard-strategy = "module"
}
```

- **In `module-shards.conf` define a new module shard:**

```
{
    name = "bgp_rib"
    shards = [
        {
            name="bgp_rib"
            replicas = [
                "member-1"
            ]
        }
    ]
}
```

Note: Use correct member name in module shard configuration.

Failover scenario

Following section presents a basic BGP speaker failover scenario on 3-node OpenDaylight cluster setup.

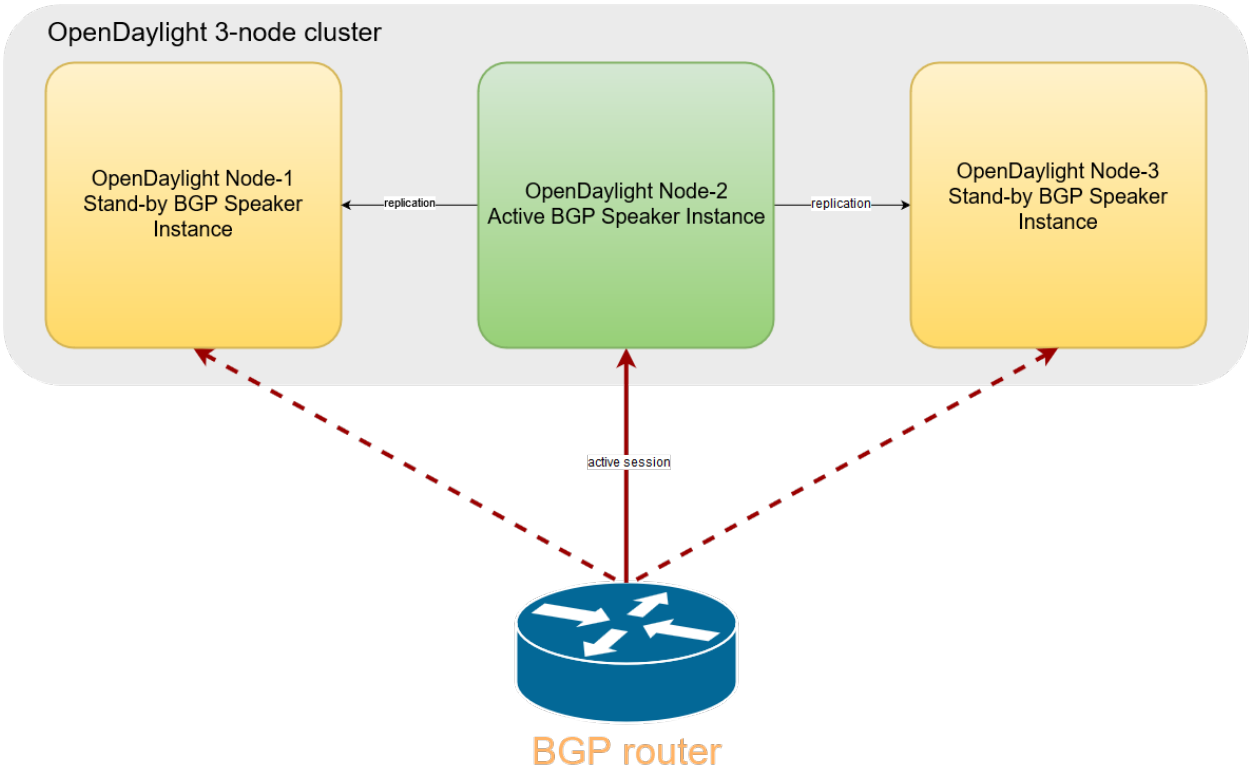

Fig. 21: Once the OpenDaylight BGP is configured, the speaker become active on one of the cluster nodes. Remote peer can establish connection with this BGP instance. Routes advertised by remote peer are replicated, hence RIBs state on all nodes is the same.

Topology Provider

This section provides an overview of the BGP topology provider service. It shows how to configure and use all available BGP topology providers. Providers are building topology view of BGP routes stored in local BGP speaker's Loc-RIB. Output topologies are rendered in a form of standardised IETF network topology model.

Contents

- *Inet Reachability Topology*
  - *Configuration*
  - *Usage*
- *BGP Linkstate Topology*

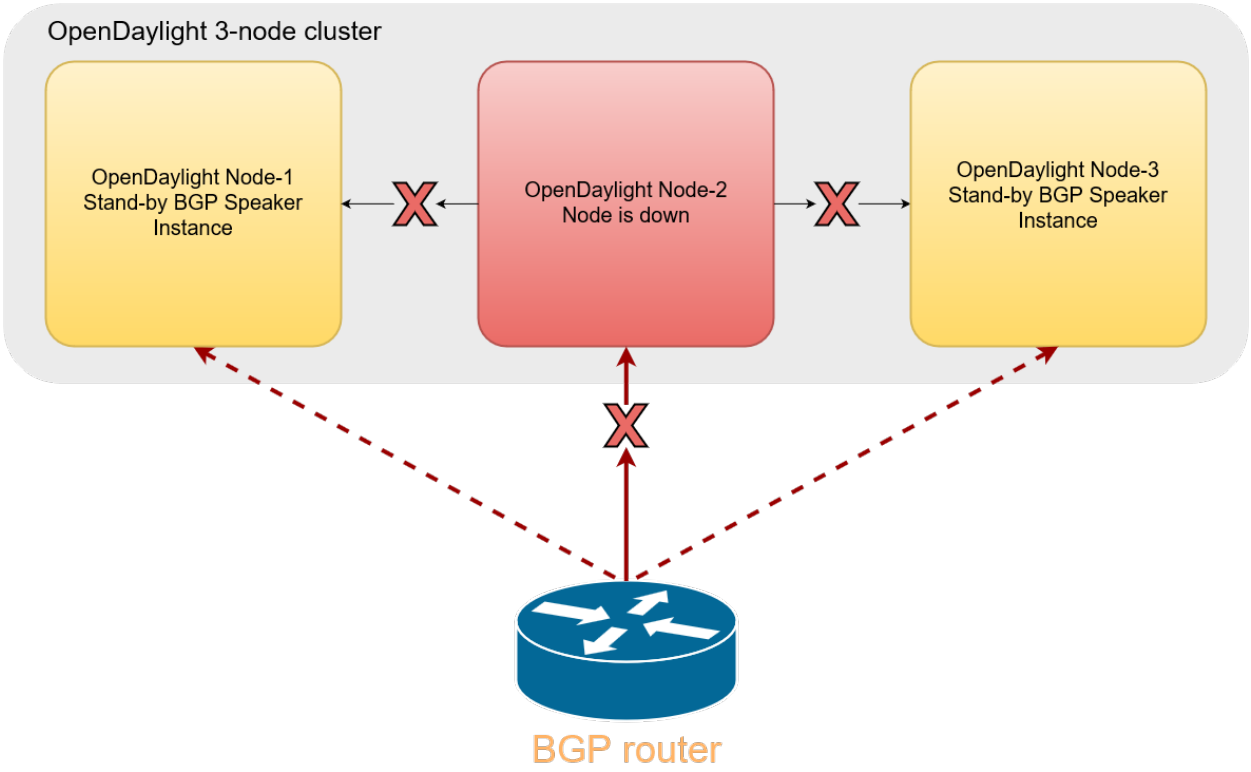

Fig. 22: In a case a cluster node, where BGP instance is running, goes down (unexpected failure, restart), active BGP session is dropped.

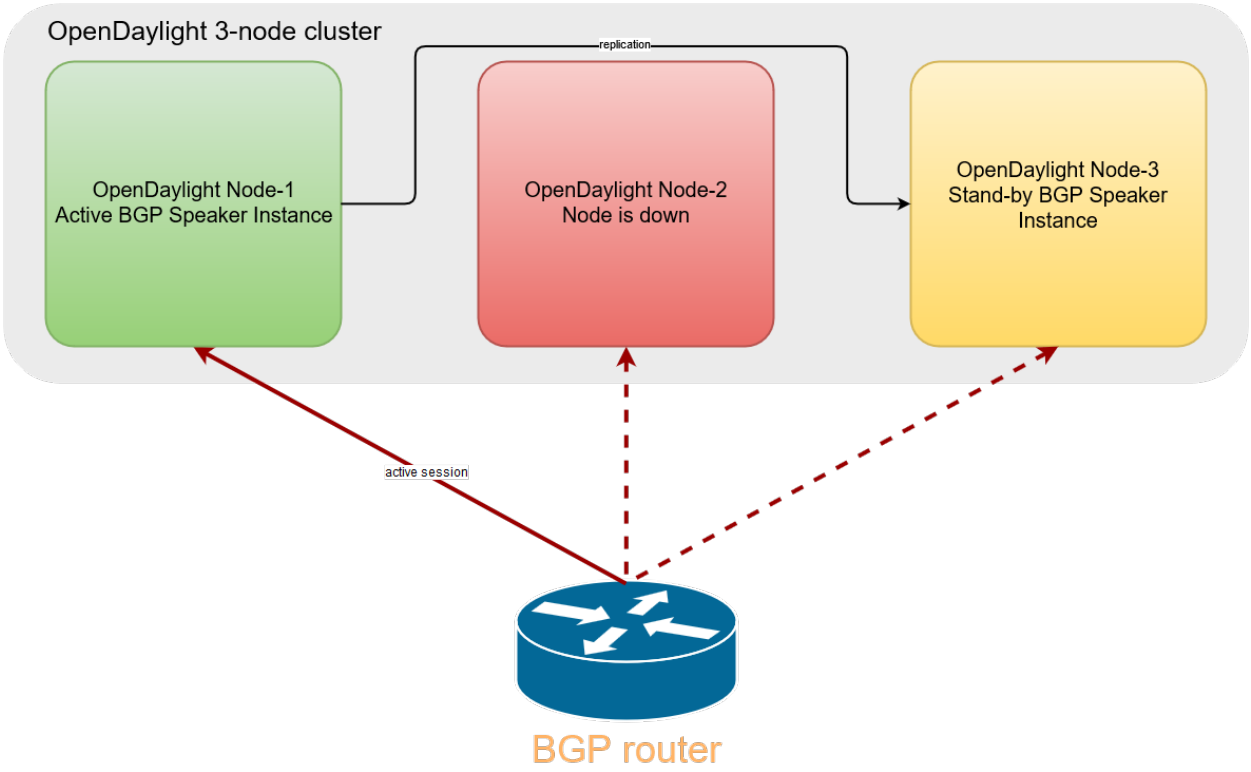

Fig. 23: Now, one of the stand-by BGP speaker instances become active. Remote peer establishes new connection and advertises routes again.

- *Configuration*
- *Usage*

Inet Reachability Topology

Inet reachability topology exporter offers a mapping service from IPv4/6 routes to network topology nodes.

Configuration

Following example shows how to create a new instance of IPv4 BGP topology exporter:

URL: `/restconf/config/network-topology:network-topology`

Method: `POST`

Content-Type: `application/xml`

Request Body:

```
1 <topology xmlns="urn:TBD:params:xml:ns:yang:network-topology">
2     <topology-id>bgp-example-ipv4-topology</topology-id>
3     <topology-types>
4         <bgp-ipv4-reachability-topology xmlns=
  ↪"urn:opendaylight:params:xml:ns:yang:odl-bgp-topology-types"></bgp-ipv4-
  ↪reachability-topology>
5     </topology-types>
6     <rib-id xmlns="urn:opendaylight:params:xml:ns:yang:odl-bgp-topology-config">bgp-
  ↪example</rib-id>
7 </topology>
```

@line 2: An identifier for a topology.

@line 4: Used to identify type of the topology. In this case BGP IPv4 reachability topology.

@line 6: A name of the local BGP speaker instance.

The topology exporter instance can be removed in a following way:

URL: `/restconf/config/network-topology:network-topology/topology/bgp-example-ipv4-topology`

Method: `DELETE`

Following example shows how to create a new instance of IPv6 BGP topology exporter:

URL: `/restconf/config/network-topology:network-topology`

Method: `POST`

Content-Type: `application/xml`

Request Body:

```
<topology xmlns="urn:TBD:params:xml:ns:yang:network-topology">
    <topology-id>bgp-example-ipv6-topology</topology-id>
    <topology-types>
        <bgp-ipv6-reachability-topology xmlns=
↪"urn:opendaylight:params:xml:ns:yang:odl-bgp-topology-types"></bgp-ipv6-
↪reachability-topology>
    </topology-types>
    <rib-id xmlns="urn:opendaylight:params:xml:ns:yang:odl-bgp-topology-config">bgp-
↪example</rib-id>
</topology>
```

Usage

Operational state of the topology can be verified via REST:

URL: `/restconf/operational/network-topology:network-topology/topology/bgp-example-ipv4-topology`

Method: `GET`

Response Body:

```
1  <topology xmlns="urn:TBD:params:xml:ns:yang:network-topology">
2      <topology-id>bgp-example-ipv4-topology</topology-id>
3      <server-provided>true</server-provided>
4      <topology-types>
5          <bgp-ipv4-reachability-topology xmlns=
↪"urn:opendaylight:params:xml:ns:yang:odl-bgp-topology-types"></bgp-ipv4-
↪reachability-topology>
6      </topology-types>
7      <node>
8          <node-id>10.10.1.1</node-id>
9          <igp-node-attributes xmlns="urn:TBD:params:xml:ns:yang:nt:l3-unicast-igp-
↪topology">
10             <prefix>
11                 <prefix>10.0.0.10/32</prefix>
12             </prefix>
13         </igp-node-attributes>
14     </node>
15 </topology>
```

@line 8: The identifier of a node in a topology. Its value is mapped from route's NEXT_HOP attribute.

@line 11: The IP prefix attribute of the node. Its value is mapped from routes's destination IP prefix.

BGP Linkstate Topology

BGP linkstate topology exporter offers a mapping service from BGP-LS routes to network topology nodes and links.

Configuration

Following example shows how to create a new instance of linkstate BGP topology exporter:

URL: `/restconf/config/network-topology:network-topology`

Method: `POST`

Content-Type: `application/xml`

Request Body:

```
<topology  xmlns="urn:TBD:params:xml:ns:yang:network-topology">
    <topology-id>bgp-example-linkstate-topology</topology-id>
    <topology-types>
        <bgp-linkstate-topology xmlns="urn:opendaylight:params:xml:ns:yang:odl-bgp-
↪topology-types"></bgp-linkstate-topology>
    </topology-types>
    <rib-id xmlns="urn:opendaylight:params:xml:ns:yang:odl-bgp-topology-config">bgp-
↪example</rib-id>
</topology>
```

Usage

Operational state of the topology can be verified via REST. A sample output below represents a two node topology with two unidirectional links interconnecting those nodes.

URL: `/restconf/operational/network-topology:network-topology/topology/bgp-example-linkstate-topology`

Method: `GET`

Response Body:

```
<topology xmlns="urn:TBD:params:xml:ns:yang:network-topology">
    <topology-id>bgp-example-linkstate-topology</topology-id>
    <server-provided>true</server-provided>
    <topology-types>
        <bgp-linkstate-topology xmlns="urn:opendaylight:params:xml:ns:yang:odl-bgp-
↪topology-types"></bgp-linkstate-topology>
    </topology-types>
    <node>
        <node-id>bgpls://IsisLevel2:1/type=node&amp;as=65000&amp;domain=673720360&amp;
↪router=0000.0000.0040</node-id>
        <termination-point>
            <tp-id>bgpls://IsisLevel2:1/type=tp&amp;ipv4=203.20.160.40</tp-id>
            <igp-termination-point-attributes xmlns="urn:TBD:params:xml:ns:yang:nt:l3-
↪unicast-igp-topology"/>
        </termination-point>
        <igp-node-attributes xmlns="urn:TBD:params:xml:ns:yang:nt:l3-unicast-igp-
↪topology">
            <prefix>
                <prefix>40.40.40.40/32</prefix>
                <metric>10</metric>
            </prefix>
            <prefix>
                <prefix>203.20.160.0/24</prefix>
                <metric>10</metric>
            </prefix>
            <name>node1</name>
            <router-id>40.40.40.40</router-id>
            <isis-node-attributes xmlns="urn:TBD:params:xml:ns:yang:network:isis-
↪topology">
                <ted>
```

(continues on next page)

(continued from previous page)

```
                    <te-router-id-ipv4>40.40.40.40</te-router-id-ipv4>
                </ted>
                <iso>
                    <iso-system-id>MDAwMDAwMDAwMDY0</iso-system-id>
                </iso>
            </isis-node-attributes>
        </igp-node-attributes>
    </node>
    <node>
        <node-id>bgpls://IsisLevel2:1/type=node&amp;as=65000&amp;domain=673720360&amp;
↪router=0000.0000.0039</node-id>
        <termination-point>
            <tp-id>bgpls://IsisLevel2:1/type=tp&amp;ipv4=203.20.160.39</tp-id>
            <igp-termination-point-attributes xmlns="urn:TBD:params:xml:ns:yang:nt:l3-
↪unicast-igp-topology"/>
        </termination-point>
        <igp-node-attributes xmlns="urn:TBD:params:xml:ns:yang:nt:l3-unicast-igp-
↪topology">
            <prefix>
                <prefix>39.39.39.39/32</prefix>
                <metric>10</metric>
            </prefix>
            <prefix>
                <prefix>203.20.160.0/24</prefix>
                <metric>10</metric>
            </prefix>
            <name>node2</name>
            <router-id>39.39.39.39</router-id>
            <isis-node-attributes xmlns="urn:TBD:params:xml:ns:yang:network:isis-
↪topology">
                <ted>
                    <te-router-id-ipv4>39.39.39.39</te-router-id-ipv4>
                </ted>
                <iso>
                    <iso-system-id>MDAwMDAwMDAwMDg3</iso-system-id>
                </iso>
            </isis-node-attributes>
        </igp-node-attributes>
    </node>
    <link>
        <destination>
            <dest-node>bgpls://IsisLevel2:1/type=node&amp;as=65000&amp;
↪domain=673720360&amp;router=0000.0000.0039</dest-node>
            <dest-tp>bgpls://IsisLevel2:1/type=tp&amp;ipv4=203.20.160.39</dest-tp>
        </destination>
        <link-id>bgpls://IsisLevel2:1/type=link&amp;local-as=65000&amp;local-
↪domain=673720360&amp;local-router=0000.0000.0040&amp;remote-as=65000&amp;remote-
↪domain=673720360&amp;remote-router=0000.0000.0039&amp;ipv4-iface=203.20.160.40&amp;
↪ipv4-neigh=203.20.160.39</link-id>
        <source>
            <source-node>bgpls://IsisLevel2:1/type=node&amp;as=65000&amp;
↪domain=673720360&amp;router=0000.0000.0040</source-node>
            <source-tp>bgpls://IsisLevel2:1/type=tp&amp;ipv4=203.20.160.40</source-tp>
        </source>
        <igp-link-attributes xmlns="urn:TBD:params:xml:ns:yang:nt:l3-unicast-igp-
↪topology">
            <metric>10</metric>
```

(continues on next page)

(continued from previous page)

```
            <isis-link-attributes xmlns="urn:TBD:params:xml:ns:yang:network:isis-
↪topology">
                <ted>
                    <color>0</color>
                    <max-link-bandwidth>1250000.0</max-link-bandwidth>
                    <max-resv-link-bandwidth>12500.0</max-resv-link-bandwidth>
                    <te-default-metric>0</te-default-metric>
                    <unreserved-bandwidth>
                        <bandwidth>12500.0</bandwidth>
                        <priority>0</priority>
                    </unreserved-bandwidth>
                    <unreserved-bandwidth>
                        <bandwidth>12500.0</bandwidth>
                        <priority>1</priority>
                    </unreserved-bandwidth>
                    <unreserved-bandwidth>
                        <bandwidth>12500.0</bandwidth>
                        <priority>2</priority>
                    </unreserved-bandwidth>
                    <unreserved-bandwidth>
                        <bandwidth>12500.0</bandwidth>
                        <priority>3</priority>
                    </unreserved-bandwidth>
                    <unreserved-bandwidth>
                        <bandwidth>12500.0</bandwidth>
                        <priority>4</priority>
                    </unreserved-bandwidth>
                    <unreserved-bandwidth>
                        <bandwidth>12500.0</bandwidth>
                        <priority>5</priority>
                    </unreserved-bandwidth>
                    <unreserved-bandwidth>
                        <bandwidth>12500.0</bandwidth>
                        <priority>6</priority>
                    </unreserved-bandwidth>
                    <unreserved-bandwidth>
                        <bandwidth>12500.0</bandwidth>
                        <priority>7</priority>
                    </unreserved-bandwidth>
                </ted>
            </isis-link-attributes>
        </igp-link-attributes>
    </link>
    <link>
        <destination>
            <dest-node>bgpls://IsisLevel2:1/type=node&amp;as=65000&amp;
↪domain=673720360&amp;router=0000.0000.0040</dest-node>
            <dest-tp>bgpls://IsisLevel2:1/type=tp&amp;ipv4=203.20.160.40</dest-tp>
        </destination>
        <link-id>bgpls://IsisLevel2:1/type=link&amp;local-as=65000&amp;local-
↪domain=673720360&amp;local-router=0000.0000.0039&amp;remote-as=65000&amp;remote-
↪domain=673720360&amp;remote-router=0000.0000.0040&amp;ipv4-iface=203.20.160.39&amp;
↪ipv4-neigh=203.20.160.40</link-id>
        <source>
            <source-node>bgpls://IsisLevel2:1/type=node&amp;as=65000&amp;
↪domain=673720360&amp;router=0000.0000.0039</source-node>
            <source-tp>bgpls://IsisLevel2:1/type=tp&amp;ipv4=203.20.160.39</source-tp>
```

(continues on next page)

(continued from previous page)

```
        </source>
        <igp-link-attributes xmlns="urn:TBD:params:xml:ns:yang:nt:l3-unicast-igp-
↪topology">
            <metric>10</metric>
            <isis-link-attributes xmlns="urn:TBD:params:xml:ns:yang:network:isis-
↪topology">
                <ted>
                    <color>0</color>
                    <max-link-bandwidth>1250000.0</max-link-bandwidth>
                    <max-resv-link-bandwidth>12500.0</max-resv-link-bandwidth>
                    <te-default-metric>0</te-default-metric>
                    <unreserved-bandwidth>
                        <bandwidth>12500.0</bandwidth>
                        <priority>0</priority>
                    </unreserved-bandwidth>
                    <unreserved-bandwidth>
                        <bandwidth>12500.0</bandwidth>
                        <priority>1</priority>
                    </unreserved-bandwidth>
                    <unreserved-bandwidth>
                        <bandwidth>12500.0</bandwidth>
                        <priority>2</priority>
                    </unreserved-bandwidth>
                    <unreserved-bandwidth>
                        <bandwidth>12500.0</bandwidth>
                        <priority>3</priority>
                    </unreserved-bandwidth>
                    <unreserved-bandwidth>
                        <bandwidth>12500.0</bandwidth>
                        <priority>4</priority>
                    </unreserved-bandwidth>
                    <unreserved-bandwidth>
                        <bandwidth>12500.0</bandwidth>
                        <priority>5</priority>
                    </unreserved-bandwidth>
                    <unreserved-bandwidth>
                        <bandwidth>12500.0</bandwidth>
                        <priority>6</priority>
                    </unreserved-bandwidth>
                    <unreserved-bandwidth>
                        <bandwidth>12500.0</bandwidth>
                        <priority>7</priority>
                    </unreserved-bandwidth>
                </ted>
            </isis-link-attributes>
        </igp-link-attributes>
    </link>
</topology>
```

Test Tools

BGP test tools serves to test basic BGP functionality, scalability and performance.

Contents

- *BGP Test Tool*
  - *Usage*
- *BGP Application Peer Benchmark*
  - *Configuration*
  - *Inject routes*
  - *Remove routes*

BGP Test Tool

The BGP Test Tool is a stand-alone Java application purposed to simulate remote BGP peers, that are capable to advertise sample routes. This application is not part of the OpenDaylight Karaf distribution, however it can be downloaded from OpenDaylight's Nexus (use latest release version):

`https://nexus.opendaylight.org/content/repositories/opendaylight.release/org/opendaylight/bgpcep/bgp-testtool`

Usage

The application can be run from command line:

```
java -jar bgp-testtool-*-executable.jar
```

with optional input parameters:

```
-i <BOOLEAN>, --active <BOOLEAN>
   Active initialisation of the connection, by default false.

-ho <N>, --holdtimer <N>
   In seconds, value of the desired holdtimer, by default 90.

-sc <N>, --speakersCount <N>
   Number of simulated BGP speakers, when creating each speaker, uses incremented␣
↪local-address for binding, by default 0.

-ra <IP_ADDRESS:PORT,...>, --remoteAddress <IP_ADDRESS:PORT,...>
   A list of IP addresses of remote BGP peers, that the tool can accept or initiate␣
↪connect to that address (based on the mode), by default 192.0.2.2:1790.

-la <IP_ADDRESS:PORT>, --localAddress <IP_ADDRESS:PORT>
   IP address of BGP speakers which the tools simulates, by default 192.0.2.2:0.

-pr <N>, --prefixes <N>
   Number of prefixes to be advertised by each simulated speaker

-mp <BOOLEAN>, --multiPathSupport <BOOLEAN>
   Active ADD-PATH support, by default false.

-as <N>, --as <N>
   Local AS Number, by default 64496.
```

(continues on next page)

(continued from previous page)

```
-ec <EXTENDED_COMMUNITIES>, --extended_communities <EXTENDED_COMMUNITIES>
   Extended communities to be send. Format: x,x,x  where  x  is  each  extended␣
↪community from bgp-types.yang, by default empty.

-ll <LOG_LEVEL>, --log_level <LOG_LEVEL>
   Log level for console output, by default INFO.
```

BGP Application Peer Benchmark

It is a simple OpenDaylight application which is capable to inject and remove specific amount of IPv4 routes. This application is part of the OpenDaylight Karaf distribution.

Configuration

As a first step install BGP, RESTCONF and NETCONF connector plugin, then configure *Application Peer*. Install `odl-bgpcep-bgp-benchmark` feature and reconfigure BGP Application Peer Benchmark application as per following:

URL: `/restconf/config/network-topology:network-topology/topology/topology-netconf/node/controller-config/yang-ext:mount/config:modules/module/odl-bgp-benchmark-cfg:app-peer-benchmark/bgp-app-peer-benchmark`

Method: `PUT`

Content-Type: `application/xml`

Request Body:

```
1  <module xmlns="urn:opendaylight:params:xml:ns:yang:controller:config">
2      <type xmlns:x="urn:opendaylight:params:xml:ns:yang:controller:odl-bgp-benchmark-
   ↪cfg">x:app-peer-benchmark</type>
3      <name>bgp-app-peer-benchmark</name>
4      <binding-data-broker xmlns="urn:opendaylight:params:xml:ns:yang:controller:odl-
   ↪bgp-benchmark-cfg">
5          <type xmlns:x="urn:opendaylight:params:xml:ns:yang:controller:md:sal:binding">
   ↪x:binding-async-data-broker</type>
6          <name>pingpong-binding-data-broker</name>
7      </binding-data-broker>
8      <rpc-registry xmlns="urn:opendaylight:params:xml:ns:yang:controller:odl-bgp-
   ↪benchmark-cfg">
9          <type xmlns:x="urn:opendaylight:params:xml:ns:yang:controller:md:sal:binding">
   ↪x:binding-rpc-registry</type>
10         <name>binding-rpc-broker</name>
11     </rpc-registry>
12     <app-rib-id xmlns="urn:opendaylight:params:xml:ns:yang:controller:odl-bgp-
   ↪benchmark-cfg">10.25.1.9</app-rib-id>
13 </module>
```

@line 12: The *Application Peer* identifier.

Warning: This configuration will be moved to configuration datastore in Carbon release.

Inject routes

Routes injection can be invoked via RPC:

URL: `/restconf/operations/odl-bgp-app-peer-benchmark:add-prefix`

Method: `POST`

Content-Type: `application/xml`

Request Body:

```
<input xmlns="urn:opendaylight:params:xml:ns:yang:odl-bgp-app-peer-benchmark">
    <prefix>1.1.1.1/32</prefix>
    <count>100000</count>
    <batchsize>2000</batchsize>
    <nexthop>192.0.2.2</nexthop>
</input>
```

@line 2: A initial IPv4 prefix carried in route. Value is incremented for following routes.

@line 3: An amount of routes to be added to *Application Peer's* programmable RIB.

@line 4: A size of the transaction batch.

@line 5: A NEXT_HOP attribute value used in all injected routes.

Response Body:

```
<output xmlns="urn:opendaylight:params:xml:ns:yang:odl-bgp-app-peer-benchmark">
    <result>
        <duration>4301</duration>
        <rate>25000</rate>
        <count>100000</count>
    </result>
</output>
```

@line 3: Request duration in milliseconds.

@line 4: Writes per second rate.

@line 5: An amount of routes added to *Application Peer's* programmable RIB.

Remove routes

Routes deletion can be invoked via RPC:

URL: `/restconf/operations/odl-bgp-app-peer-benchmark:delete-prefix`

Method: `POST`

Content-Type: `application/xml`

Request Body:

```
<input xmlns="urn:opendaylight:params:xml:ns:yang:odl-bgp-app-peer-benchmark">
    <prefix>1.1.1.1/32</prefix>
    <count>100000</count>
    <batchsize>2000</batchsize>
</input>
```

@line 2: A initial IPv4 prefix carried in route to be removed. Value is incremented for following routes.

@line 3: An amount of routes to be removed from *Application Peer's* programmable RIB.

@line 4: A size of the transaction batch.

Response Body:

```
<output xmlns="urn:opendaylight:params:xml:ns:yang:odl-bgp-app-peer-benchmark">
    <result>
        <duration>1837</duration>
        <rate>54500</rate>
        <count>100000</count>
    </result>
</output>
```

Troubleshooting

This section offers advices in a case OpenDaylight BGP plugin is not working as expected.

Contents

- *BGP is not working...*
- *Bug reporting*

BGP is not working...

- First of all, ensure that all required features are installed, local and remote BGP configuration is correct.
- Check OpenDaylight Karaf logs:

From Karaf console:

```
log:tail
```

or open log file: `data/log/karaf.log`

Possibly, a reason/hint for a cause of the problem can be found there.

- Try to minimise effect of other OpenDaylight features, when searching for a reason of the problem.
- Try to set DEBUG severity level for BGP logger via Karaf console commands, in order to collect more information:

```
log:set DEBUG org.opendaylight.protocol.bgp
```

```
log:set DEBUG org.opendaylight.bgpcep.bgp
```

Bug reporting

Before you report a bug, check BGPCEP Bugzilla to ensure same/similar bug is not already filed there.

Write an e-mail to bgpcep-users@lists.opendaylight.org and provide following information:

1. State OpenDaylight version
2. Describe your use-case and provide as much details related to BGP as possible
3. Steps to reproduce
4. Attach Karaf log files, optionally packet captures, REST input/output

BGP Monitoring Protocol User Guide

This guide contains information on how to use the OpenDaylight BGP Monitoring Protocol (BMP) plugin. It covers BMP basic concepts, supported capabilities, configuration and operations.

Contents

- *Overview*
- *Running BMP*
- *BMP Monitoring Station*
- *Test tools*
- *Troubleshooting*

Overview

This section provides high-level overview of the BMP plugin, OpenDaylight implementation and BMP usage for SDN.

Contents

- *BGP Monitoring Protocol*
- *BMP in SDN*
- *OpenDaylight BMP plugin*
- *List of supported capabilities*

BGP Monitoring Protocol

The BGP Monitoring Protocol (BMP) serves to monitor BGP sessions. The BMP can be used to obtain route view instead of screen scraping. The BMP provides access to unprocessed routing information (Adj-RIB-In) and processed routes (applied inbound policy) of monitored router's peer. In addition, monitored router can provide periodic dump of statistics.

The BMP runs over TCP. Both monitored router and monitoring station can be configured as active or passive party of the connection. The passive party listens at particular port. The router can be monitored by multiple monitoring stations. BMP messages are sent by monitored router only, monitoring station supposed to collect and process data received over BMP.

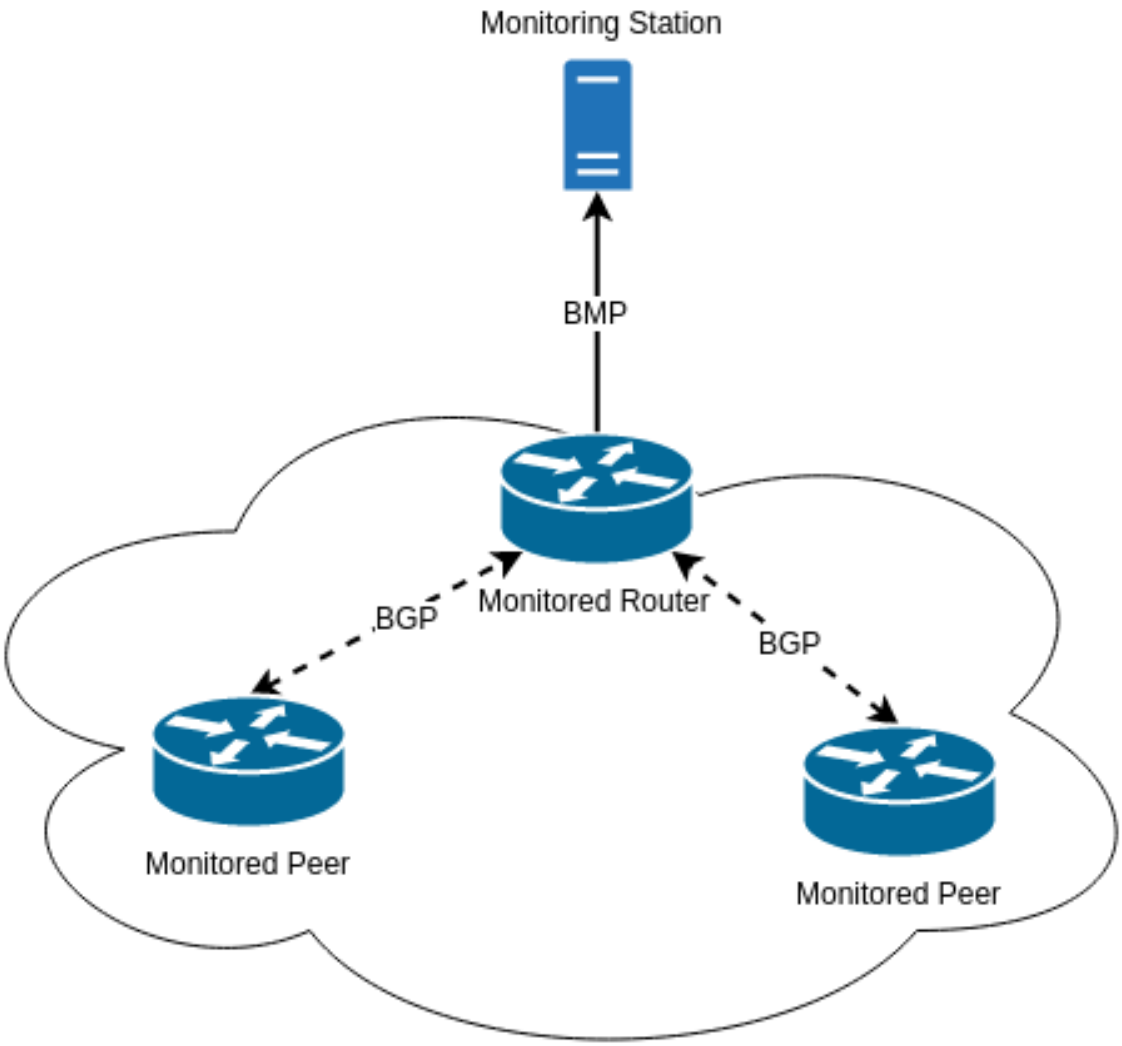

Fig. 24: The BMP overview - Monitoring Station, Monitored Router and Monitored Peers.

BMP in SDN

The main concept of BMP is to monitor BGP sessions - monitoring station is aware of monitored peer's status, collects statistics and analyzes them in order to provide valuable information for network operators.

Moreover, BMP provides provides peer RIBs visibility, without need to establish BGP sessions. Unprocessed routes may serve as a source of information for software-driven routing optimization. In this case, SDN controller, a BMP monitoring station, collects routing information from monitored routers. The routes are used in subsequent optimization procedures.

OpenDaylight BMP plugin

The OpenDaylight BMP plugin provides monitoring station implementation. The plugin can establish BMP session with one or more monitored routers in order to collect routing and statistical information.

- Runtime configurable monitoring station
- Read-only routes and statistics view
- Supports various routing information types

Important: The BMP plugin is not storing historical data, it provides current snapshot only.

List of supported capabilities

The BMP plugin implementation is based on Internet standards:

- RFC7854 - BGP Monitoring Protocol (BMP)

Note: The BMP plugin is capable to process various types of routing information (IP Unicast, EVPN, L3VPN, Link-State,...). Please, see complete list in BGP user guide.

Running BMP

This section explains how to install BMP plugin.

1. Install BMP feature - `odl-bgpcep-bmp`. Also, for sake of this sample, it is required to install RESTCONF. In the Karaf console, type command:

   ```
   feature:install odl-restconf odl-bgpcep-bmp
   ```

2. The BMP plugin contains a default configuration, which is applied after the feature starts up. One instance of BMP monitoring station is created (named *example-bmp-monitor*), and its presence can be verified via REST:

   URL: `/restconf/operational/bmp-monitor:bmp-monitor/monitor/example-bmp-monitor`

   Method: `GET`

   Response Body:

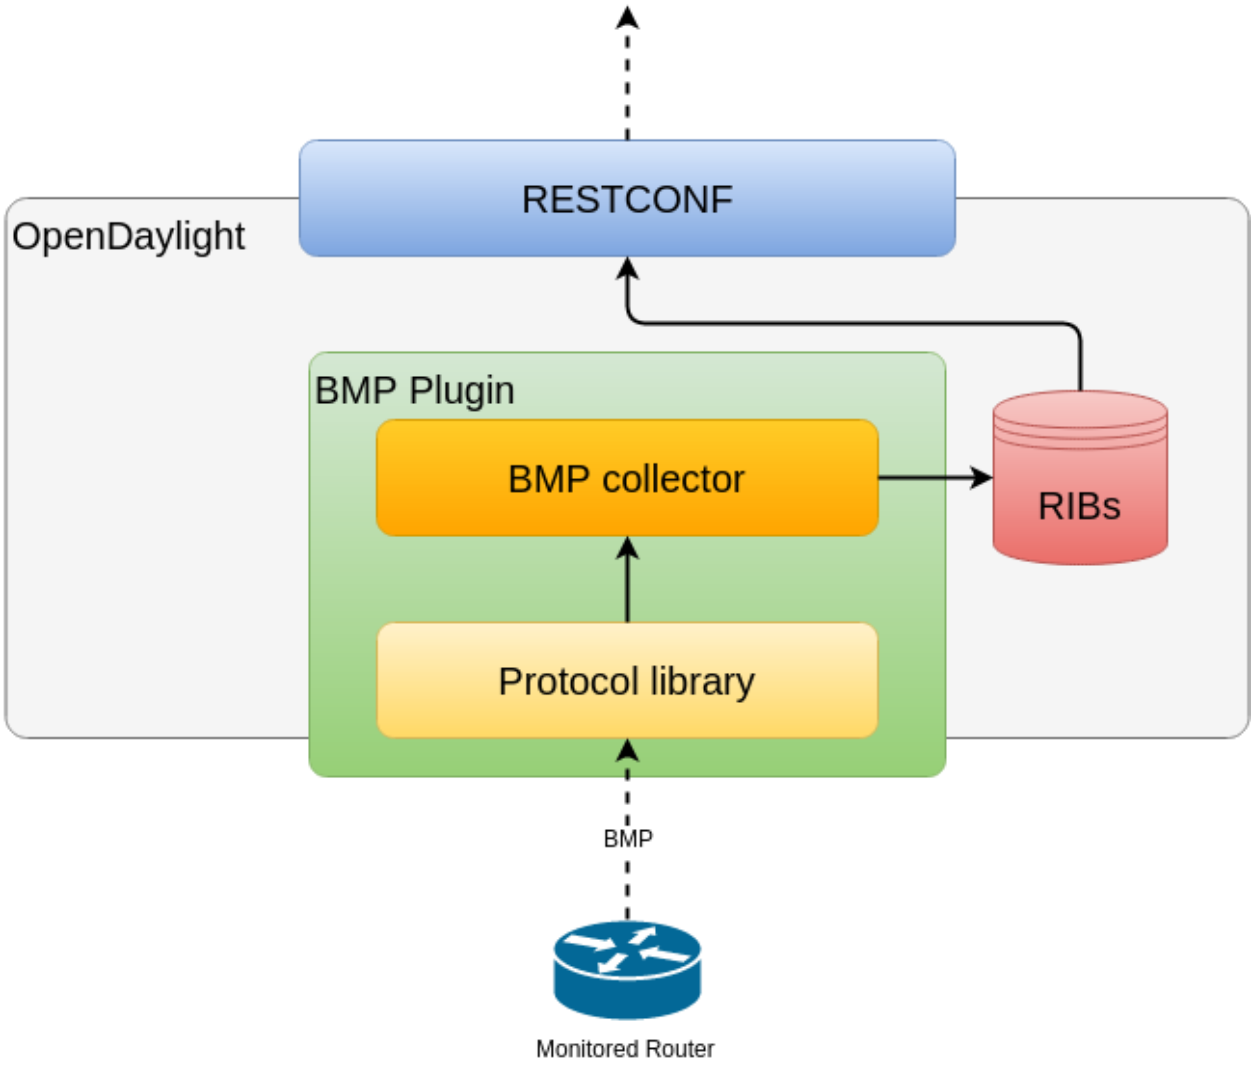

Fig. 25: OpenDaylight BMP plugin overview.

```
<monitor xmlns="urn:opendaylight:params:xml:ns:yang:bmp-monitor">
    <monitor-id>example-bmp-monitor</monitor-id>
</monitor>
```

BMP Monitoring Station

The following section shows how to configure BMP basics, how to verify functionality and presents essential components of the plugin. Next samples demonstrate the plugin's runtime configuration capability.

The monitoring station is responsible for received BMP PDUs processing and storage. The default BMP server is listening at port *12345*.

Contents

- *Configuration*
  - *Monitoring station configuration*
  - *Active mode configuration*
  - *MD5 authentication configuration*
- *Collector DB Tree*
- *Operations*

Configuration

This section shows the way to configure the BMP monitoring station via REST API.

Warning: The BMP monitoring station configuration is going to be changed in Carbon. This user-guide will be updated accordingly.

Monitoring station configuration

In order to change default's BMP monitoring station configuration, use following request. It is required to install `odl-netconf-connector-ssh` feature first.

URL: `/restconf/config/network-topology:network-topology/topology/topology-netconf/node/controller-config/yang-ext:mount/config:modules/config:module/odl-bmp-impl-cfg:bmp-monitor-impl/example-bmp-monitor`

Method: `PUT`

Content-Type: `application/xml`

Request Body:

```
1 <module xmlns="urn:opendaylight:params:xml:ns:yang:controller:config">
2   <name>example-bmp-monitor</name>
3   <type xmlns:x="urn:opendaylight:params:xml:ns:yang:controller:bmp:impl">x:bmp-
↪monitor-impl</type>
```

(continues on next page)

(continued from previous page)

```
4    <binding-port xmlns="urn:opendaylight:params:xml:ns:yang:controller:bmp:impl">12355
↪</binding-port>
5    <binding-address xmlns="urn:opendaylight:params:xml:ns:yang:controller:bmp:impl">0.
↪0.0.0</binding-address>
6    <bmp-dispatcher xmlns="urn:opendaylight:params:xml:ns:yang:controller:bmp:impl">
7      <type>bmp-dispatcher</type>
8      <name>global-bmp-dispatcher</name>
9    </bmp-dispatcher>
10   <codec-tree-factory xmlns="urn:opendaylight:params:xml:ns:yang:controller:bmp:impl
↪">
11     <type xmlns:x="urn:opendaylight:params:xml:ns:yang:controller:md:sal:binding">
↪x:binding-codec-tree-factory</type>
12     <name>runtime-mapping-singleton</name>
13   </codec-tree-factory>
14   <extensions xmlns="urn:opendaylight:params:xml:ns:yang:controller:bmp:impl">
15     <type xmlns:x="urn:opendaylight:params:xml:ns:yang:controller:bgp:rib:spi">
↪x:extensions</type>
16     <name>global-rib-extensions</name>
17   </extensions>
18   <dom-data-provider xmlns="urn:opendaylight:params:xml:ns:yang:controller:bmp:impl">
19     <type xmlns:x="urn:opendaylight:params:xml:ns:yang:controller:md:sal:dom">x:dom-
↪async-data-broker</type>
20     <name>pingpong-broker</name>
21   </dom-data-provider>
22 </module>
```

@line 4: **binding-port** - The BMP server listening port.

@line 5: **binding-address** - The BMP server biding address.

Note: User may create multiple BMP monitoring station instances at runtime.

Active mode configuration

In order to enable active connection, use following request. It is required to install `odl-netconf-connector-ssh` feature first.

URL: `/restconf/config/network-topology:network-topology/topology/topology-netconf/node/controller-config/yang-ext:mount/config:modules/config:module/odl-bmp-impl-cfg:bmp-monitor-impl/example-bmp-monitor`

Method: `PUT`

Content-Type: `application/xml`

Request Body:

```
1  <module xmlns="urn:opendaylight:params:xml:ns:yang:controller:config">
2    <name>example-bmp-monitor</name>
3    <type xmlns:x="urn:opendaylight:params:xml:ns:yang:controller:bmp:impl">x:bmp-
↪monitor-impl</type>
4    <bmp-dispatcher xmlns="urn:opendaylight:params:xml:ns:yang:controller:bmp:impl">
5      <type>bmp-dispatcher</type>
6      <name>global-bmp-dispatcher</name>
7    </bmp-dispatcher>
```

(continues on next page)

(continued from previous page)

```
8     <codec-tree-factory xmlns="urn:opendaylight:params:xml:ns:yang:controller:bmp:impl
↪">
9       <type xmlns:x="urn:opendaylight:params:xml:ns:yang:controller:md:sal:binding">
↪x:binding-codec-tree-factory</type>
10      <name>runtime-mapping-singleton</name>
11    </codec-tree-factory>
12    <extensions xmlns="urn:opendaylight:params:xml:ns:yang:controller:bmp:impl">
13      <type xmlns:x="urn:opendaylight:params:xml:ns:yang:controller:bgp:rib:spi">
↪x:extensions</type>
14      <name>global-rib-extensions</name>
15    </extensions>
16    <binding-address xmlns="urn:opendaylight:params:xml:ns:yang:controller:bmp:impl">0.
↪0.0.0</binding-address>
17        <dom-data-provider xmlns=
↪"urn:opendaylight:params:xml:ns:yang:controller:bmp:impl">
18      <type xmlns:x="urn:opendaylight:params:xml:ns:yang:controller:md:sal:dom">x:dom-
↪async-data-broker</type>
19      <name>pingpong-broker</name>
20    </dom-data-provider>
21    <binding-port xmlns="urn:opendaylight:params:xml:ns:yang:controller:bmp:impl">12345
↪</binding-port>
22    <monitored-router xmlns="urn:opendaylight:params:xml:ns:yang:controller:bmp:impl">
23      <address xmlns="urn:opendaylight:params:xml:ns:yang:controller:bmp:impl">10.10.
↪10.10</address>
24      <port xmlns="urn:opendaylight:params:xml:ns:yang:controller:bmp:impl">1234</port>
25      <active xmlns="urn:opendaylight:params:xml:ns:yang:controller:bmp:impl">true</
↪active>
26    </monitored-router>
27  </module>
```

@line 23: **address** - The monitored router's IP address.

@line 24: **port** - The monitored router's port.

@line 25: **active** - Active mode set.

Note: User may configure active session establishment for multiple monitored routers.

MD5 authentication configuration

In order to enable active connection, use following request. It is required to install `odl-netconf-connector-ssh` feature first.

URL: `/restconf/config/network-topology:network-topology/topology/topology-netconf/node/controller-config/yang-ext:mount/config:modules/config:module/odl-bmp-impl-cfg:bmp-monitor-impl/example-bmp-monitor`

Method: `PUT`

Content-Type: `application/xml`

Request Body:

```
1  <module xmlns="urn:opendaylight:params:xml:ns:yang:controller:config">
2    <name>example-bmp-monitor</name>
```

(continues on next page)

(continued from previous page)

```
3     <type xmlns:x="urn:opendaylight:params:xml:ns:yang:controller:bmp:impl">x:bmp-
  ↪monitor-impl</type>
4     <bmp-dispatcher xmlns="urn:opendaylight:params:xml:ns:yang:controller:bmp:impl">
5       <type>bmp-dispatcher</type>
6       <name>global-bmp-dispatcher</name>
7     </bmp-dispatcher>
8     <codec-tree-factory xmlns="urn:opendaylight:params:xml:ns:yang:controller:bmp:impl
  ↪">
9       <type xmlns:x="urn:opendaylight:params:xml:ns:yang:controller:md:sal:binding">
  ↪x:binding-codec-tree-factory</type>
10      <name>runtime-mapping-singleton</name>
11    </codec-tree-factory>
12    <extensions xmlns="urn:opendaylight:params:xml:ns:yang:controller:bmp:impl">
13      <type xmlns:x="urn:opendaylight:params:xml:ns:yang:controller:bgp:rib:spi">
  ↪x:extensions</type>
14      <name>global-rib-extensions</name>
15    </extensions>
16    <binding-address xmlns="urn:opendaylight:params:xml:ns:yang:controller:bmp:impl">0.
  ↪0.0.0</binding-address>
17        <dom-data-provider xmlns=
  ↪"urn:opendaylight:params:xml:ns:yang:controller:bmp:impl">
18      <type xmlns:x="urn:opendaylight:params:xml:ns:yang:controller:md:sal:dom">x:dom-
  ↪async-data-broker</type>
19      <name>pingpong-broker</name>
20    </dom-data-provider>
21    <binding-port xmlns="urn:opendaylight:params:xml:ns:yang:controller:bmp:impl">12345
  ↪</binding-port>
22    <monitored-router xmlns="urn:opendaylight:params:xml:ns:yang:controller:bmp:impl">
23      <address xmlns="urn:opendaylight:params:xml:ns:yang:controller:bmp:impl">11.11.
  ↪11.11</address>
24      <password xmlns="urn:opendaylight:params:xml:ns:yang:controller:bmp:impl">
  ↪topsecret</password>
25    </monitored-router>
26  </module>
```

@line 23: **address** - The monitored router's IP address.

@line 24: **password** - The TCP MD5 signature.

Collector DB Tree

```
module: bmp-monitor
   +--rw bmp-monitor
      +--ro monitor* [monitor-id]
         +--ro monitor-id    monitor-id
         +--ro router* [router-id]
            +--ro name?          string
            +--ro description?   string
            +--ro info?          string
            +--ro router-id      router-id
            +--ro status?        status
            +--ro peer* [peer-id]
               +--ro peer-id                 rib:peer-id
               +--ro type                    peer-type
               x--ro distinguisher
```

(continues on next page)

(continued from previous page)

```
|  +--ro distinguisher-type?   distinguisher-type
|  +--ro distinguisher?        string
+--ro peer-distinguisher?     union
+--ro address                 inet:ip-address
+--ro as                      inet:as-number
+--ro bgp-id                  inet:ipv4-address
+--ro router-distinguisher?   string
+--ro peer-session
|  +--ro local-address      inet:ip-address
|  +--ro local-port         inet:port-number
|  +--ro remote-port        inet:port-number
|  +--ro sent-open
|  |  +--ro version?          protocol-version
|  |  +--ro my-as-number?     uint16
|  |  +--ro hold-timer        uint16
|  |  +--ro bgp-identifier    inet:ipv4-address
|  |  +--ro bgp-parameters*
|  |     +--ro optional-capabilities*
|  |        +--ro c-parameters
|  |           +--ro as4-bytes-capability
|  |           |  +--ro as-number?   inet:as-number
|  |           +--ro bgp-extended-message-capability!
|  |           +--ro multiprotocol-capability
|  |           |  +--ro afi?    identityref
|  |           |  +--ro safi?   identityref
|  |           +--ro graceful-restart-capability
|  |           |  +--ro restart-flags    bits
|  |           |  +--ro restart-time     uint16
|  |           |  +--ro tables* [afi safi]
|  |           |     +--ro afi          identityref
|  |           |     +--ro safi         identityref
|  |           |     +--ro afi-flags    bits
|  |           +--ro add-path-capability
|  |           |  +--ro address-families*
|  |           |     +--ro afi?            identityref
|  |           |     +--ro safi?           identityref
|  |           |     +--ro send-receive?   send-receive
|  |           +--ro route-refresh-capability!
|  +--ro received-open
|  |  +--ro version?          protocol-version
|  |  +--ro my-as-number?     uint16
|  |  +--ro hold-timer        uint16
|  |  +--ro bgp-identifier    inet:ipv4-address
|  |  +--ro bgp-parameters*
|  |     +--ro optional-capabilities*
|  |        +--ro c-parameters
|  |           +--ro as4-bytes-capability
|  |           |  +--ro as-number?   inet:as-number
|  |           +--ro bgp-extended-message-capability!
|  |           +--ro multiprotocol-capability
|  |           |  +--ro afi?    identityref
|  |           |  +--ro safi?   identityref
|  |           +--ro graceful-restart-capability
|  |           |  +--ro restart-flags    bits
|  |           |  +--ro restart-time     uint16
|  |           |  +--ro tables* [afi safi]
|  |           |     +--ro afi          identityref
```

(continues on next page)

(continued from previous page)

```
|  |           |     +--ro safi         identityref
|  |           |     +--ro afi-flags    bits
|  |           +--ro add-path-capability
|  |           |  +--ro address-families*
|  |           |     +--ro afi?            identityref
|  |           |     +--ro safi?           identityref
|  |           |     +--ro send-receive?   send-receive
|  |           +--ro route-refresh-capability!
|  +--ro information
|  |  +--ro string-information*
|  |     +--ro string-tlv
|  |        +--ro string-info?   string
|  +--ro status?            status
|  +--ro timestamp-sec?     yang:timestamp
|  +--ro timestamp-micro?   yang:timestamp
+--ro stats
|  +--ro rejected-prefixes?                 yang:counter32
|  +--ro duplicate-prefix-advertisements?   yang:counter32
|  +--ro duplicate-withdraws?               yang:counter32
|  +--ro invalidated-cluster-list-loop?     yang:counter32
|  +--ro invalidated-as-path-loop?          yang:counter32
|  +--ro invalidated-originator-id?         yang:counter32
|  +--ro invalidated-as-confed-loop?        yang:counter32
|  +--ro adj-ribs-in-routes?                yang:gauge64
|  +--ro loc-rib-routes?                    yang:gauge64
|  +--ro per-afi-safi-adj-rib-in-routes
|  |  +--ro afi-safi* [afi safi]
|  |     +--ro afi      identityref
|  |     +--ro safi     identityref
|  |     +--ro count?   yang:gauge64
|  +--ro per-afi-safi-loc-rib-routes
|  |  +--ro afi-safi* [afi safi]
|  |     +--ro afi      identityref
|  |     +--ro safi     identityref
|  |     +--ro count?   yang:gauge64
|  +--ro updates-treated-as-withdraw?       yang:counter32
|  +--ro prefixes-treated-as-withdraw?      yang:counter32
|  +--ro duplicate-updates?                 yang:counter32
|  +--ro timestamp-sec?                     yang:timestamp
|  +--ro timestamp-micro?                   yang:timestamp
+--ro pre-policy-rib
|  +--ro tables* [afi safi]
|     +--ro afi           identityref
|     +--ro safi          identityref
|     +--ro attributes
|     |  +--ro uptodate?   boolean
|     +--ro (routes)?
+--ro post-policy-rib
|  +--ro tables* [afi safi]
|     +--ro afi           identityref
|     +--ro safi          identityref
|     +--ro attributes
|     |  +--ro uptodate?   boolean
|     +--ro (routes)?
+--ro mirrors
   +--ro information?       bmp-msg:mirror-information-code
   +--ro timestamp-sec?     yang:timestamp
```

(continues on next page)

(continued from previous page)

```
                  +--ro timestamp-micro?   yang:timestamp
```

Operations

The BMP plugin offers view of collected routes and statistical information from monitored peers. To get top-level view of monitoring station:

URL: `/restconf/operational/bmp-monitor:bmp-monitor/monitor/example-bmp-monitor`

Method: `GET`

Response Body:

```
<bmp-monitor xmlns="urn:opendaylight:params:xml:ns:yang:bmp-monitor">
   <monitor>
      <monitor-id>example-bmp-monitor</monitor-id>
         <router>
            <router-id>10.10.10.10</router-id>
            <name>name</name>
            <description>monitored-router</description>
            <info>monitored router;</info>
            <status>up</status>
            <peer>
                <peer-id>20.20.20.20</peer-id>
                <address>20.20.20.20</address>
                <bgp-id>20.20.20.20</bgp-id>
                <as>65000</as>
                <type>global</type>
                <peer-session>
                  <remote-port>1790</remote-port>
                  <timestamp-sec>0</timestamp-sec>
                  <status>up</status>
                  <local-address>10.10.10.10</local-address>
                  <local-port>2200</local-port>
                  <received-open>
                     <hold-timer>180</hold-timer>
                     <my-as-number>65000</my-as-number>
                     <bgp-identifier>20.20.20.20</bgp-identifier>
                  </received-open>
                  <sent-open>
                     <hold-timer>180</hold-timer>
                     <my-as-number>65000</my-as-number>
                     <bgp-identifier>65000</bgp-identifier>
                  </sent-open>
                </peer-session>
                <pre-policy-rib>
                  <tables>
                     <afi xmlns:x="urn:opendaylight:params:xml:ns:yang:bgp-types">
↪x:ipv4-address-family</afi>
                     <safi xmlns:x="urn:opendaylight:params:xml:ns:yang:bgp-types">
↪x:unicast-subsequent-address-family</safi>
                     <ipv4-routes xmlns="urn:opendaylight:params:xml:ns:yang:bgp-inet
↪">
                        <ipv4-route>
                           <prefix>10.10.10.0/24</prefix>
                           <attributes>
```

(continues on next page)

(continued from previous page)

```
41                             ...
42                             </attributes>
43                          </ipv4-route>
44                       </ipv4-routes>
45                       <attributes>
46                          <uptodate>true</uptodate>
47                       </attributes>
48                    </tables>
49                 </pre-policy-rib>
50                 <post-policy-rib>
51                    ...
52                 </post-policy-rib>
53                 <stats>
54                    <timestamp-sec>0</timestamp-sec>
55                    <invalidated-cluster-list-loop>0</invalidated-cluster-list-loop>
56                    <duplicate-prefix-advertisements>0</duplicate-prefix-advertisements>
57                    <loc-rib-routes>100</loc-rib-routes>
58                    <duplicate-withdraws>0</duplicate-withdraws>
59                    <invalidated-as-confed-loop>0</invalidated-as-confed-loop>
60                    <adj-ribs-in-routes>10</adj-ribs-in-routes>
61                    <invalidated-as-path-loop>0</invalidated-as-path-loop>
62                    <invalidated-originator-id>0</invalidated-originator-id>
63                    <rejected-prefixes>8</rejected-prefixes>
64                 </stats>
65              </peer>
66           </router>
67        </monitor>
68     </bmp-monitor>
```

@line 3: **monitor-id** - The BMP monitoring station instance identifier.

@line 5: **router-id** - The monitored router IP address, serves as an identifier.

@line 11: **peer-id** - The monitored peer's BGP identifier, serves a an identifier.

@line 12: **address** - The IP address of the peer, associated with the TCP session.

@line 13: **bgp-id** - The BGP Identifier of the peer.

@line 14: **as** - The Autonomous System number of the peer.

@line 15: **type** - Identifies type of the peer - *Global Instance*, *RD Instance* or *Local Instance*

@line 17: **remote-port** - The peer's port number associated with TCP session.

@line 20: **local-address** - The IP address of the monitored router associated with the peering TCP session.

@line 21: **local-port** - The port number of the monitored router associated with the peering TCP session.

@line 22: **received-open** - The full OPEN message received by monitored router from the peer.

@line 27: **sent-open** - The full OPEN message send by monitored router to the peer.

@line 33: **pre-policy-rib** - The Adj-RIB-In that contains unprocessed routing information.

@line 50: **post-policy-rib** - The Post-Policy Ad-RIB-In that contains routes filtered by inbound policy.

@line 53: **stats** - Contains various statistics, periodically updated by the router.

- **To view collected information from particular monitored router: URL:** `/restconf/operational/bmp-monitor:bmp-monitor/monitor/example-bmp-monitor/router/10.10.10.10`

- **To view collected information from particular monitored peer: URL:** `/restconf/operational/bmp-monitor:bmp-monitor/monitor/example-bmp-monitor/router/10.10.10.10/peer/20.20.20.20`

Test tools

BMP test tool serves to test basic BMP functionality, scalability and performance.

BMP mock

The BMP mock is a stand-alone Java application purposed to simulate a BMP-enabled router(s) and peers. The simulator is capable to report dummy routes and statistics. This application is not part of the OpenDaylight Karaf distribution, however it can be downloaded from OpenDaylight's Nexus (use latest release version):

`https://nexus.opendaylight.org/content/repositories/opendaylight.release/org/opendaylight/bgpcep/bgp-bmp-mock`

Usage

The application can be run from command line:

```
java -jar bgp-bmp-mock-*-executable.jar
```

with optional input parameters:

```
--local_address <address> (optional, default 127.0.0.1)
   The IPv4 address where BMP mock is bind to.

--remote_address <address:port> (optional, default 127.0.0.1:12345)
   The remote IPv4 Address and port number of BMP monitoring station.

--passive (optional, not present by default)
   This flags enables passive mode for simulated routers.

--routers_count <0..N> (optional, default 1)
    An amount of BMP routers to be connected to the BMP monitoring station.

--peers_count <0..N> (optional, default 0)
   An amount of peers reported by each BMP router.

--pre_policy_routes <0..N> (optional, default 0)
   An amount of "pre-policy" simple IPv4 routes reported by each peer.

--post_policy_routes <0..N> (optional, default 0)
   An amount of "post-policy" simple IPv4 routes reported by each peer.

--log_level <FATAL|ERROR|INFO|DEBUG|TRACE> (optional, default INFO)
   Set logging level for BMP mock.
```

Troubleshooting

This section offers advices in a case OpenDaylight BMP plugin is not working as expected.

Contents

- *BMP is not working...*
- *Bug reporting*

BMP is not working...

- First of all, ensure that all required features are installed, local monitoring station and monitored router/peers configuration is correct.

  To list all installed features in OpenDaylight use the following command at the Karaf console:

```
feature:list -i
```

- Check OpenDaylight Karaf logs:

  From Karaf console:

```
log:tail
```

  or open log file: `data/log/karaf.log`

  Possibly, a reason/hint for a cause of the problem can be found there.

- Try to minimize effect of other OpenDaylight features, when searching for a reason of the problem.
- Try to set DEBUG severity level for BMP logger via Karaf console commands, in order to collect more information:

```
log:set DEBUG org.opendaylight.protocol.bmp
```

Bug reporting

Before you report a bug, check BGPCEP Bugzilla to ensure same/similar bug is not already filed there.

Write an e-mail to bgpcep-users@lists.opendaylight.org and provide following information:

1. State OpenDaylight version
2. Describe your use-case and provide as much details related to BMP as possible
3. Steps to reproduce
4. Attach Karaf log files, optionally packet captures, REST input/output

CAPWAP User Guide

This document describes how to use the Control And Provisioning of Wireless Access Points (CAPWAP) feature in OpenDaylight. This document contains configuration, administration, and management sections for the feature.

Overview

CAPWAP feature fills the gap OpenDaylight Controller has with respect to managing CAPWAP compliant wireless termination point (WTP) network devices present in enterprise networks. Intelligent applications (e.g. centralized firmware management, radio planning) can be developed by tapping into the WTP network device's operational states via REST APIs.

CAPWAP Architecture

The CAPWAP feature is implemented as an MD-SAL based provider module, which helps discover WTP devices and update their states in MD-SAL operational datastore.

Scope of CAPWAP Project

In this release, CAPWAP project aims to only detect the WTPs and store their basic attributes in the operational data store, which is accessible via REST and JAVA APIs.

Installing CAPWAP

To install CAPWAP, download OpenDaylight and use the Karaf console to install the following feature:

odl-capwap-ac-rest

Configuring CAPWAP

As of this release, there are no configuration requirements.

Administering or Managing CAPWAP

After installing the odl-capwap-ac-rest feature from the Karaf console, users can administer and manage CAPWAP from the APIDOCS explorer.

Go to http://\protect\T1\textdollar\protect\T1\textbraceleftipaddress\protect\T1\textbraceright:8181/apidoc/explorer/index.html, sign in, and expand the capwap-impl panel. From there, users can execute various API calls.

Tutorials

Viewing Discovered WTPs

Overview

This tutorial can be used as a walk through to understand the steps for starting the CAPWAP feature, detecting CAPWAP WTPs, accessing the operational states of WTPs.

Prerequisites

It is assumed that user has access to at least one hardware/software based CAPWAP compliant WTP. These devices should be configured with OpenDaylight controller IP address as a CAPWAP Access Controller (AC) address. It is also assumed that WTPs and OpenDaylight controller share the same ethernet broadcast domain.

Instructions

1. Run the OpenDaylight distribution and install odl-capwap-ac-rest from the Karaf console.
2. Go to http://\protect\T1\textdollar\protect\T1\textbraceleftipaddress\protect\T1\textbraceright:8181/apidoc/explorer/index.html
3. Expand capwap-impl
4. Click /operational/capwap-impl:capwap-ac-root/
5. Click "Try it out"
6. The above step should display list of WTPs discovered using ODL CAPWAP feature.

Cardinal: OpenDaylight Monitoring as a Service

This section describes how to use the Cardinal feature in OpenDaylight and contains configuration, administration, and management sections for the feature.

Overview

Cardinal (OpenDaylight Monitoring as a Service) enables OpenDaylight and the underlying software defined network to be remotely monitored by deployed Network Management Systems (NMS) or Analytics suite. In the Boron release, Cardinal will add:

1. OpenDaylight MIB.
2. Enable ODL diagnostics/monitoring to be exposed across SNMP (v2c, v3) and REST north-bound.
3. Extend ODL System health, Karaf parameter and feature info, ODL plugin scalability and network parameters.
4. Support autonomous notifications (SNMP Traps).

Cardinal Architecture

The Cardinal architecture can be found at the below link:

https://wiki.opendaylight.org/images/8/89/Cardinal-ODL_Monitoring_as_a_Service_V2.pdf

Configuring Cardinal feature

To start Cardinal feature, start karaf and type the following command:

```
feature:install odl-cardinal
```

After this Cardinal should be up and working with SNMP daemon running on port 161.

Tutorials

Below are tutorials for Cardinal.

Using Cardinal

These tutorials are intended for any user who wants to monitor three basic component in OpenDaylight

1. System Info in which controller is running.
2. Karaf Info
3. Project Specific Information.

Prerequisites

There is no as such specific prerequisite. Cardinal can work without installing any third party software. However If one wants to see the output of a snmpget/snmpwalk on the CLI prompt, than one can install the SNMP using the below link:

https://www.digitalocean.com/community/tutorials/how-to-install-and-configure-an-snmp-daemon-and-client-on-ubuntu-14-04

Using the above command line utility one can get the same result as the cardinal APIs will give for the snmpget/snmpwalk request.

Target Environment

This tutorial is developed considering the following environment:

controller-Linux(Ubuntu 14.02).

Instructions

Install Cardinal feature

Open karaf and install the cardinal feature using the following command:

```
feature:install odl-cardinal
```

Please verify that SNMP daemon is up on port 161 using the following command on the terminal window of Linux machine:

```
netstat -anp | grep "161"
```

If the grep on the ‘‘snmpd‘‘ port is successful than SNMP daemon is up and working.

APIs Reference

Please see Developer guide for usage of Cardinal APIs.

CLI commands to do snmpget/walk

One can do snmpget/walk on the ODL-CARDINAL-MIB. Open the linux terminal and type the below command:

```
snmpget -v2c -c public localhost Oid_Of_the_mib_variable
```

Or

```
snmpget -v2c -c public localhost ODL-CARDINAL-MIB::mib_variable_name
```

For snmpwalk use the below command:

```
snmpwalk -v2c -c public localhost SNMPv2-SMI::experimental
```

Centinel User Guide

The Centinel project aims at providing a distributed, reliable framework for efficiently collecting, aggregating and sinking streaming data across Persistence DB and stream analyzers (example: Graylog, Elastic search, Spark, Hive etc.). This document contains configuration, administration, management, using sections for the feature.

Overview

In this release of Centinel, this framework enables SDN applications/services to receive events from multiple streaming sources (e.g., Syslog, Thrift, Avro, AMQP, Log4j, HTTP/REST) and execute actions like network configuration/batch processing/real-time analytics. It also provides a Log Service to assist operators running SDN ecosystem by installing the feature odl-centinel-all.

With the configurations development of "Log Service" and plug-in for log analyzer (e.g., Graylog) will take place. Log service will do processing of real time events coming from log analyzer. Additionally, stream collector (Flume and Sqoop based) that will collect logs from OpenDaylight and sink it to persistence service (integrated with TSDR). Also includes RESTCONF interface to inject events to north bound applications for real-time analytic/network configuration. Centinel User Interface (web interface) will be available to operators to enable rules/alerts/dashboard.

Centinel core features

The core features of the Centinel framework are:

Stream collector Collecting, aggregating and sinking streaming data

Log Service Listen log stream events coming from log analyzer

Log Service Enables user to configure rules (e.g., alerts, diagnostic, health, dashboard)

Log Service Performs event processing/analytics

User Interface Enable set-rule, search, visualize, alert, diagnostic, dashboard etc.

Adaptor Log analyzer plug-in to Graylog and a generic data-model to extend to other stream analyzers (e.g., Logstash)

REST Service Northbound APIs for Log Service and Steam collector framework

Leverages TSDR persistence service, data query, purging and elastic search

Centinel Architecture

The following wiki pages capture the Centinel Model/Architecture

1. https://wiki.opendaylight.org/view/Centinel:Main
2. https://wiki.opendaylight.org/view/Project_Proposals:Centinel
3. https://wiki.opendaylight.org/images/0/09/Centinel-08132015.pdf

Administering or Managing Centinel with default configuration

Prerequisites

1. Check whether Graylog is up and running and plugins deployed as mentioned in installation guide.
2. Check whether HBase is up and respective tables and column families as mentioned in installation guide are created.
3. Check if apache flume is up and running.
4. Check if apache drill is up and running.

Running Centinel

The following steps should be followed to bring up the controller:

1. Download the Centinel OpenDaylight distribution release from below link: http://www.opendaylight.org/software/downloads
2. Run Karaf of the distribution from bin folder

```
./karaf
```

3. Install the centinel features using below command:

```
feature:install odl-centinel-all
```

4. Give some time for the centinel to come up.

User Actions

1. **Log In:** User logs into the Centinel with required credentials using following URL: http://localhost:8181/index.html
2. **Create Rule:**
   (a) Select Centinel sub-tree present in left side and go to Rule tab.
   (b) Create Rule with title and description.
   (c) Configure flow rule on the stream to filter the logs accordingly for, e.g., `bundle_name=org.opendaylight.openflow-plugin`
3. **Set Alarm Condition:** Configure alarm condition, e.g., message-count-rule such that if 10 messages comes on a stream (e.g., The OpenFlow Plugin) in last 1 minute with an alert is generated.

4. **Subscription:** User can subscribe to the rule and alarm condition by entering the http details or email-id in subscription textfield by clicking on the subscribe button.
5. **Create Dashboard:** Configure dashboard for stream and alert widgets. Alarm and Stream count will be updated in corresponding widget in Dashboard.
6. **Event Tab:** Intercepted Logs, Alarms and Raw Logs in Event Tab will be displayed by selecting the appropriate radio button. User can also filter the searched data using SQL query in the search box.

DIDM User Guide

Overview

The Device Identification and Driver Management (DIDM) project addresses the need to provide device-specific functionality. Device-specific functionality is code that performs a feature, and the code is knowledgeable of the capability and limitations of the device. For example, configuring VLANs and adjusting FlowMods are features, and there may be different implementations for different device types. Device-specific functionality is implemented as Device Drivers. Device Drivers need to be associated with the devices they can be used with. To determine this association requires the ability to identify the device type.

DIDM Architecture

The DIDM project creates the infrastructure to support the following functions:

- **Discovery** - Determination that a device exists in the controller management domain and connectivity to the device can be established. For devices that support the OpenFlow protocol, the existing discovery mechanism in OpenDaylight suffices. Devices that do not support OpenFlow will be discovered through manual means such as the operator entering device information via GUI or REST API.
- **Identification** – Determination of the device type.
- **Driver Registration** – Registration of Device Drivers as routed RPCs.
- **Synchronization** – Collection of device information, device configuration, and link (connection) information.
- **Data Models for Common Features** – Data models will be defined to perform common features such as VLAN configuration. For example, applications can configure a VLAN by writing the VLAN data to the data store as specified by the common data model.
- **RPCs for Common Features** – Configuring VLANs and adjusting FlowMods are example of features. RPCs will be defined that specify the APIs for these features. Drivers implement features for specific devices and support the APIs defined by the RPCs. There may be different Driver implementations for different device types.

Atrium Support

The Atrium implements an open source router that speaks BGP to other routers, and forwards packets received on one port/vlan to another, based on the next-hop learnt via BGP peering. A BGP peering application for the Open Daylight Controller and a new model for flow objective drivers for switches integrated with the Open Daylight Atrium distribution was developed for this project. The implementation has the same level of feature partly that was introduced by the Atrium 2015/A distribution on the ONOS controller. An overview of the architecture is available at here (https://github.com/onfsdn/atrium-docs/wiki/ODL-Based-Atrium-Router-16A).

Atrium stack is implemented in OpenDaylight using Atrium and DIDM project. Atrium project provides the application implementation for BGP peering and the DIDM project provides implementation for FlowObjectives. FlowObjective provides an abstraction layer and present the pipeline agnostic api to application to consume.

FlowObjective

Flow Objectives describe an SDN application's objective (or intention) behind a flow it is sending to a device.

Application communicates the flow installation requirement using Flow Objectives. DIDM drivers translates the Flow Objectives to device specific flows as per the device pipeline.

There are three FlowObjectives (already implemented in ONOS controller) :

- Filtering Objective
- Next Objective
- Forwarding Objective

Installing DIDM

To install DIDM, download OpenDaylight and use the Karaf console to install the following features:

- odl-openflowplugin-all
- odl-didm-all

odl-didm-all installs the following required features:

- odl-didm-ovs-all
- odl-didm-ovs-impl
- odl-didm-util
- odl-didm-identification
- odl-didm-drivers
- odl-didm-hp-all

Configuring DIDM

This section shows an example configuration steps for installing a driver (HP 3800 OpenFlow switch driver).

Install DIDM features:

```
feature:install odl-didm-identification-api
feature:install odl-didm-drivers
```

In order to identify the device, device driver needs to be installed first. Identification Manager will be notified when a new device connects to the controller.

Install HP driver

feature:install odl-didm-hp-all installs the following features

- odl-didm-util
- odl-didm-identification
- odl-didm-drivers

- odl-didm-hp-all
- odl-didm-hp-impl

Now at this point, the driver has written all of the identification information in to the MD-SAL datastore. The identification manager should have that information so that it can try to identify the HP 3800 device when it connects to the controller.

Configure the switch and connect it to the controller from the switch CLI.

Run REST GET command to verify the device details:

http://<CONTROLLER-IP:8181>/restconf/operational/opendaylight-inventory:nodes

Run REST adjust-flow command to adjust flows and push to the device

Flow mod driver for HP 3800 device

This driver adjusts the flows and push the same to the device. This API takes the flow to be adjusted as input and displays the adjusted flow as output in the REST output container. Here is the REST API to adjust and push flows to HP 3800 device:

http://<CONTROLLER-IP:8181>/restconf/operations/openflow-feature:adjust-flow

FlowObjectives API

FlowObjective presents the OpenFlow pipeline agnostic API to Application to consume. Application communicate their intent behind installation of flow to Drivers using the FlowObjective. Driver translates the FlowObjective in device specific flows and uses the OpenFlowPlugin to install the flows to the device.

Filter Objective

http://<CONTROLLER-IP>:8181/restconf/operations/atrium-flow-objective:filter

Next Objective

http://<CONTROLLER-IP>:8181/restconf/operations/atrium-flow-objective:next

Forward Objective

http://<CONTROLLER-IP>:8181/restconf/operations/atrium-flow-objective:forward

Fabric As A Service

This document describes, from a user's or application's perspective, how to use the Fabric As A Service (FaaS) feature in OpenDaylight. This document contains configuration, administration, and management sections for the FaaS feature.

Overview

Currently network applications and network administrators mostly rely on lower level interfaces such as CLI, SNMP, OVSDB, NETCONF or OpenFlow to directly configure individual device for network service provisioning. In general, those interfaces are:

- Technology oriented, not application oriented.
- Vendor specific
- Individual device oriented, not network oriented.
- Not declarative, complicated and procedure oriented.

To address the gap between application needs and network interface, there are a few application centric language proposed in OpenDaylight including Group Based Policy (GBP), Network Intent Composition (NIC), and NEtwork MOdeling (NEMO) trying to replace traditional southbound interface to application. Those languages are top-down abstractions and model application requirements in a more application-oriented way.

After being involved with GBP development for a while, we feel the top down model still has a quite gap between the model and the underneath network since the existing interfaces to network devices lack of abstraction makes it very hard to map high level abstractions to the physical network. Often the applications built with these low level interfaces are coupled tightly with underneath technology and make the application's architecture monolithic, error prone and hard to maintain.

We think a bottom-up abstraction of network can simplify, reduce the gap, and make it easy to implement the application centric model. Moreover in some uses cases, an interface with network service oriented are still desired for example from network monitoring/troubleshooting perspective. That's where the Fabric as a Service comes in.

FaaS Architecture

Fabric and its controller (Fabric Controller) The Fabric object provides an abstraction of a homogeneous network or portion of the network and also has a built in Fabric controller which provides management plane and control plane for the fabric. The fabric controller implements the services required in Fabric Service and monitor and control the fabric operation.

Fabric Manager Fabric Manager manages all the fabric objects. also Fabric manager acts as a Unified Fabric Controller which provides inter-connect fabric control and configuration Also Fabric Manager is FaaS API service via Which FaaS user level logical network API (the top level API as mentioned previously) exposed and implemented.

FaaS render for GBP (Group Based Policy) FaaS render for GBP is an application of FaaS and provides the rendering service between GBP model and logical network model provided by Fabric Manager.

FaaS RESTCONF API

FaaS Provides two layers API:

- The top layer is a **user level logical network** API which defines CRUD services operating on the following abstracted constructs:
  - vcontainer - virtual container
  - logical Port - a input/out/access point of a logical device
  - logical link - connects ports
  - logical switch - a layer 2 forwarding device

  - logical router - a layer 3 forwarding device
  - Subnet
  - Rule/ACL
  - End Points Registration
  - End Points Attachment

Through these constructs, a logical network can be described without users knowing too much details about the physical network device and technology, i.e. users' network services is decoupled from underneath physical infrastructure. This decoupling brought the benefit that the users defined service is not locked in with any specific technology or physical devices. FaaS maps the logical network to the physical network configuration automatically which in maximum eliminates the manual provisioning work. As a result, human error is avoided and OPEX for network operators is massively reduced. Moreover migration from one technology to another is much easier to do and transparent to users' services.

- The second layer defines an abstraction layer called **Fabric** API. The idea is to abstract network into a topology formed by a collections of fabric objects other than varies of physical devices.Each Fabric object provides a collection of unified services. The top level API enables application developers or users to write applications to map high level model such as GBP, Intent etc... into a logical network model, while the lower level gives the application more control to individual fabric object level. More importantly the Fabric API is more like SPI (Service Provider API) a fabric provider or vendor can implement the SPI based on its own Fabric technique such as TRILL, SPB etc ...

This document is focused on the top layer API. For how to use second level API operation, please refer to FaaS developer guide for more details.

Note: that for any JSON data or link not described here, please go to http://\protect\T1\textdollar\protect\T1\textbraceleftipaddress\protect\T1\textbraceright:8181/apidoc/explorer/index.htm for details or clarification.

Resource Management API

The FaaS Resource Management API provides services to allocate and reclaim the network resources provided by Fabric object. Those APIs can be accessed via RESTCONF rendered from YANG in MD-SAL.

- Name: Create virtual container
  - virtual container is an collections of resources allocated to a tenant, for example, a list of physical ports, bandwidth or L2 network or logical routers quantity. etc...
  - http://\protect\T1\textdollar\protect\T1\textbraceleftipaddress\protect\T1\textbraceright:8181/restconf/operations/vcontainer-topology:create-vcontainer
  - Description: create a given virtual container.
- Name: assign or remove fabric resource to a virtual container
  - http://\protect\T1\textdollar\protect\T1\textbraceleftipaddress\protect\T1\textbraceright:8181/restconf/operations/vc-ld-node:add-vfabric-to-ld-node
  - http://\protect\T1\textdollar\protect\T1\textbraceleftipaddress\protect\T1\textbraceright:8181/restconf/operations/vc-ld-node:rm-vfabric-to-ld-node
- Name: assign or remove appliance to a virtual container
  - http://\protect\T1\textdollar\protect\T1\textbraceleftipaddress\protect\T1\textbraceright:8181/restconf/operations/vc-ld-node:add-appliance-to-ld-node

  - http://\protect\T1\textdollar\protect\T1\textbraceleftipaddress\protect\T1\textbraceright:8181/restconf/operations/vc-ld-node:rm-appliance-to-ld-node

- Name: create or remove a child container
  - http://\protect\T1\textdollar\protect\T1\textbraceleftipaddress\protect\T1\textbraceright:8181/restconf/operations/vc-ld-node:create-child-ld-node
  - http://\protect\T1\textdollar\protect\T1\textbraceleftipaddress\protect\T1\textbraceright:8181/restconf/operations/vc-ld-node:rm-child-ld-node
- RESTCONF path for Virtual Container Datastore query (note: vcontainer-id equals tenantID for now):
  - http://\protect\T1\textdollar\protect\T1\textbraceleftipaddress\protect\T1\textbraceright:8181/restconf/config/network-topology/topology/\protect\T1\textbraceleftvcontainer-id\protect\T1\textbraceright/vc-topology-attributes/
  - http://\protect\T1\textdollar\protect\T1\textbraceleftipaddress\protect\T1\textbraceright:8181/restconf/config/network-topology/topology/\protect\T1\textbraceleftvcontainer-id\protect\T1\textbraceright/node/\protect\T1\textbraceleftnet-node-id\protect\T1\textbraceright/vc-net-node-attributes
  - http://\protect\T1\textdollar\protect\T1\textbraceleftipaddress\protect\T1\textbraceright:8181/restconf/config/network-topology/topology/\protect\T1\textbraceleftvcontainer-id\protect\T1\textbraceright/node/\protect\T1\textbraceleftld-node-id\protect\T1\textbraceright/vc-ld-node-attributes

Installing Fabric As A Service

To install FaaS, download OpenDaylight and use the Karaf console to install the following feature: **odl-restconf odl-faas-all odl-groupbasedpolicy-faas** (if needs to use FaaS to render GBP)

Configuring FaaS

This section gives details about the configuration settings for various components in FaaS.

The FaaS configuration files for the Karaf distribution are located in distribution/karaf/target/assembly/etc/faas

- akka.conf
  - This file contains configuration related to clustering. Potential configuration properties can be found on the akka website at http://doc.akka.io
- fabric-factory.xml
- vxlan-fabric.xml
- vxlan-fabric-ovs-adapter.xml
  - Those 3 files are used to initialize fabric module and located under distribution/karaf/target/assembly/etc/opendaylight/karaf

Managing FaaS

Start opendaylight karaf distribution

- *>bin/karaf* Then From karaf console,Install features in the following order:
- *>feature:Install odl-restconf*
- *>feature:install odl-faas-all*

- *>feature install odl-groupbasedpolicy-faas*

After installing features above, users can manage Fabric resource and FaaS logical network channels from the API-DOCS explorer via RESTCONF

Go to http://\protect\T1\textdollar\protect\T1\textbraceleftipaddress\protect\T1\textbraceright:8181/apidoc/explorer/index.html, sign in, and expand the FaaS panel. From there, users can execute various API calls to test their FaaS deployment such as create virtual container, create fabric, delete fabric, create/edit logical network elements.

Tutorials

Below are tutorials for 4 major use cases.

1. to create and provision a fabric
2. to allocate resource from the fabric to a tenant
3. to define a logical network for a tenant. Currently there are two ways to create a logical network
   (a) Create a GBP (Group Based Policy) profile for a tenant and then convert it to a logical network via GBP FaaS render Or
   (b) Manually create a logical network via RESTCONF APIs.
4. to attach or detach an Endpoint to a logical switch or logical router

Create a fabric

Overview

This tutorial walks users through the process of create a Fabric object

Prerequisites

A set of virtual switches (OVS) have to be registered or discovered by ODL. Mininet is recommended to create a OVS network. After an OVS network is created, set up the controller IP pointing to ODL IP address in each of the OVS. From ODL, a physical topology can be viewed via ODL DLUX UI or retrieved via RESTCONF API.

Instructions

- Run the OpenDaylight distribution and install odl-faas-all from the Karaf console.
- Go to http://\protect\T1\textdollar\protect\T1\textbraceleftipaddress\protect\T1\textbraceright:8181/apidoc/explorer/index.html
- Get the network topology after OVS switches are registered in the controller
- Determine the nodes and links to be included in the to-be-defined Fabric object.
- Execute create-fabric RESTCONF API with the corresponding JSON data as required.

Create virtual container for a tenant

The purpose of this tutorial is to allocate network resources to a tenant

Overview

This tutorial walks users through the process of create a Fabric

Prerequisites

1 or more fabric objects have been created.

Instructions

- Run the OpenDaylight karaf distribution and install odl-faas-all feature from the Karaf console. >feature:install odl-rest-conf odl-faas-all odl-mdsal-apidoc
- Go to http://\protect\T1\textdollar\protect\T1\textbraceleftipaddress\protect\T1\textbraceright:8181/apidoc/explorer/index.html
- Execute create-vcontainer with the following restconf API with corresponding JSON data > http://\protect\T1\textdollar\protect\T1\textbraceleftipaddress\protect\T1\textbraceright:8181/restconf/operations/vcontainer-topology:create-vcontainer

After a virtual container is created, fabric resource and appliance resource can be assigned to the container object via the following RESTConf API.

- http://\protect\T1\textdollar\protect\T1\textbraceleftipaddress\protect\T1\textbraceright:8181/restconf/operations/vc-ld-node:add-vfabric-to-ld-node
- http://\protect\T1\textdollar\protect\T1\textbraceleftipaddress\protect\T1\textbraceright:8181/restconf/operations/vc-ld-node:add-appliance-to-ld-node

Create a logical network

Overview

This tutorial walks users through the process of create a logical network for a tenant

Prerequisites

a virtual container has been created and assigned to the tenant

Instructions

Currently there are two ways to create a logical network.

- Option 1 is to use logical network RESTConf REST API and directly create individual network elements and connect them into a network
- Option 2 is to define a GBP model and FaaS can map GBP model automatically into a logical network. Notes that for option 2, if the generated network requires some modification, we recommend modify the GBP model rather than change the network directly due to there is no synchronization from network back to GBP model in current release.

Manual Provisioning

To create a logical switch

- http://\protect\T1\textdollar\protect\T1\textbraceleftipaddress\protect\T1\textbraceright:8181/restconf/configuration/faas-logical-networks:tenant-logical-networks:logical-switches:logical-switches To create a logical router
- http://\protect\T1\textdollar\protect\T1\textbraceleftipaddress\protect\T1\textbraceright:8181/restconf/configuration/faas-logical-networks:tenant-logical-networks:logical-routers:logical-routers To attach a logical switch to a router
  - Step 1: updating/adding a port A on the logical switch http://\protect\T1\textdollar\protect\T1\textbraceleftipaddress\protect\T1\textbraceright:8181/restconf/configuration/faas-logical-networks:tenant-logical-networks:logical-switches:logical-switches
  - Step 2: updating/adding a port B on the logical router http://\protect\T1\textdollar\protect\T1\textbraceleftipaddress\protect\T1\textbraceright:8181/restconf/configuration/faas-logical-networks:tenant-logical-networks:logical-routers:logical-routers
  - Step 3; create a link between the port A and B http://\protect\T1\textdollar\protect\T1\textbraceleftipaddress\protect\T1\textbraceright:8181/restconf/configuration/faas-logical-networks:tenant-logical-networks:logical-edges:logical-edges
- To add security policies (ACL or SFC) on a port http://\protect\T1\textdollar\protect\T1\textbraceleftipaddress\protect\T1\textbraceright:8181/restconf/configuration/faas-logical-networks:tenant-logical-networks:faas-security-rules
- To query the logical network just created http://\protect\T1\textdollar\protect\T1\textbraceleftipaddress\protect\T1\textbraceright:8181/restconf/configuration/faas-logical-networks:tenant-logical-networks

Provision via GBP FaaS Render

- Run the OpenDaylight distribution and install odl-faas-all and GBP faas render feature from the Karaf console.
- Go to http://\protect\T1\textdollar\protect\T1\textbraceleftipaddress\protect\T1\textbraceright:8181/apidoc/explorer/index.html
- Execute "create GBP model" via GBP REST API. The GBP model then can be automatically mapped into a logical network.

Attach/detach an end point to a logical device

Overview

This tutorial walks users through the process of registering an End Point to a logical device either logical switch or router. The purpose of this API is to inform the FaaS where an endpoint physically attach. The location information consists of the binding information between physical port identifier and logical port information. The logical port is indicated by the endpoint either Layer 2 attribute(MAC address) or Layer 3 attribute (IP address) and logical network ID (VLAN ID). The logical network ID is indirectly indicated the tenant ID since it is mutual exclusive resource allocated to a tenant.

Prerequisites

The logical switch to which those end points are attached has to be created beforehand. and the identifier of the logical switch is required for the following RESTCONF calls.

Instructions

- Run the OpenDaylight distribution and install odl-faas-all from the Karaf console.
- Go to http://\protect\T1\textdollar\protect\T1\textbraceleftipaddress\protect\T1\textbraceright:8181/apidoc/explorer/index.html
- Execute "attach end point " with the following RESTCONF API and corresponding JSON data: http://\protect\T1\textdollar\protect\T1\textbraceleftipaddress\protect\T1\textbraceright:8181/restconf/configuration/faas-logical-networks:tenant-logical-networks:faas-endpoints-locations

Genius User Guide

Overview

The Genius project provides generic network interfaces, utilities and services. Any OpenDaylight application can use these to achieve interference-free co-existence with other applications using Genius.

Modules and Interfaces

In the first phase delivered in OpenDaylight Boron release, Genius provides following modules—

- Modules providing a common view of network interfaces for different services
  - **Interface (logical port) Manager**
    - *Allows bindings/registration of multiple services to logical ports/interfaces*
    - *Ability to plug in different types of southbound protocol renderers*
  - **Overlay Tunnel Manager**
    - *Creates and maintains overlay tunnels between configured Tunnel Endpoints (TEPs)*
- Modules providing commonly used functions as shared services to avoid duplication of code and waste of resources
  - **Liveness Monitor**
    - *Provides tunnel/nexthop liveness monitoring services*
  - **ID Manager**
    - *Generates persistent unique integer IDs*
  - **MD-SAL Utils**
    - *Provides common generic APIs for interaction with MD-SAL*

Interface Manager Operations

Creating interfaces

The YANG file Data Model odl-interface.yang contains the interface configuration data-model.

You can create interfaces at the MD-SAL Data Node Path **/config/if:interfaces/interface**, with the following attributes—

Common attributes

- **name**—unique interface name, can be any unique string (e.g., UUID string)
- **type**—interface type, currently supported *iana-if-type:l2vlan and iana-if-type:tunnel*
- **enabled**—admin status, possible values *true* or *false*
- **parent-refs** : used to specify references to parent interface/port feeding to this interface
- datapath-node-identifier—identifier for a fixed/physical dataplane node, can be physical switch identifier
- parent-interface—can be a physical switch port (in conjunction of above), virtual switch port (e.g., neutron port) or another interface
- list node-identifier—identifier of the dependant underlying configuration protocol
  - *topology-id*—can be ovsdb configuration protocol
  - *node-id*—can be hwvtep node-id

Type specific attributes

- when type = l2vlan
  - **vlan-id**—VLAN id for trunk-member l2vlan interfaces
  - **l2vlan-mode**—currently supported ones are *transparent*, *trunk* or *trunk-member*
- when type = stacked_vlan (Not supported yet)
  - **stacked-vlan-id**—VLAN-Id for additional/second VLAN tag
- when type = tunnel
  - **tunnel-interface-type**—tunnel type, currently supported ones are:
    - tunnel-type-vxlan
    - tunnel-type-gre
    - tunnel-type-mpls-over-gre
  - **tunnel-source**—tunnel source IP address
  - **tunnel-destination**—tunnel destination IP address
  - **tunnel-gateway**—gateway IP address
  - **monitor-enabled**—tunnel monitoring enable control
  - **monitor-interval**—tunnel monitoring interval in millisiconds
- when type = mpls (Not supported yet)
  - **list labelStack**—list of lables
  - **num-labels**—number of lables configured

Supported REST calls are **GET, PUT, DELETE, POST**

Creating L2 port interfaces

Interfaces on normal L2 ports (e.g. Neutron tap ports) are created with type *l2vlan* and *l2vlan-mode* as *transparent*. This type of interface classifies packets passing through a particular L2 (OpenFlow) port. In dataplane, packets belonging to this interface are classified by matching in-port against the of-port-id assigned to the base port as specified in parent-interface.

URL: /restconf/config/ietf-interfaces:interfaces

Sample JSON data

```
"interfaces": {
    "interface": [
        {
            "name": "4158408c-942b-487c-9a03-0b603c39d3dd",
            "type": "iana-if-type:l2vlan",                        <--- interface type
↪'l2vlan' for normal L2 port
            "odl-interface:l2vlan-mode": "transparent",           <--- 'transparent'␣
↪VLAN port mode allows any (tagged, untagged) ethernet packet
            "odl-interface:parent-interface": "tap4158408c-94",   <--- port-name as it␣
↪appears on southbound interface
            "enabled": true
        }
    ]
}
```

Creating VLAN interfaces

A VLAN interface is created as a *l2vlan* interface in *trunk-member* mode, by configuring a VLAN-Id and a particular L2 (vlan trunk) interface. Parent VLAN trunk interface is created in the same way as the *transparent* interface as specified above. A *trunk-member* interface defines a flow on a particular L2 port and having a particular VLAN tag. On ingress, after classification the VLAN tag is popped out and corresponding unique dataplane-id is associated with the packet, before delivering the packet to service processing. When a service module delivers the packet to this interface for egress, it pushes corresponding VLAN tag and sends the packet out of the parent L2 port.

URL: /restconf/config/ietf-interfaces:interfaces

Sample JSON data

```
"interfaces": {
    "interface": [
        {
            "name": "4158408c-942b-487c-9a03-0b603c39d3dd:100",
            "type": "iana-if-type:l2vlan",
            "odl-interface:l2vlan-mode": "trunk-member",        <--- for 'trunk-member
↪', flow is classified with particular vlan-id on an l2 port
            "odl-interface:parent-interface": "4158408c-942b-487c-9a03-0b603c39d3dd",␣
↪ <--- Parent 'trunk' iterface name
            "odl-interface:vlan-id": "100",
            "enabled": true
        }
    ]
}
```

Creating Overlay Tunnel Interfaces

An overlay tunnel interface is created with type *tunnel* and particular *tunnel-interface-type*. Tunnel interfaces are created on a particular data plane node (virtual switches) with a pair of (local, remote) IP addresses. Currently supported tunnel interface types are VxLAN, GRE and MPLSoverGRE.

URL: /restconf/config/ietf-interfaces:interfaces

Sample JSON data

```
"interfaces": {
    "interface": [
        {
            "name": "MGRE_TUNNEL:1",
            "type": "iana-if-type:tunnel",
            "odl-interface:tunnel-interface-type": "odl-interface:tunnel-type-mpls-
↪over-gre",
            "odl-interface:datapath-node-identifier": 156613701272907,
            "odl-interface:tunnel-source": "11.0.0.43",
            "odl-interface:tunnel-destination": "11.0.0.66",
            "odl-interface:monitor-enabled": false,
            "odl-interface:monitor-interval": 10000,
            "enabled": true
        }
    ]
}
```

Binding services on interface

The YANG file odl-interface-service-bindings.yang contains the service binding configuration data model.

An application can bind services to a particular interface by configuring MD-SAL data node at path /config/interface-service-binding. Binding services on interface allows particular service to pull traffic arriving on that interface depending upon the service priority. Service modules can specify openflow-rules to be applied on the packet belonging to the interface. Usually these rules include sending the packet to specific service table/pipeline. Service modules are responsible for sending the packet back (if not consumed) to service dispatcher table, for next service to process the packet.

URL:/restconf/config/interface-service-bindings:service-bindings/

Sample JSON data

```
"service-bindings": {
  "services-info": [
    {
      "interface-name": "4152de47-29eb-4e95-8727-2939ac03ef84",
      "bound-services": [
        {
          "service-name": "ELAN",
          "service-type": "interface-service-bindings:service-type-flow-based"
          "service-priority": 3,
          "flow-priority": 5,
          "flow-cookie": 134479872,
          "instruction": [
            {
              "order": 2,
              "go-to-table": {
```

(continues on next page)

(continued from previous page)

```
                "table_id": 50
              }
            },
            {
              "order": 1,
              "write-metadata": {
                "metadata": 83953188864,
                "metadata-mask": 1099494850560
              }
            }
          ],
        },
        {
         "service-name": "L3VPN",
         "service-type": "interface-service-bindings:service-type-flow-based"
         "service-priority": 2,
         "flow-priority": 10,
         "flow-cookie": 134217729,
         "instruction": [
            {
              "order": 2,
              "go-to-table": {
                "table_id": 21
              }
            },
            {
              "order": 1,
              "write-metadata": {
                "metadata": 100,
                "metadata-mask": 4294967295
              }
            }
          ],
        }
      ]
    }
  ]
}
```

Interface Manager RPCs

In addition to the above defined configuration interfaces, Interface Manager also provides several RPCs to access interface operational data and other helpful information. Interface Manger RPCs are defined in odl-interface-rpc.yang

The following RPCs are available—

get-dpid-from-interface

This RPC is used to retrieve dpid/switch hosting the root port from given interface name.

```
rpc get-dpid-from-interface {
    description "used to retrieve dpid from interface name";
    input {
        leaf intf-name {
```

(continues on next page)

(continued from previous page)

```
            type string;
        }
    }
    output {
        leaf dpid {
            type uint64;
        }
    }
}
```

get-port-from-interface

This RPC is used to retrieve south bound port attributes from the interface name.

```
rpc get-port-from-interface {
    description "used to retrieve south bound port attributes from the interface name
↪";
    input {
        leaf intf-name {
            type string;
        }
    }
    output {
        leaf dpid {
            type uint64;
        }
        leaf portno {
            type uint32;
        }
        leaf portname {
            type string;
        }
    }
}
```

get-egress-actions-for-interface

This RPC is used to retrieve group actions to use from interface name.

```
rpc get-egress-actions-for-interface {
    description "used to retrieve group actions to use from interface name";
    input {
        leaf intf-name {
            type string;
            mandatory true;
        }
        leaf tunnel-key {
            description "It can be VNI for VxLAN tunnel ifaces, Gre Key for GRE␣
↪tunnels, etc.";
            type uint32;
            mandatory false;
        }
    }
```

(continues on next page)

(continued from previous page)

```
    output {
        uses action:action-list;
    }
}
```

get-egress-instructions-for-interface

This RPC is used to retrieve flow instructions to use from interface name.

```
rpc get-egress-instructions-for-interface {
    description "used to retrieve flow instructions to use from interface name";
    input {
        leaf intf-name {
            type string;
            mandatory true;
        }
        leaf tunnel-key {
            description "It can be VNI for VxLAN tunnel ifaces, Gre Key for GRE␣
↪tunnels, etc.";
            type uint32;
            mandatory false;
        }
    }
    output {
        uses offlow:instruction-list;
    }
}
```

get-endpoint-ip-for-dpn

This RPC is used to get the local ip of the tunnel/trunk interface on a particular DPN (Data Plane Node).

```
rpc get-endpoint-ip-for-dpn {
    description "to get the local ip of the tunnel/trunk interface";
    input {
        leaf dpid {
            type uint64;
        }
    }
    output {
        leaf-list local-ips {
            type inet:ip-address;
        }
    }
}
```

get-interface-type

This RPC is used to get the type of the interface (vlan/vxlan or gre).

```
rpc get-interface-type {
description "to get the type of the interface (vlan/vxlan or gre)";
    input {
        leaf intf-name {
            type string;
        }
    }
    output {
        leaf interface-type {
            type identityref {
                base if:interface-type;
            }
        }
    }
}
```

get-tunnel-type

This RPC is used to get the type of the tunnel interface(vxlan or gre).

```
rpc get-tunnel-type {
description "to get the type of the tunnel interface (vxlan or gre)";
    input {
        leaf intf-name {
            type string;
        }
    }
    output {
        leaf tunnel-type {
            type identityref {
                base odlif:tunnel-type-base;
            }
        }
    }
}
```

get-nodeconnector-id-from-interface

This RPC is used to get node-connector-id associated with an interface.

```
rpc get-nodeconnector-id-from-interface {
description "to get nodeconnector id associated with an interface";
    input {
        leaf intf-name {
            type string;
        }
    }
    output {
        leaf nodeconnector-id {
            type inv:node-connector-id;
        }
    }
}
```

get-interface-from-if-index

This RPC is used to get interface associated with an if-index (dataplane interface id).

```
rpc get-interface-from-if-index {
    description "to get interface associated with an if-index";
        input {
            leaf if-index {
                type int32;
            }
        }
        output {
            leaf interface-name {
                type string;
            }
        }
    }
```

create-terminating-service-actions

This RPC is used to create the tunnel termination service table entries.

```
rpc create-terminating-service-actions {
description "create the ingress terminating service table entries";
    input {
         leaf dpid {
             type uint64;
         }
         leaf tunnel-key {
             type uint64;
         }
         leaf interface-name {
             type string;
         }
         uses offlow:instruction-list;
    }
}
```

remove-terminating-service-actions

This RPC is used to remove the tunnel termination service table entries.

```
rpc remove-terminating-service-actions {
description "remove the ingress terminating service table entries";
    input {
         leaf dpid {
             type uint64;
         }
         leaf interface-name {
             type string;
         }
         leaf tunnel-key {
             type uint64;
         }
```

(continues on next page)

(continued from previous page)

```
        }
}
```

ID Manager

TBD.

Group Based Policy User Guide

Overview

OpenDaylight Group Based Policy allows users to express network configuration in a declarative versus imperative way.

This is often described as asking for **"what you want"**, rather than **"how to do it"**.

In order to achieve this Group Based Policy (herein referred to as **GBP**) is an implementation of an **Intent System**.

An **Intent System**:

- is a process around an intent driven data model
- contains no domain specifics
- is capable of addressing multiple semantic definitions of intent

To this end, **GBP** Policy views an **Intent System** visually as:

- **expressed intent** is the entry point into the system.
- **operational constraints** provide policy for the usage of the system which modulates how the system is consumed. For instance *"All Financial applications must use a specific encryption standard"*.
- **capabilities and state** are provided by *renderers*. *Renderers* dynamically provide their capabilities to the core model, allowing the core model to remain non-domain specific.
- **governance** provides feedback on the delivery of the *expressed intent*. i.e. *"Did we do what you asked us?"*

In summary **GBP is about the Automation of Intent**.

By thinking of **Intent Systems** in this way, it enables:

- **automation of intent**

  By focusing on **Model. Process. Automation**, a consistent policy resolution process enables for mapping between the **expressed intent** and renderers responsible for providing the capabilities of implementing that intent.

- recursive/intent level-independent behaviour.

  Where *one person's concrete is another's abstract*, intent can be fulfilled through a hierarchical implementation of non-domain specific policy resolution. Domain specifics are provided by the *renderers*, and exposed via the API, at each policy resolution instance. For example:

  - To DNS: The name "www.foo.com" is *abstract*, and it's IPv4 address 10.0.0.10 is *concrete*,
  - To an IP stack: 10.0.0.10 is *abstract* and the MAC 08:05:04:03:02:01 is *concrete*,
  - To an Ethernet switch: The MAC 08:05:04:03:02:01 is *abstract*, the resolution to a port in it's CAM table is *concrete*,

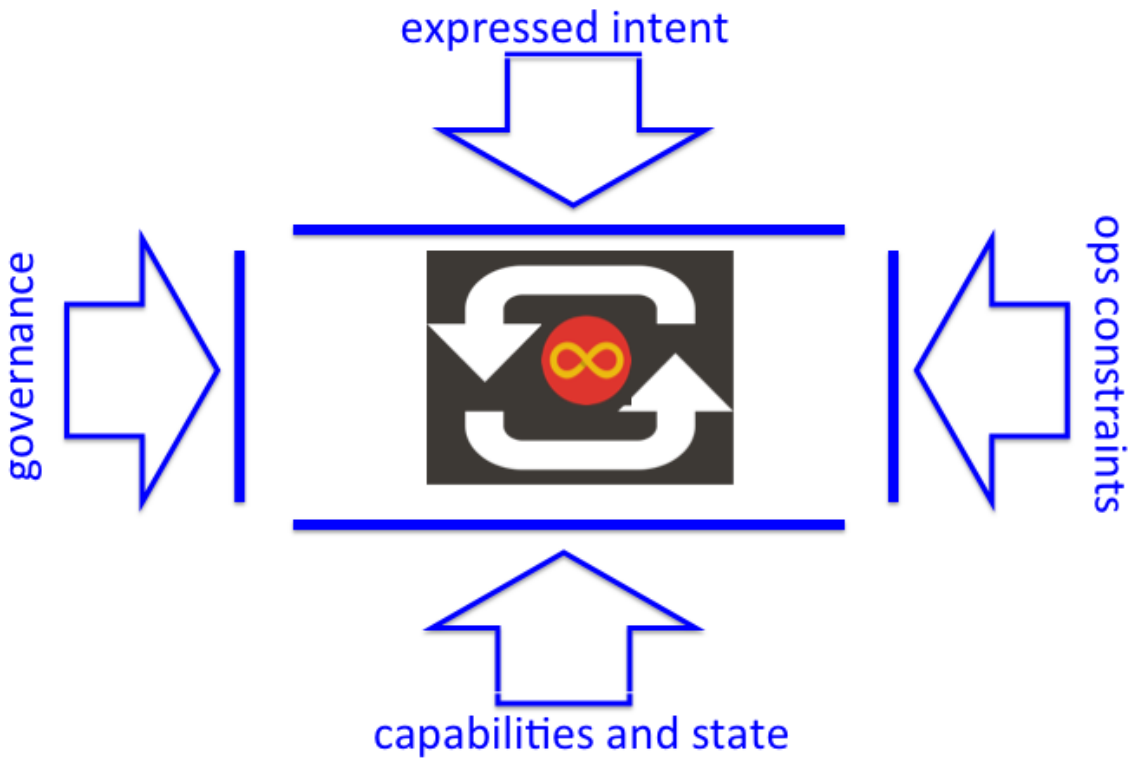

Fig. 26: Intent System Process and Policy Surfaces

- To an optical network: The port maybe *abstract*, yet the optical wavelength is *concrete*.

Note: *This is a very domain specific analogy, tied to something most readers will understand. It in no way implies the **GBP* should be implemented in an OSI type fashion. The premise is that by implementing a full **Intent System**, the user is freed from a lot of the constraints of how the expressed intent is realised.*

It is important to show the overall philosophy of **GBP** as it sets the project's direction.

In this release of OpenDaylight, **GBP** focused on **expressed intent**, **refactoring of how renderers consume and publish Subject Feature Definitions for multi-renderer support**.

GBP Base Architecture and Value Proposition

Terminology

In order to explain the fundamental value proposition of **GBP**, an illustrated example is given. In order to do that some terminology must be defined.

The Access Model is the core of the **GBP** Intent System policy resolution process.

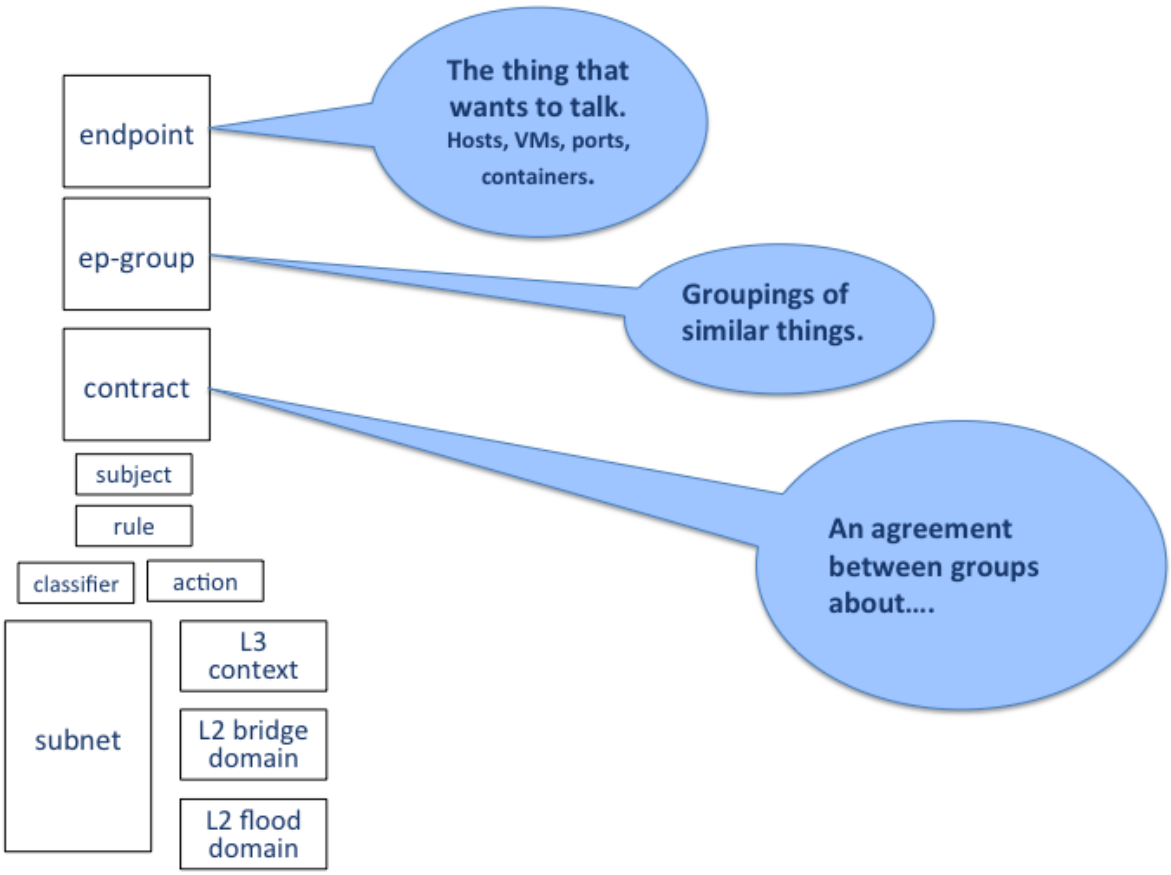

Fig. 27: GBP Access Model Terminology - Endpoints, EndpointGroups, Contract

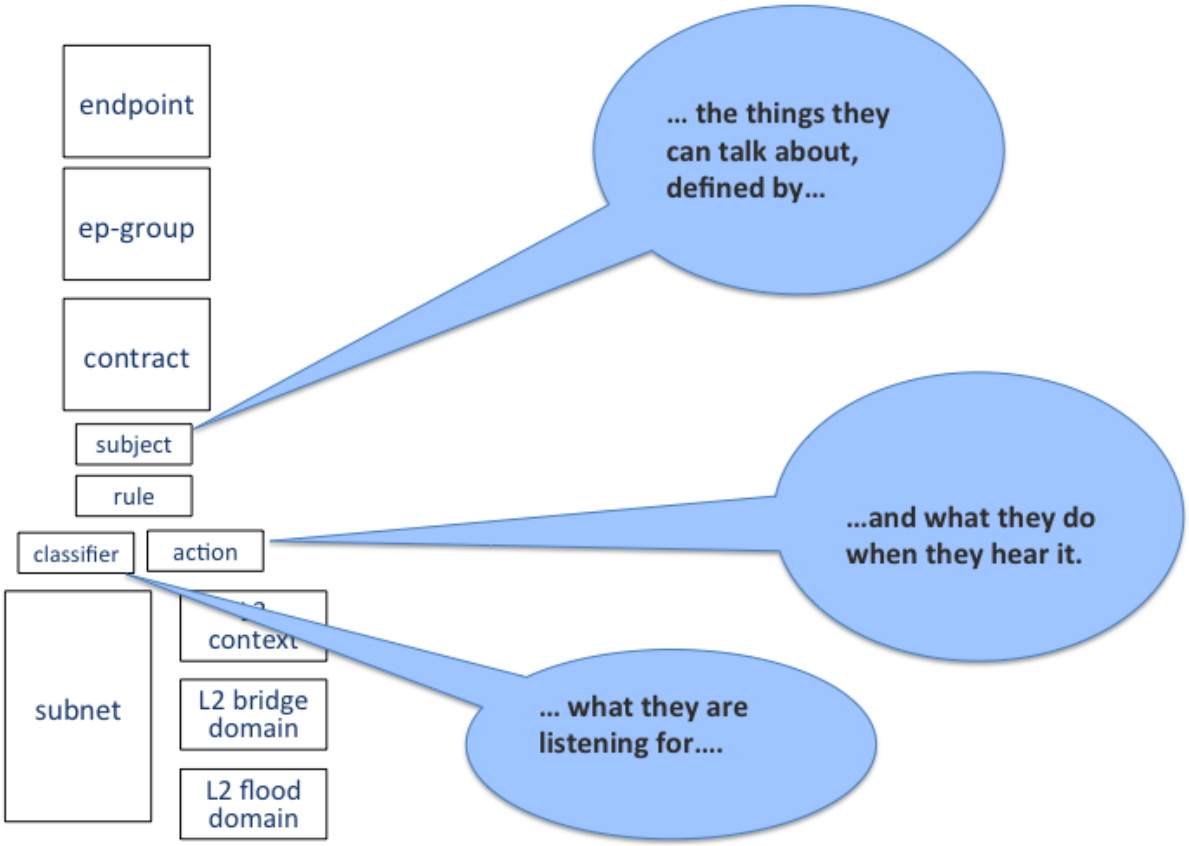

Fig. 28: GBP Access Model Terminology - Subject, Classifier, Action

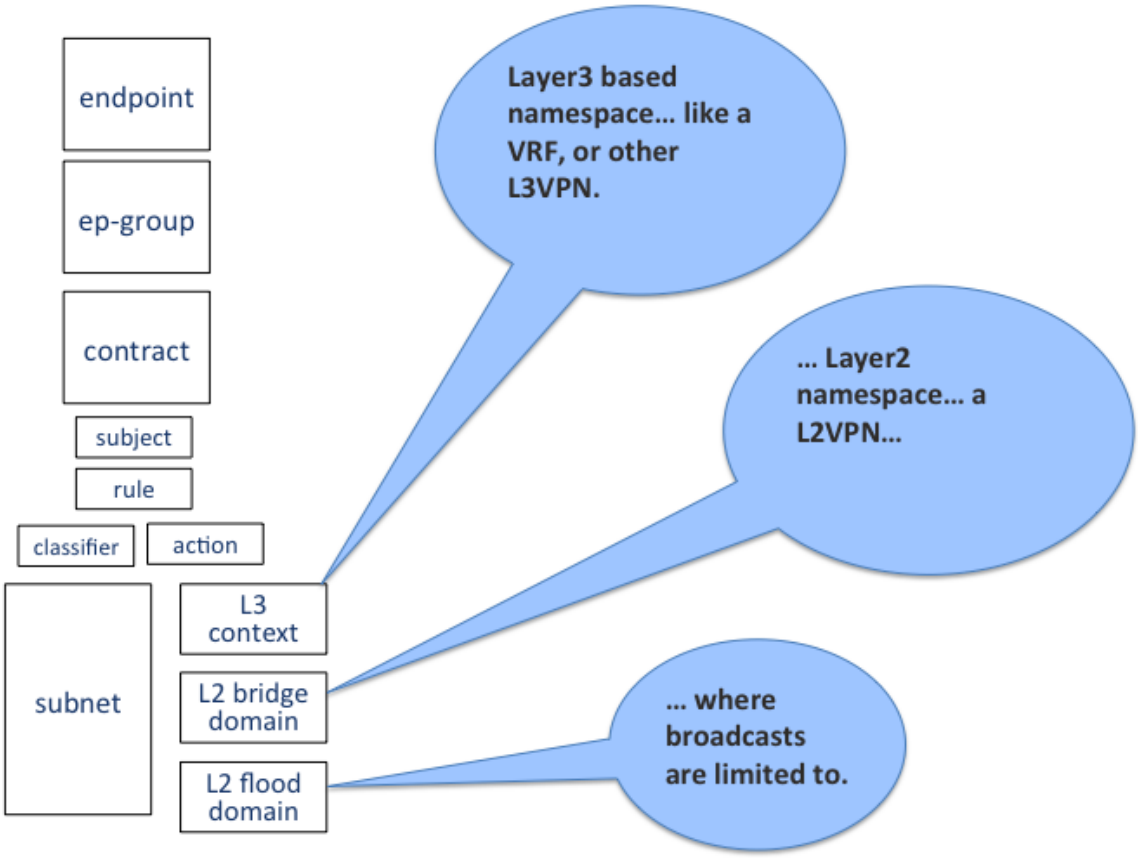

Fig. 29: GBP Forwarding Model Terminology - L3 Context, L2 Bridge Context, L2 Flood Context/Domain, Subnet

- Endpoints:

  Define concrete uniquely identifiable entities. In this release, examples could be a Docker container, or a Neutron port

- EndpointGroups:

  EndpointGroups are sets of endpoints that share a common set of policies. EndpointGroups can participate in contracts that determine the kinds of communication that are allowed. EndpointGroups *consume* and *provide* contracts. They also expose both *requirements and capabilities*, which are labels that help to determine how contracts will be applied. An EndpointGroup can specify a parent EndpointGroup from which it inherits.

- Contracts:

  Contracts determine which endpoints can communicate and in what way. Contracts between pairs of EndpointGroups are selected by the contract selectors defined by the EndpointGroup. Contracts expose qualities, which are labels that can help EndpointGroups to select contracts. Once the contract is selected, contracts have clauses that can match against requirements and capabilities exposed by EndpointGroups, as well as any conditions that may be set on endpoints, in order to activate subjects that can allow specific kinds of communication. A contract is allowed to specify a parent contract from which it inherits.

- Subject

  Subjects describe some aspect of how two endpoints are allowed to communicate. Subjects define an ordered list of rules that will match against the traffic and perform any necessary actions on that traffic. No communication is allowed unless a subject allows that communication.

- Clause

  Clauses are defined as part of a contract. Clauses determine how a contract should be applied to particular endpoints and EndpointGroups. Clauses can match against requirements and capabilities exposed by EndpointGroups, as well as any conditions that may be set on endpoints. Matching clauses define some set of subjects which can be applied to the communication between the pairs of endpoints.

Architecture and Value Proposition

GBP offers an intent based interface, accessed via the *UX*, via the *REST API* or directly from a domain-specific-language such as *Neutron* through a mapping interface.

There are two models in **GBP**:

- the access (or core) model
- the forwarding model

The *classifier* and *action* portions of the model can be thought of as hooks, with their definition provided by each *renderer* about its domain specific capabilities. In **GBP** for this release, there is one renderer, the *OpenFlow Overlay renderer (OfOverlay).*

These hooks are filled with *definitions* of the types of *features* the renderer can provide the *subject*, and are called **subject-feature-definitions**.

This means an *expressed intent* can be fulfilled by, and across, multiple renderers simultaneously, without any specific provisioning from the consumer of **GBP**.

Since **GBP** is implemented in OpenDaylight, which is an SDN controller, it also must address networking. This is done via the *forwarding model*, which is domain specific to networking, but could be applied to many different *types* of networking.

Each endpoint is provisioned with a *network-containment*. This can be a:

- subnet

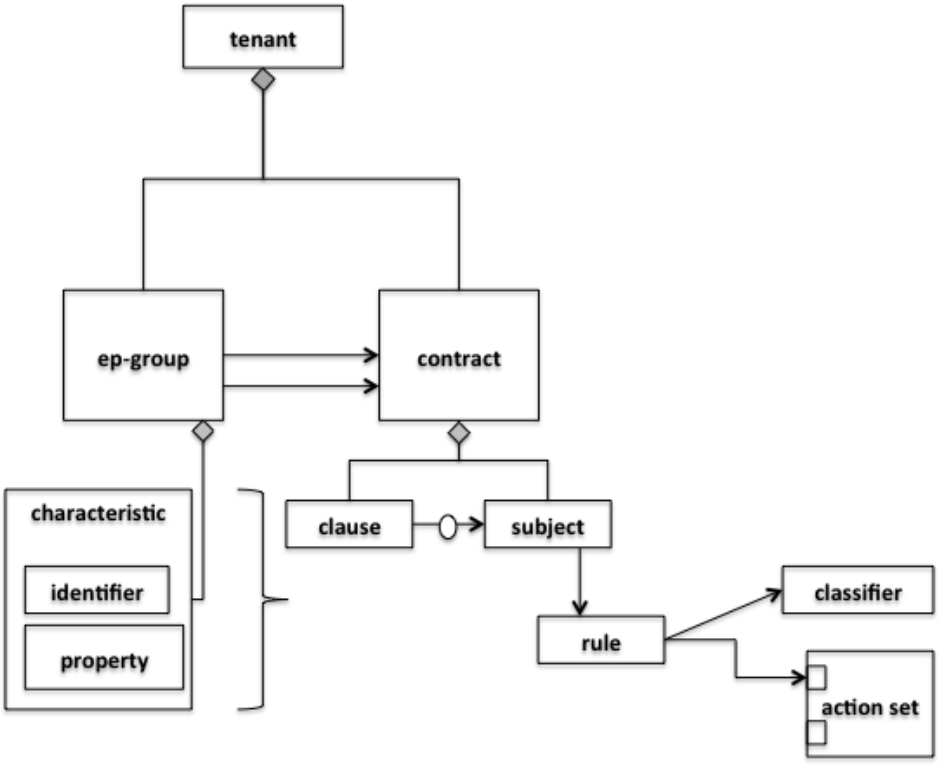

Fig. 30: GBP Access (or Core) Model

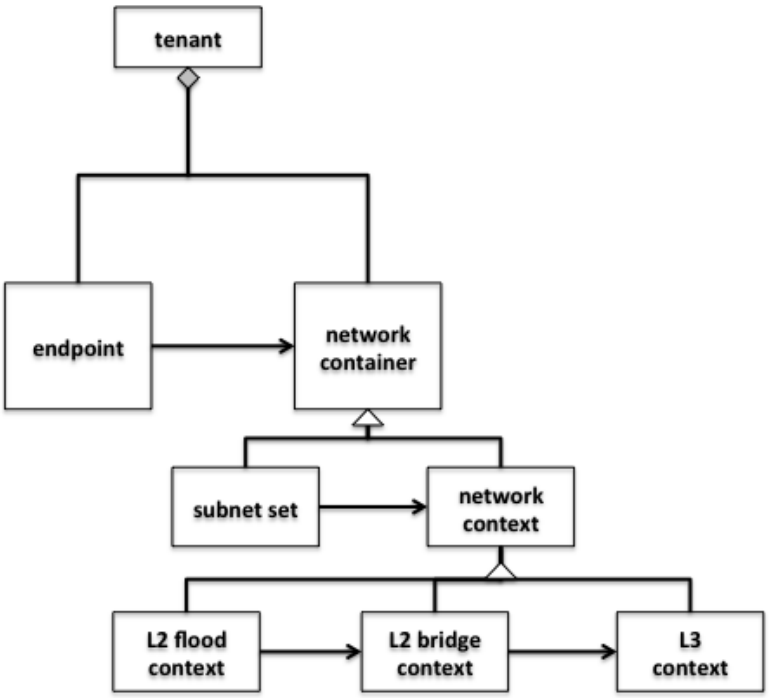

Fig. 31: GBP Forwarding Model

  - normal IP stack behaviour, where ARP is performed in subnet, and for out of subnet, traffic is sent to default gateway.
  - a subnet can be a child of any of the below forwarding model contexts, but typically would be a child of a flood-domain

- L2 flood-domain
  - allows flooding behaviour.
  - is a n:1 child of a bridge-domain
  - can have multiple children
- L2 bridge-domain
  - is a layer2 namespace
  - is the realm where traffic can be sent at layer 2
  - is a n:1 child of a L3 context
  - can have multiple children
- L3 context
  - is a layer3 namespace
  - is the realm where traffic is passed at layer 3
  - is a n:1 child of a tenant
  - can have multiple children

A simple example of how the access and forwarding models work is as follows:

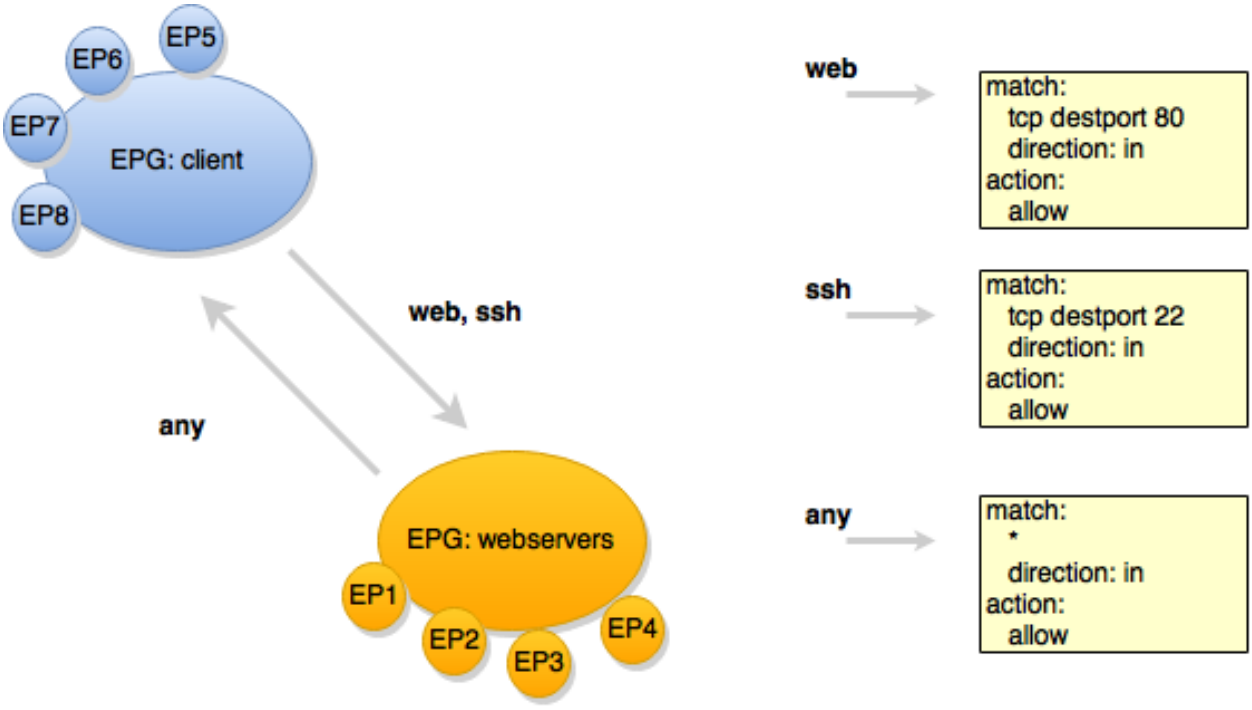

Fig. 32: GBP Endpoints, EndpointGroups and Contracts

In this example, the **EPG:webservers** is *providing* the *web* and *ssh* contracts. The **EPG:client** is consuming those contracts. **EPG:client** is providing the *any* contract, which is consumed by **EPG:webservers**.

The *direction* keyword is always from the perspective of the *provider* of the contract. In this case contract *web*, being *provided* by **EPG:webservers**, with the classifier to match TCP destination port 80, means:

- packets with a TCP destination port of 80
- sent to (*in*) endpoints in the **EPG:webservers**
- will be *allowed*.

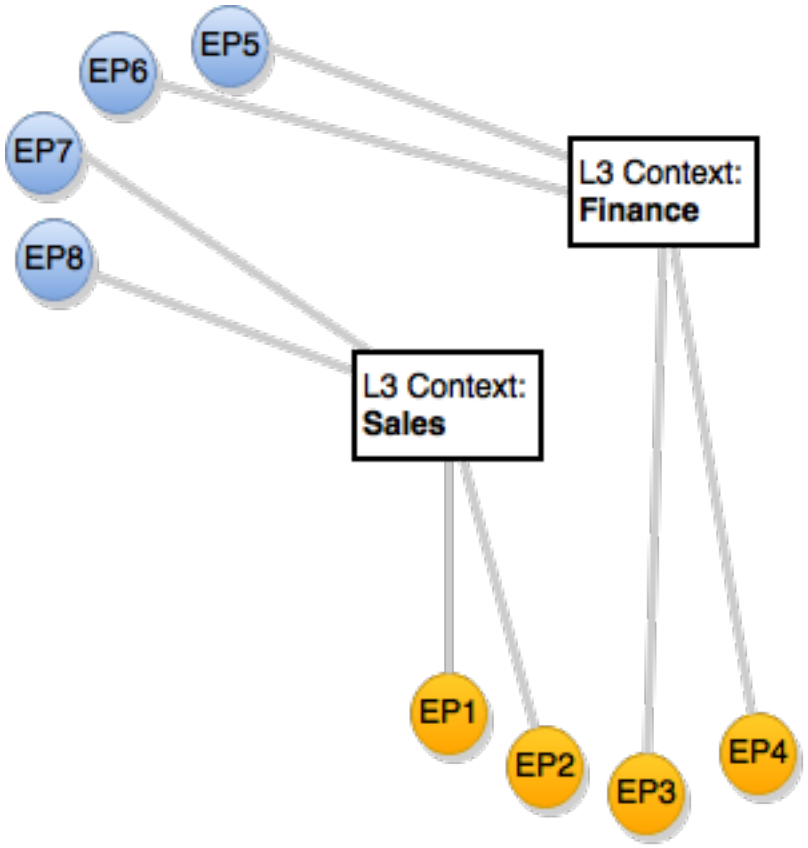

Fig. 33: GBP Endpoints and the Forwarding Model

When the forwarding model is considered in the figure above, it can be seen that even though all endpoints are communicating using a common set of contracts, their forwarding is *contained* by the forwarding model contexts or namespaces. In the example shown, the endpoints associated with a *network-containment* that has an ultimate parent of *L3Context:Sales* can only communicate with other endpoints within this L3Context. In this way L3VPN services can be implemented without any impact to the **Intent** of the contract.

High-level implementation Architecture

The overall architecture, including *Neutron* domain specific mapping, and the *OpenFlow Overlay renderer* looks as so:

The major benefit of this architecture is that the mapping of the domain-specific-language is completely separate and independent of the underlying renderer implementation.

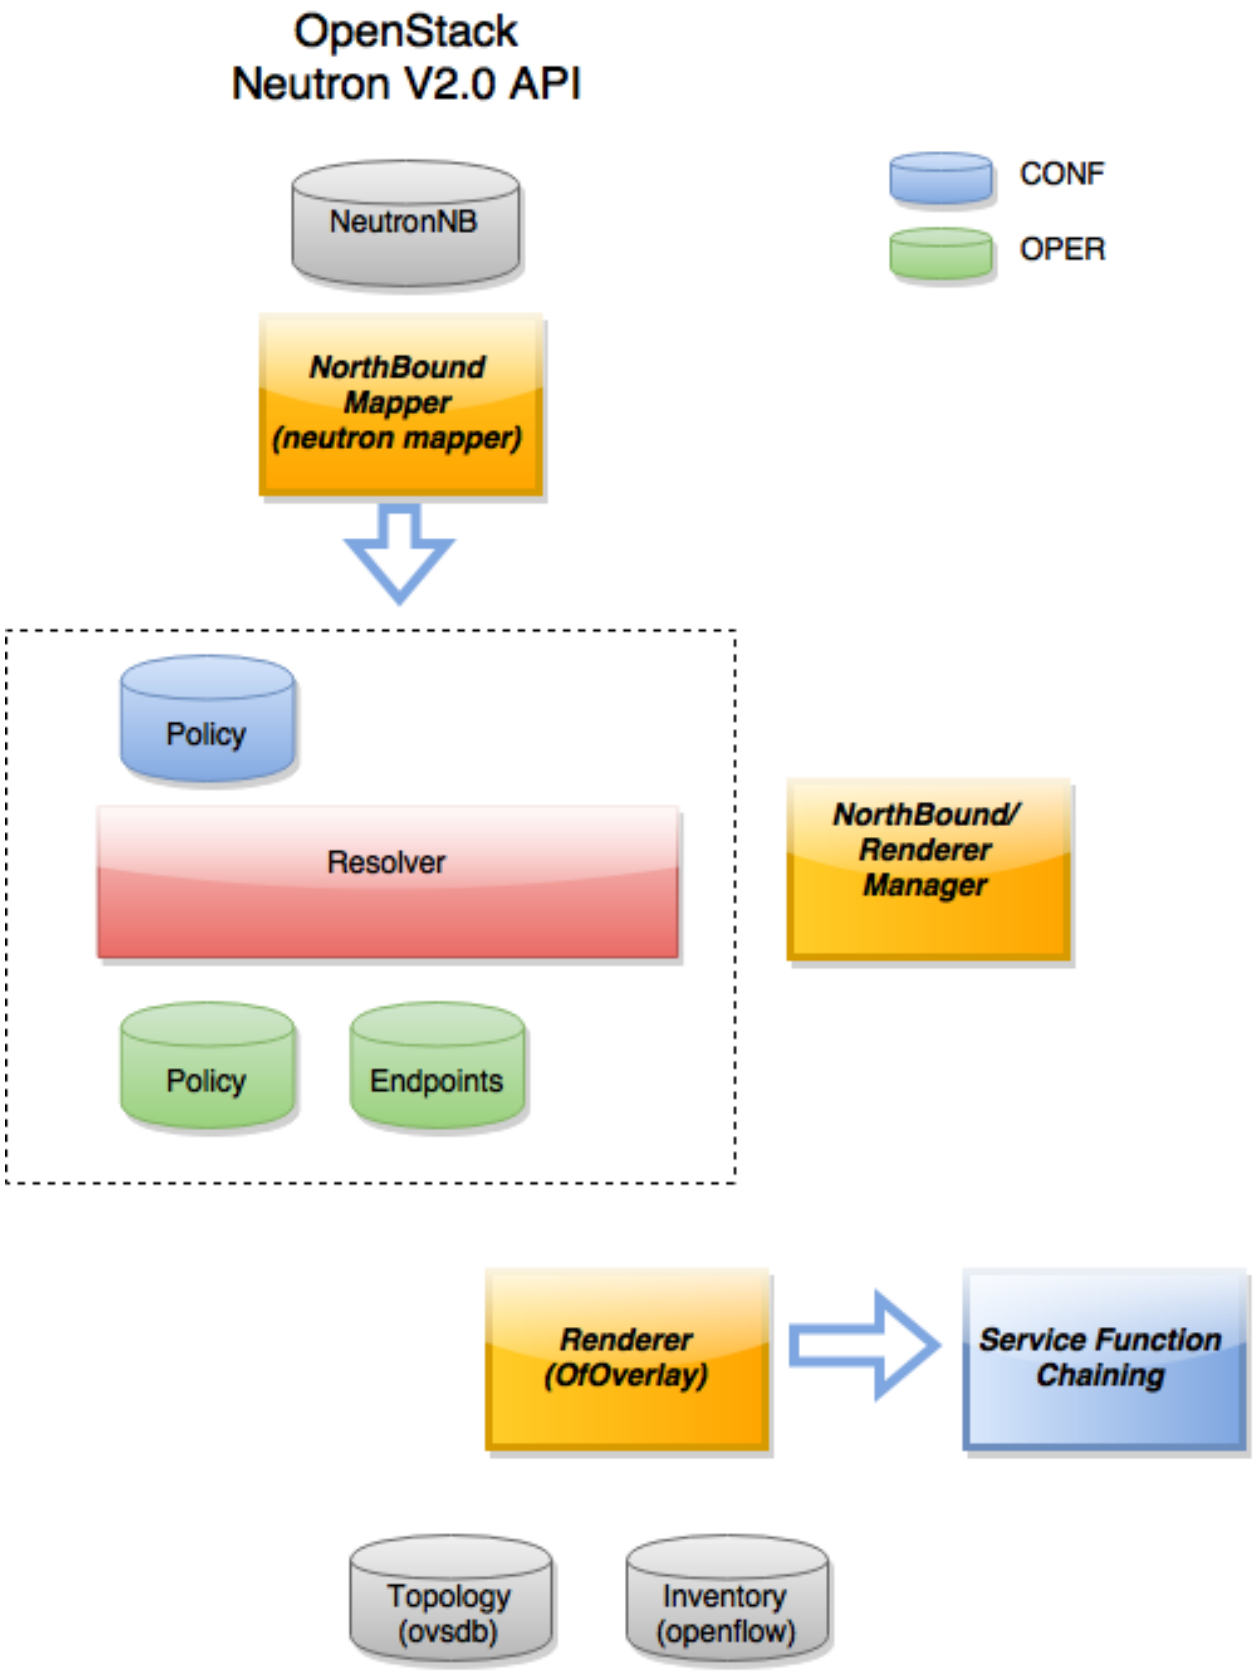
OpenStack
Neutron V2.0 API
NeutronNB
CONF
OPER
NorthBound
Mapper
(neutron mapper)
Policy
Resolver
Policy
Endpoints
NorthBound/
Renderer
Manager
Renderer
(OfOverlay)
Service Function
Chaining
Topology
(ovsdb)
Inventory
(openflow)

For instance, using the *Neutron Mapper*, which maps the Neutron API to the **GBP** core model, any contract automatically generated from this mapping can be augmented via the *UX* to use *Service Function Chaining*, a capability not currently available in OpenStack Neutron.

When another renderer is added, for instance, NetConf, the same policy can now be leveraged across NetConf devices simultaneously:

As other domain-specific mappings occur, they too can leverage the same renderers, as the renderers only need to implement the **GBP** access and forwarding models, and the domain-specific mapping need only manage mapping to the access and forwarding models. For instance:

In summary, the **GBP** architecture:

- separates concerns: the Expressed Intent is kept completely separated from the underlying renderers.
- is cohesive: each part does it's part and it's part only
- is scalable: code can be optimised around model mapping/implementation, and functionality re-used

Policy Resolution

Contract Selection

The first step in policy resolution is to select the contracts that are in scope.

EndpointGroups participate in contracts either as a *provider* or as a *consumer* of a contract. Each EndpointGroup can participate in many contracts at the same time, but for each contract it can be in only one role at a time. In addition, there are two ways for an EndpointGroup to select a contract: either with either a:

- *named selector*

  Named selectors simply select a specific contract by its contract ID.

- target selector.

  Target selectors allow for additional flexibility by matching against *qualities* of the contract's *target.*

Thus, there are a total of 4 kinds of contract selector:

- provider named selector

  Select a contract by contract ID, and participate as a provider.

- provider target selector

  Match against a contract's target with a quality matcher, and participate as a provider.

- consumer named selector

  Select a contract by contract ID, and participate as a consumer.

- consumer target selector

  Match against a contract's target with a quality matcher, and participate as a consumer.

To determine which contracts are in scope, contracts are found where either the source EndpointGroup selects a contract as either a provider or consumer, while the destination EndpointGroup matches against the same contract in the corresponding role. So if endpoint *x* in EndpointGroup *X* is communicating with endpoint *y* in EndpointGroup *Y*, a contract *C* is in scope if either *X* selects *C* as a provider and *Y* selects *C* as a consumer, or vice versa.

The details of how quality matchers work are described further in *Matchers*. Quality matchers provide a flexible mechanism for contract selection based on labels.

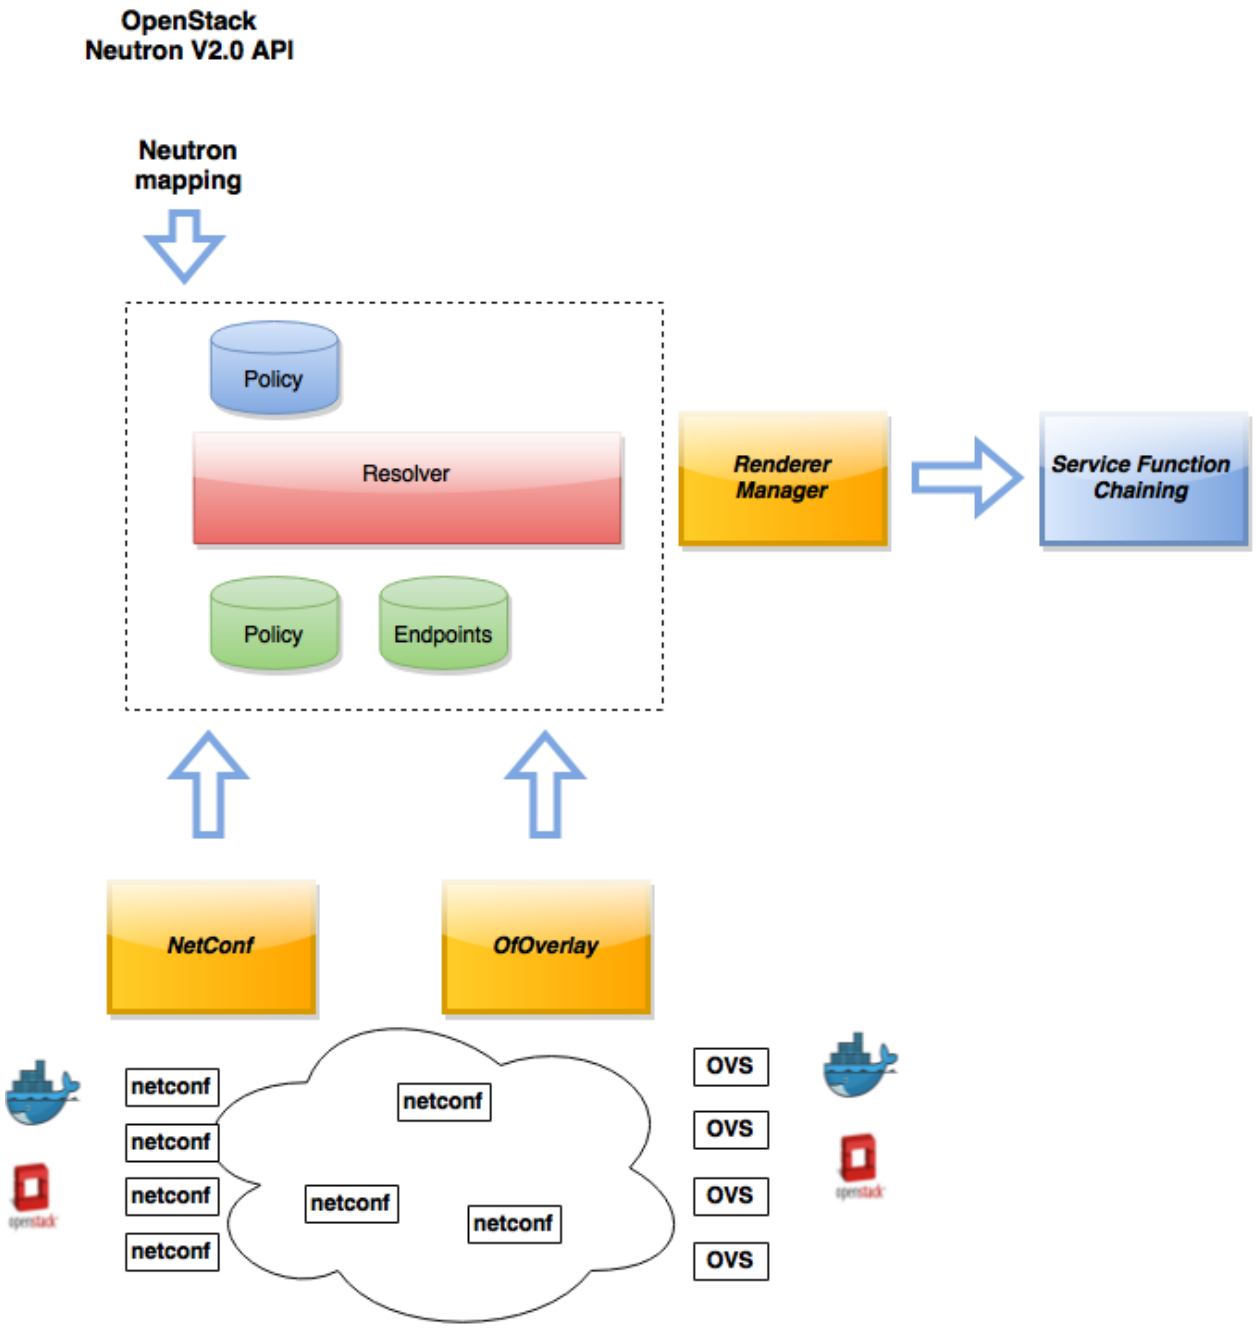

Fig. 35: GBP High Level Architecture - adding a renderer

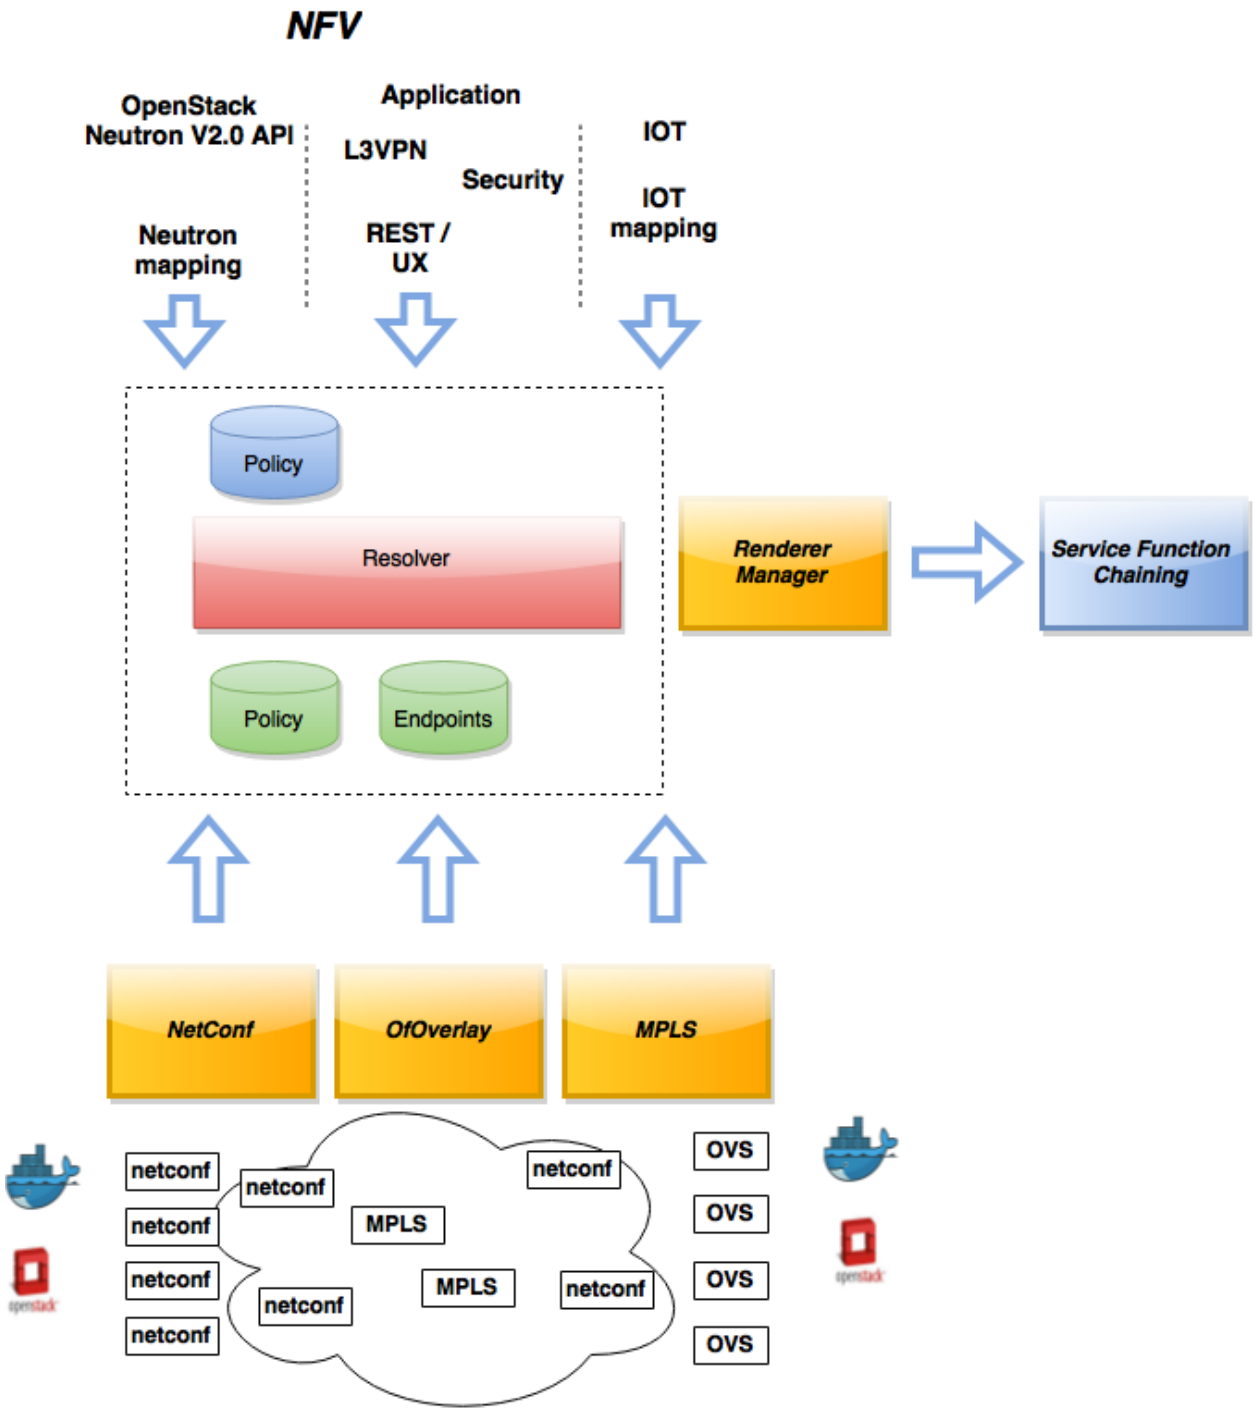

Fig. 36: GBP High Level Architecture - adding a renderer

The end result of the contract selection phase can be thought of as a set of tuples representing selected contract scopes. The fields of the tuple are:

- Contract ID
- The provider EndpointGroup ID
- The name of the selector in the provider EndpointGroup that was used to select the contract, called the *matching provider selector.*
- The consumer EndpointGroup ID
- The name of the selector in the consumer EndpointGroup that was used to select the contract, called the *matching consumer selector.*

The result is then stored in the datastore under **Resolved Policy**.

Subject Selection

The second phase in policy resolution is to determine which subjects are in scope. The subjects define what kinds of communication are allowed between endpoints in the EndpointGroups. For each of the selected contract scopes from the contract selection phase, the subject selection procedure is applied.

Labels called, capabilities, requirements and conditions are matched against to bring a Subject *into scope*. EndpointGroups have capabilities and requirements, while endpoints have conditions.

Requirements and Capabilities

When acting as a *provider*, EndpointGroups expose *capabilities,* which are labels representing specific pieces of functionality that can be exposed to other EndpointGroups that may meet functional requirements of those EndpointGroups.

When acting as a *consumer*, EndpointGroups expose *requirements*, which are labels that represent that the EndpointGroup requires some specific piece of functionality.

As an example, we might create a capability called "user-database" which indicates that an EndpointGroup contains endpoints that implement a database of users.

We might create a requirement also called "user-database" to indicate an EndpointGroup contains endpoints that will need to communicate with the endpoints that expose this service.

Note that in this example the requirement and capability have the same name, but the user need not follow this convention.

The matching provider selector (that was used by the provider EndpointGroup to select the contract) is examined to determine the capabilities exposed by the provider EndpointGroup for this contract scope.

The provider selector will have a list of capabilities either directly included in the provider selector or inherited from a parent selector or parent EndpointGroup. (See *Inheritance*).

Similarly, the matching consumer selector will expose a set of requirements.

Conditions

Endpoints can have *conditions*, which are labels representing some relevant piece of operational state related to the endpoint.

An example of a condition might be “malware-detected,” or “authentication-succeeded.” Conditions are used to affect how that particular endpoint can communicate.

To continue with our example, the “malware-detected” condition might cause an endpoint’s connectivity to be cut off, while “authentication-succeeded” might open up communication with services that require an endpoint to be first authenticated and then forward its authentication credentials.

Clauses

Clauses perform the actual selection of subjects. A clause has lists of matchers in two categories. In order for a clause to become active, all lists of matchers must match. A matching clause will select all the subjects referenced by the clause. Note that an empty list of matchers counts as a match.

The first category is the consumer matchers, which match against the consumer EndpointGroup and endpoints. The consumer matchers are:

- Group Idenfication Constraint: Requirement matchers

  Matches against requirements in the matching consumer selector.

- Group Identification Constraint: GroupName

  Matches against the group name

- Consumer condition matchers

  Matches against conditions on endpoints in the consumer EndpointGroup

- Consumer Endpoint Identification Constraint

  Label based criteria for matching against endpoints. In this release this can be used to label endpoints based on IpPrefix.

The second category is the provider matchers, which match against the provider EndpointGroup and endpoints. The provider matchers are:

- Group Idenfication Constraint: Capability matchers

  Matches against capabilities in the matching provider selector.

- Group Identification Constraint: GroupName

  Matches against the group name

- Consumer condition matchers

  Matches against conditions on endpoints in the provider EndpointGroup

- Consumer Endpoint Identification Constraint

  Label based criteria for matching against endpoints. In this release this can be used to label endpoints based on IpPrefix.

Clauses have a list of subjects that apply when all the matchers in the clause match. The output of the subject selection phase logically is a set of subjects that are in scope for any particular pair of endpoints.

Rule Application

Now subjects have been selected that apply to the traffic between a particular set of endpoints, policy can be applied to allow endpoints to communicate. The applicable subjects from the previous step will each contain a set of rules.

Rules consist of a set of *classifiers* and a set of *actions*. Classifiers match against traffic between two endpoints. An example of a classifier would be something that matches against all TCP traffic on port 80, or one that matches against HTTP traffic containing a particular cookie. Actions are specific actions that need to be taken on the traffic before it reaches its destination. Actions could include tagging or encapsulating the traffic in some way, redirecting the traffic, or applying a *service function chain*.

Rules, subjects, and actions have an *order* parameter, where a lower order value means that a particular item will be applied first. All rules from a particular subject will be applied before the rules of any other subject, and all actions from a particular rule will be applied before the actions from another rule. If more than item has the same order parameter, ties are broken with a lexicographic ordering of their names, with earlier names having logically lower order.

Matchers

Matchers specify a set of labels (which include requirements, capabilities, conditions, and qualities) to match against. There are several kinds of matchers that operate similarly:

- Quality matchers

  used in target selectors during the contract selection phase. Quality matchers provide a more advanced and flexible way to select contracts compared to a named selector.

- Requirement and capability matchers

  used in clauses during the subject selection phase to match against requirements and capabilities on EndpointGroups

- Condition matchers

  used in clauses during the subject selection phase to match against conditions on endpoints

A matcher is, at its heart, fairly simple. It will contain a list of label names, along with a *match type*. The match type can be either:

- "all"

  which means the matcher matches when all of its labels match

- "any"

  which means the matcher matches when any of its labels match,

- "none"

  which means the matcher matches when none of its labels match.

Note a *match all* matcher can be made by matching against an empty set of labels with a match type of "all."

Additionally each label to match can optionally include a relevant name field. For quality matchers, this is a target name. For capability and requirement matchers, this is a selector name. If the name field is specified, then the matcher will only match against targets or selectors with that name, rather than any targets or selectors.

Inheritance

Some objects in the system include references to parents, from which they will inherit definitions. The graph of parent references must be loop free. When resolving names, the resolution system must detect loops and raise an exception. Objects that are part of these loops may be considered as though they are not defined at all. Generally, inheritance works by simply importing the objects in the parent into the child object. When there are objects with the same name in the child object, then the child object will override the parent object according to rules which are specific to the type of object. We'll next explore the detailed rules for inheritance for each type of object

EndpointGroups

EndpointGroups will inherit all their selectors from their parent EndpointGroups. Selectors with the same names as selectors in the parent EndpointGroups will inherit their behavior as defined below.

Selectors

Selectors include provider named selectors, provider target selectors, consumer named selectors, and consumer target selectors. Selectors cannot themselves have parent selectors, but when selectors have the same name as a selector of the same type in the parent EndpointGroup, then they will inherit from and override the behavior of the selector in the parent EndpointGroup.

Named Selectors

Named selectors will add to the set of contract IDs that are selected by the parent named selector.

Target Selectors

A target selector in the child EndpointGroup with the same name as a target selector in the parent EndpointGroup will inherit quality matchers from the parent. If a quality matcher in the child has the same name as a quality matcher in the parent, then it will inherit as described below under Matchers.

Contracts

Contracts will inherit all their targets, clauses and subjects from their parent contracts. When any of these objects have the same name as in the parent contract, then the behavior will be as defined below.

Targets

Targets cannot themselves have a parent target, but it may inherit from targets with the same name as the target in a parent contract. Qualities in the target will be inherited from the parent. If a quality with the same name is defined in the child, then this does not have any semantic effect except if the quality has its inclusion-rule parameter set to "exclude." In this case, then the label should be ignored for the purpose of matching against this target.

Subjects

Subjects cannot themselves have a parent subject, but it may inherit from a subject with the same name as the subject in a parent contract. The order parameter in the child subject, if present, will override the order parameter in the parent subject. The rules in the parent subject will be added to the rules in the child subject. However, the rules will not override rules of the same name. Instead, all rules in the parent subject will be considered to run with a higher order than all rules in the child; that is all rules in the child will run before any rules in the parent. This has the effect of overriding any rules in the parent without the potentially-problematic semantics of merging the ordering.

Clauses

Clauses cannot themselves have a parent clause, but it may inherit from a clause with the same name as the clause in a parent contract. The list of subject references in the parent clause will be added to the list of subject references in the child clause. This is just a union operation. A subject reference that refers to a subject name in the parent contract might have that name overridden in the child contract. Each of the matchers in the clause are also inherited by the child clause. Matchers in the child of the same name and type as a matcher from the parent will inherit from and override the parent matcher. See below under Matchers for more information.

Matchers

Matchers include quality matchers, condition matchers, requirement matchers, and capability matchers. Matchers cannot themselves have parent matchers, but when there is a matcher of the same name and type in the parent object, then the matcher in the child object will inherit and override the behavior of the matcher in the parent object. The match type, if specified in the child, overrides the value specified in the parent. Labels are also inherited from the parent object. If there is a label with the same name in the child object, this does not have any semantic effect except if the label has its inclusion-rule parameter set to "exclude." In this case, then the label should be ignored for the purpose of matching. Otherwise, the label with the same name will completely override the label from the parent.

Using the GBP UX interface

Overview

These following components make up this application and are described in more detail in following sections:

- Basic view
- Governance view
- Policy Expression view
- Wizard view

The **GBP** UX is access via:

```
http://<odl controller>:8181/index.html
```

Basic view

Basic view contains 5 navigation buttons which switch user to the desired section of application:

- Governance – switch to the Governance view (middle of graphic has the same function)
- Renderer configuration – switch to the Policy expression view with Renderers section expanded
- Policy expression – switch to the Policy expression view with Policy section expanded
- Operational constraints – placeholder for development in next release

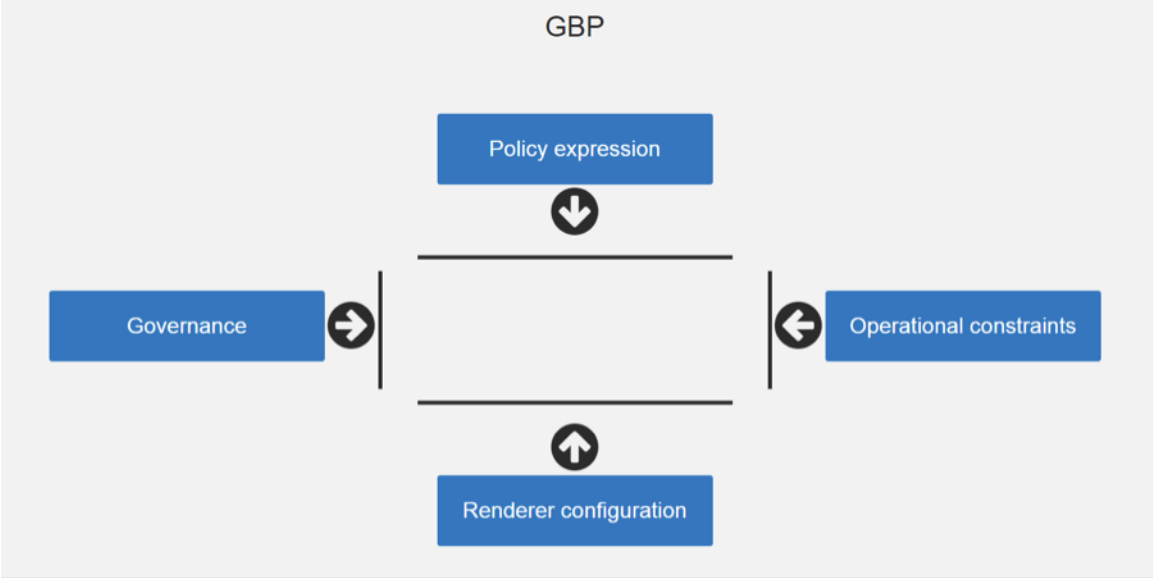

Fig. 37: Basic view

Governance view

Governance view consists from three columns.

Governance view – Basic view – Left column

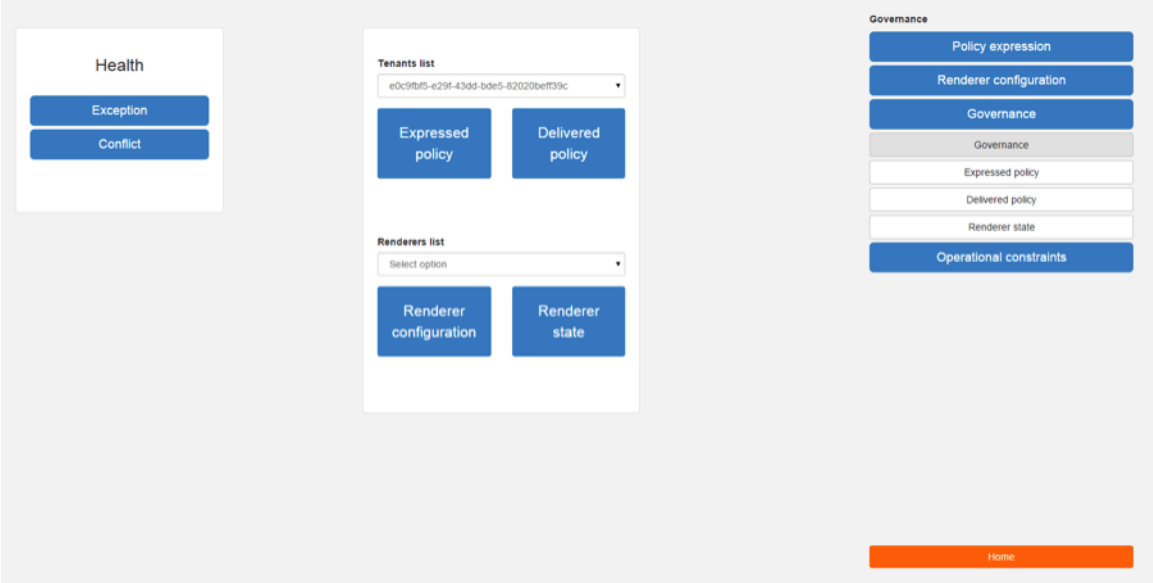

Fig. 38: Governance view

In the left column is Health section with Exception and Conflict buttons with no functionality yet. This is a placeholder for development in further releases.

Governance view – Basic view – Middle column

In the top half of this section is select box with list of tenants for select. Once the tenant is selected, all sub sections in application operate and display data with actual selected tenant.

Below the select box are buttons which display Expressed or Delivered policy of Governance section. In the bottom half of this section is select box with list of renderers for select. There is currently only *OfOverlay* renderer available.

Below the select box is Renderer configuration button, which switch the app into the Policy expression view with Renderers section expanded for performing CRUD operations. Renderer state button display Renderer state view.

Governance view – Basic view – Right column

In the bottom part of the right section of Governance view is Home button which switch the app to the Basic view.

In the top part is situated navigation menu with four main sections.

Policy expression button expand/collapse sub menu with three main parts of Policy expression. By clicking on sub menu buttons, user will be switched into the Policy expressions view with appropriate section expanded for performing CRUD operations.

Renderer configuration button switches user into the Policy expressions view.

Governance button expand/collapse sub menu with four main parts of Governance section. Sub menu buttons of Governance section display appropriate section of Governance view.

Operational constraints have no functionality yet, and is a placeholder for development in further releases.

Below the menu is place for view info section which displays info about actual selected element from the topology (explained below).

Governance view – Expressed policy

In this view are displayed contracts with their consumed and provided EndpointGroups of actual selected tenant, which can be changed in select box in the upper left corner.

By single-clicking on any contract or EPG, the data of actual selected element will be shown in the right column below the menu. A Manage button launches a display wizard window for managing configuration of items such as *Service Function Chaining*.

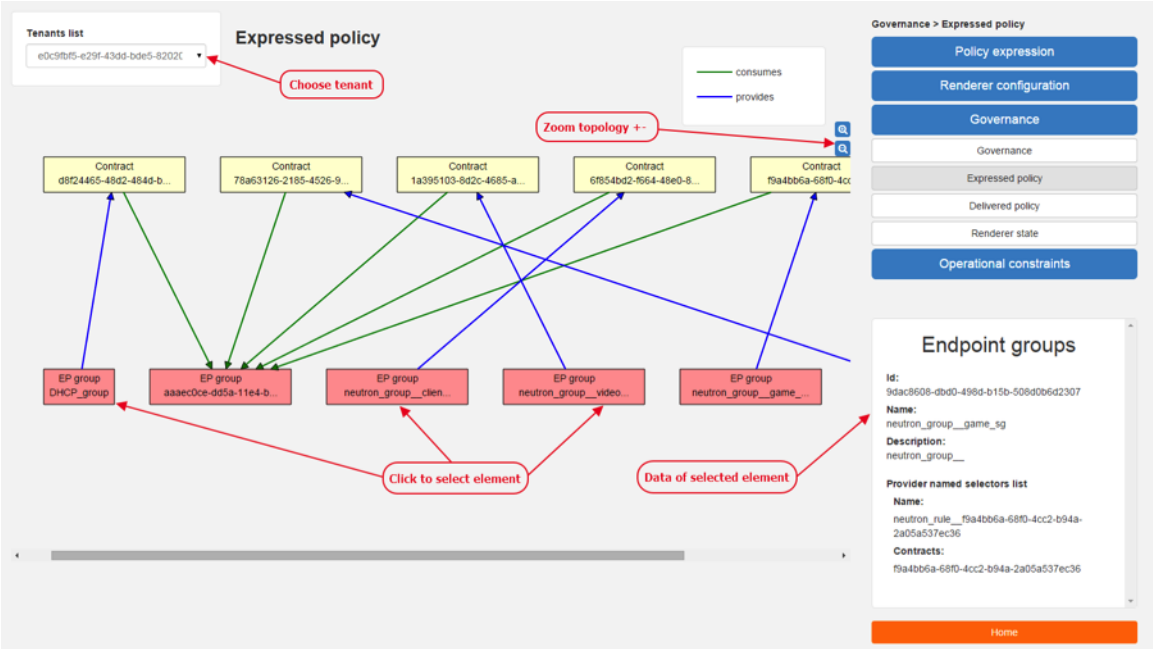

Fig. 39: Expressed policy

Governance view – Delivered policy In this view are displayed subjects with their consumed and provided Endpoint-Groups of actual selected tenant, which can be changed in select box in the upper left corner.

By single-clicking on any subject or EPG, the data of actual selected element will be shown in the right column below the menu.

By double-click on subject the subject detail view will be displayed with subject's rules of actual selected subject, which can be changed in select box in the upper left corner.

By single-clicking on rule or subject, the data of actual selected element will be shown in the right column below the menu.

By double-clicking on EPG in Delivered policy view, the EPG detail view will be displayed with EPG's endpoints of actual selected EPG, which can be changed in select box in the upper left corner.

By single-clicking on EPG or endpoint the data of actual selected element will be shown in the right column below the menu.

Governance view – Renderer state

In this part are displayed Subject feature definition data with two main parts: Action definition and Classifier definition.

By clicking on the down/right arrow in the circle is possible to expand/hide data of appropriate container or list. Next to the list node are displayed names of list's elements where one is always selected and element's data are shown (blue line under the name).

By clicking on names of children nodes is possible to select desired node and node's data will be displayed.

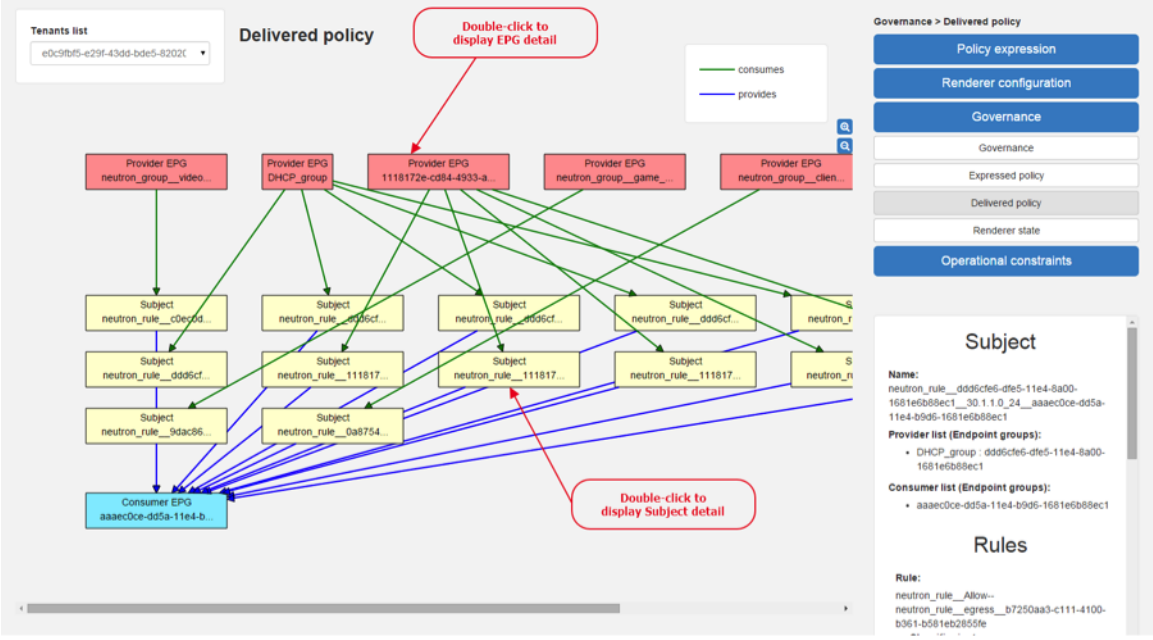

Fig. 40: Delivered policy

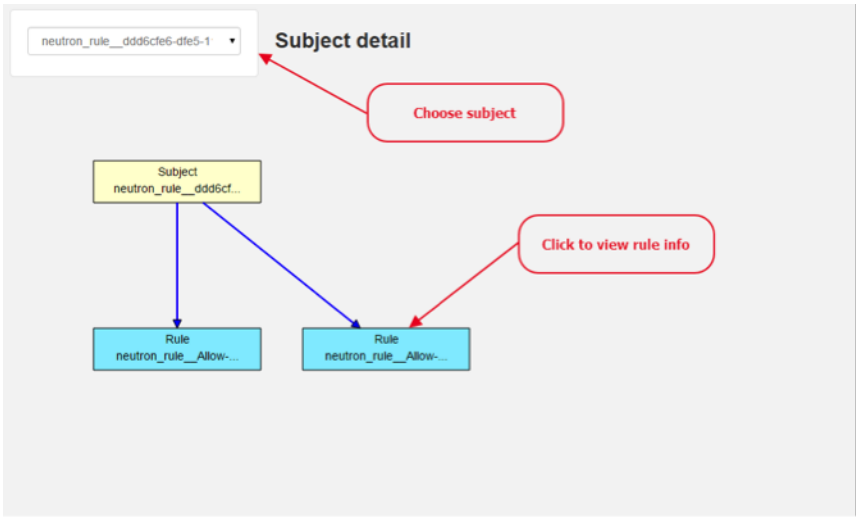

Fig. 41: Subject detail

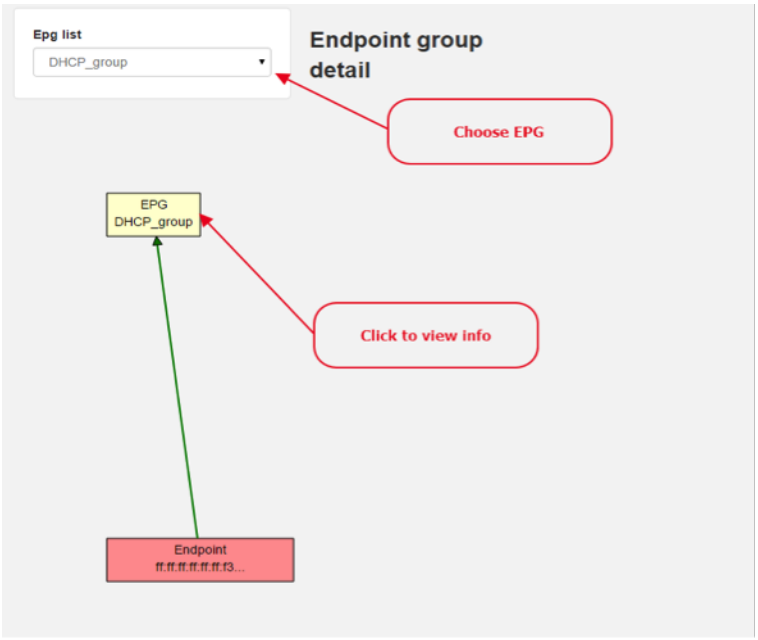

Fig. 42: EPG detail

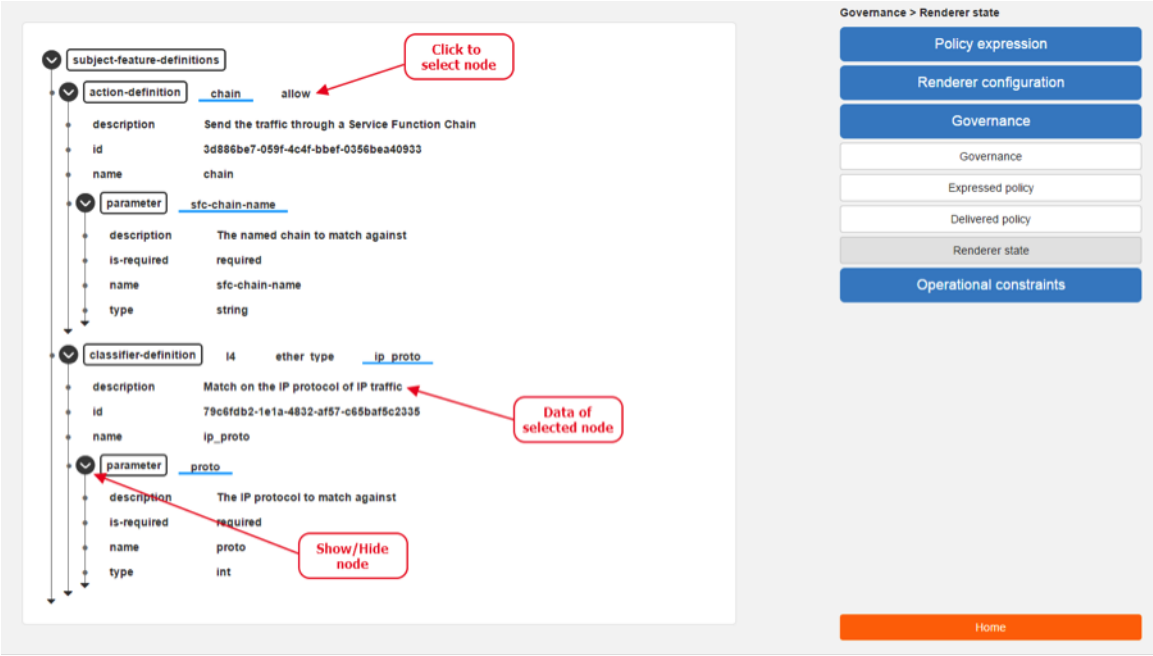

Fig. 43: Renderer state

Policy expression view

In the left part of this view is placed topology of actual selected elements with the buttons for switching between types of topology at the bottom.

Right column of this view contains four parts. At the top of this column are displayed breadcrumbs with actual position in the application.

Below the breadcrumbs is select box with list of tenants for select. In the middle part is situated navigation menu, which allows switch to the desired section for performing CRUD operations.

At the bottom is quick navigation menu with Access Model Wizard button which display Wizard view, Home button which switch application to the Basic view and occasionally Back button, which switch application to the upper section.

Policy expression - Navigation menu

To open Policy expression, select Policy expression from the GBP Home screen.

In the top of navigation box you can select the tenant from the tenants list to activate features addicted to selected tenant.

In the right menu, by default, the Policy menu section is expanded. Subitems of this section are modules for CRUD (creating, reading, updating and deleting) of tenants, EndpointGroups, contracts, L2/L3 objects.

- Section Renderers contains CRUD forms for Classifiers and Actions.
- Section Endpoints contains CRUD forms for Endpoint and L3 prefix endpoint.

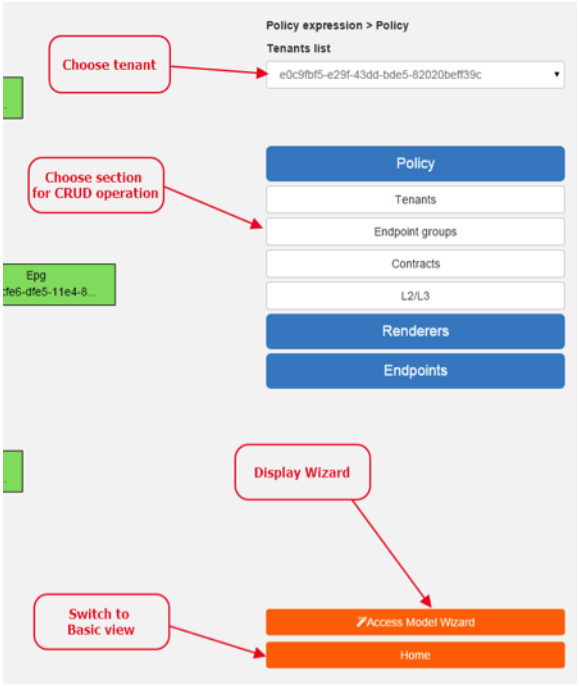

Fig. 44: Navigation menu

Policy expression - Types of topology

There are three different types of topology:

- Configured topology - EndpointGroups and contracts between them from CONFIG datastore

Fig. 45: CRUD operations

- Operational topology - displays same information but is based on operational data.
- L2/L3 - displays relationships between L3Contexts, L2 Bridge domains, L2 Flood domains and Subnets.

Policy expression - CRUD operations

In this part are described basic flows for viewing, adding, editing and deleting system elements like tenants, Endpoint-Groups etc.

Tenants

To edit tenant objects click the Tenants button in the right menu. You can see the CRUD form containing tenants list and control buttons.

To add new tenant, click the Add button This will display the form for adding a new tenant. After filling tenant attributes Name and Description click Save button. Saving of any object can be performed only if all the object attributes are filled correctly. If some attribute doesn't have correct value, exclamation mark with mouse-over tooltip will be displayed next to the label for the attribute. After saving of tenant the form will be closed and the tenants list will be set to default value.

To view an existing tenant, select the tenant from the select box Tenants list. The view form is read-only and can be closed by clicking cross mark in the top right of the form.

To edit selected tenant, click the Edit button, which will display the edit form for selected tenant. After editing the Name and Description of selected tenant click the Save button to save selected tenant. After saving of tenant the edit form will be closed and the tenants list will be set to default value.

To delete tenant select the tenant from the Tenants list and click Delete button.

To return to the Policy expression click Back button on the bottom of window.

EndpointGroups

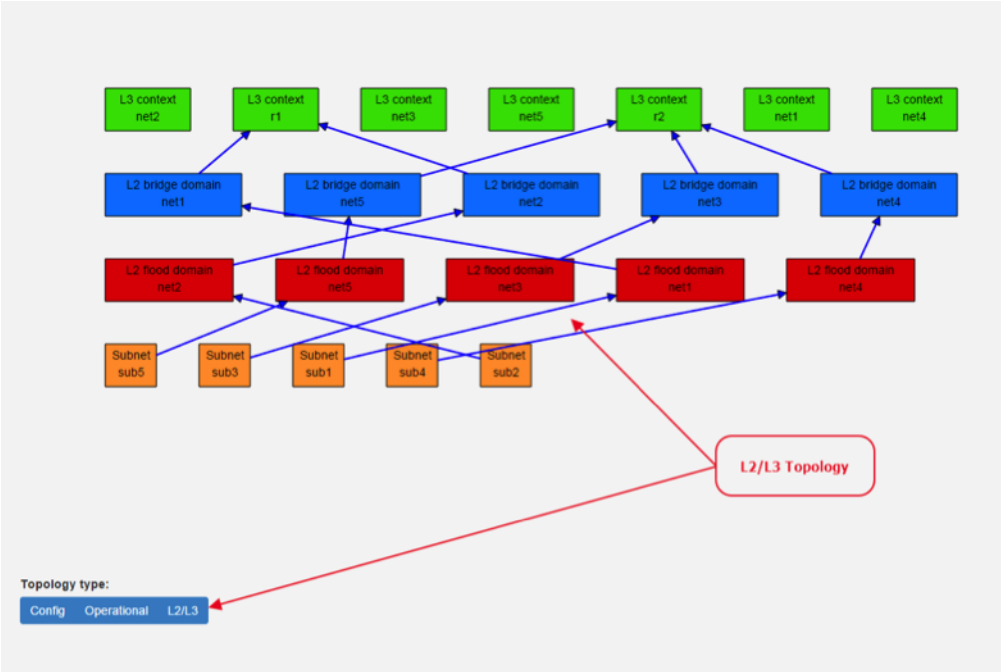

Fig. 46: L2/L3 Topology

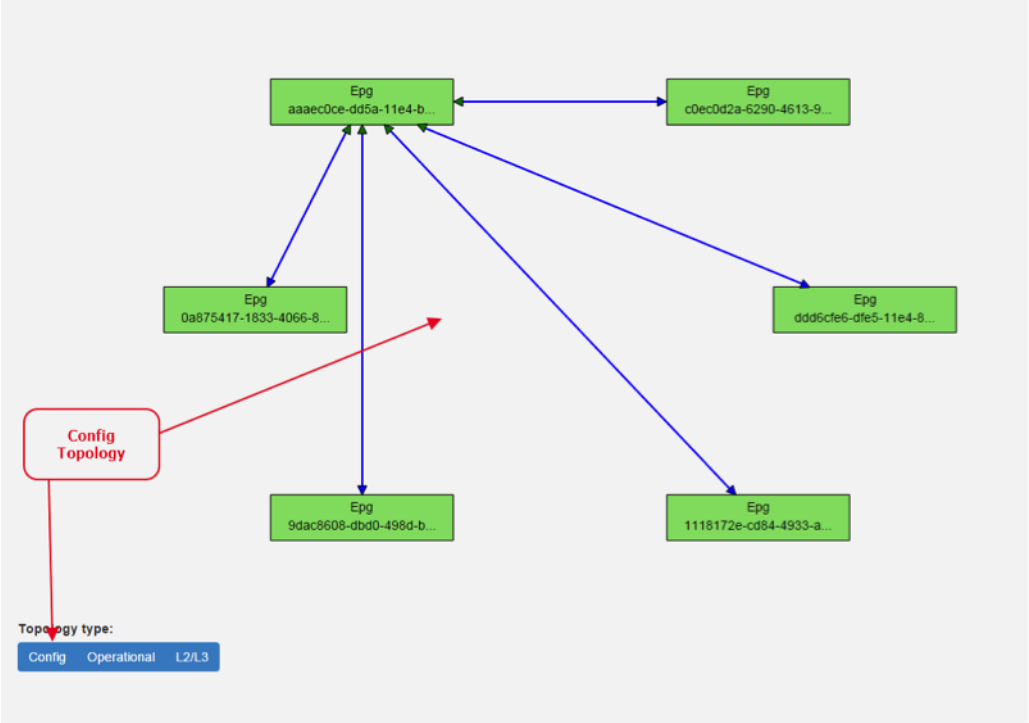

Fig. 47: Config Topology

For managing EndpointGroups (EPG) the tenant from the top Tenants list must be selected.

To add new EPG click Add button and after filling required attributes click Save button. After adding the EPG you can edit it and assign Consumer named selector or Provider named selector to it.

To edit EPG click the Edit button after selecting the EPG from Group list.

To add new Consumer named selector (CNS) click the Add button next to the Consumer named selectors list. While CNS editing you can set one or more contracts for current CNS pressing the Plus button and selecting the contract from the Contracts list. To remove the contract, click on the cross mark next to the contract. Added CNS can be viewed, edited or deleted by selecting from the Consumer named selectors list and clicking the Edit and Delete buttons like with the EPG or tenants.

To add new Provider named selector (PNS) click the Add button next to the Provider named selectors list. While PNS editing you can set one or more contracts for current PNS pressing the Plus button and selecting the contract from the Contracts list. To remove the contract, click on the cross mark next to the contract. Added PNS can be viewed, edited or deleted by selecting from the Provider named selectors list and clicking the Edit and Delete buttons like with the EPG or tenants.

To delete EPG, CNS or PNS select it in selectbox and click the Delete button next to the selectbox.

Contracts

For managing contracts the tenant from the top Tenants list must be selected.

To add new Contract click Add button and after filling required fields click Save button.

After adding the Contract user can edit it by selecting in the Contracts list and clicking Edit button.

To add new Clause click Add button next to the Clause list while editing the contract. While editing the Clause after selecting clause from the Clause list user can assign clause subjects by clicking the Plus button next to the Clause subjects label. Adding and editing action must be submitted by pressing Save button. To manage Subjects you can use CRUD form like with the Clause list.

L2/L3

For managing L2/L3 the tenant from the top Tenants list must be selected.

To add L3 Context click the Add button next to the L3 Context list ,which will display the form for adding a new L3 Context. After filling L3 Context attributes click Save button. After saving of L3 Context, form will be closed and the L3 Context list will be set to default value.

To view an existing L3 Context, select the L3 Context from the select box L3 Context list. The view form is read-only and can be closed by clicking cross mark in the top right of the form.

If user wants to edit selected L3 Context, click the Edit button, which will display the edit form for selected L3 Context. After editing click the Save button to save selected L3 Context. After saving of L3 Context, the edit form will be closed and the L3 Context list will be set to default value.

To delete L3 Context, select it from the L3 Context list and click Delete button.

To add L2 Bridge Domain, click the Add button next to the L2 Bridge Domain list. This will display the form for adding a new L2 Bridge Domain. After filling L2 Bridge Domain attributes click Save button. After saving of L2 Bridge Domain, form will be closed and the L2 Bridge Domain list will be set to default value.

To view an existing L2 Bridge Domain, select the L2 Bridge Domain from the select box L2 Bridge Domain list. The view form is read-only and can be closed by clicking cross mark in the top right of the form.

If user wants to edit selected L2 Bridge Domain, click the Edit button, which will display the edit form for selected L2 Bridge Domain. After editing click the Save button to save selected L2 Bridge Domain. After saving of L2 Bridge Domain the edit form will be closed and the L2 Bridge Domain list will be set to default value.

To delete L2 Bridge Domain select it from the L2 Bridge Domain list and click Delete button.

To add L3 Flood Domain, click the Add button next to the L3 Flood Domain list. This will display the form for adding a new L3 Flood Domain. After filling L3 Flood Domain attributes click Save button. After saving of L3 Flood Domain, form will be closed and the L3 Flood Domain list will be set to default value.

To view an existing L3 Flood Domain, select the L3 Flood Domain from the select box L3 Flood Domain list. The view form is read-only and can be closed by clicking cross mark in the top right of the form.

If user wants to edit selected L3 Flood Domain, click the Edit button, which will display the edit form for selected L3 Flood Domain. After editing click the Save button to save selected L3 Flood Domain. After saving of L3 Flood Domain the edit form will be closed and the L3 Flood Domain list will be set to default value.

To delete L3 Flood Domain select it from the L3 Flood Domain list and click Delete button.

To add Subnet click the Add button next to the Subnet list. This will display the form for adding a new Subnet. After filling Subnet attributes click Save button. After saving of Subnet, form will be closed and the Subnet list will be set to default value.

To view an existing Subnet, select the Subnet from the select box Subnet list. The view form is read-only and can be closed by clicking cross mark in the top right of the form.

If user wants to edit selected Subnet, click the Edit button, which will display the edit form for selected Subnet. After editing click the Save button to save selected Subnet. After saving of Subnet the edit form will be closed and the Subnet list will be set to default value.

To delete Subnet select it from the Subnet list and click Delete button.

Classifiers

To add Classifier, click the Add button next to the Classifier list. This will display the form for adding a new Classifier. After filling Classifier attributes click Save button. After saving of Classifier, form will be closed and the Classifier list will be set to default value.

To view an existing Classifier, select the Classifier from the select box Classifier list. The view form is read-only and can be closed by clicking cross mark in the top right of the form.

If you want to edit selected Classifier, click the Edit button, which will display the edit form for selected Classifier. After editing click the Save button to save selected Classifier. After saving of Classifier the edit form will be closed and the Classifier list will be set to default value.

To delete Classifier select it from the Classifier list and click Delete button.

Actions

To add Action, click the Add button next to the Action list. This will display the form for adding a new Action. After filling Action attributes click Save button. After saving of Action, form will be closed and the Action list will be set to default value.

To view an existing Action, select the Action from the select box Action list. The view form is read-only and can be closed by clicking cross mark in the top right of the form.

If user wants to edit selected Action, click the Edit button, which will display the edit form for selected Action. After editing click the Save button to save selected Action. After saving of Action the edit form will be closed and the Action list will be set to default value.

To delete Action select it from the Action list and click Delete button.

Endpoint

To add Endpoint, click the Add button next to the Endpoint list. This will display the form for adding a new Endpoint. To add EndpointGroup assignment click the Plus button next to the label EndpointGroups. To add Condition click Plus button next to the label Condition. To add L3 Address click the Plus button next to the L3 Addresses label. After filling Endpoint attributes click Save button. After saving of Endpoint, form will be closed and the Endpoint list will be set to default value.

To view an existing Endpoint just, the Endpoint from the select box Endpoint list. The view form is read-only and can be closed by clicking cross mark in the top right of the form.

If you want to edit selected Endpoint, click the Edit button, which will display the edit form for selected Endpoint. After editing click the Save button to save selected Endpoint. After saving of Endpoint the edit form will be closed and the Endpoint list will be set to default value.

To delete Endpoint select it from the Endpoint list and click Delete button.

L3 prefix endpoint

To add L3 prefix endpoint, click the Add button next to the L3 prefix endpoint list. This will display the form for adding a new Endpoint. To add EndpointGroup assignment, click the Plus button next to the label EndpointGroups. To add Condition, click Plus button next to the label Condition. To add L2 gateway click the Plus button next to the L2 gateways label. To add L3 gateway, click the Plus button next to the L3 gateways label. After filling L3 prefix endpoint attributes click Save button. After saving of L3 prefix endpoint, form will be closed and the Endpoint list will be set to default value.

To view an existing L3 prefix endpoint, select the Endpoint from the select box L3 prefix endpoint list. The view form is read-only and can be closed by clicking cross mark in the top right of the form.

If you want to edit selected L3 prefix endpoint, click the Edit button, which will display the edit form for selected L3 prefix endpoint. After editing click the Save button to save selected L3 prefix endpoint. After saving of Endpoint the edit form will be closed and the Endpoint list will be set to default value.

To delete Endpoint select it from the L3 prefix endpoint list and click Delete button.

Wizard

Wizard provides quick method to send basic data to controller necessary for basic usage of GBP application. It is useful in the case that there aren't any data in controller. In the first tab is form for create tenant. The second tab is for CRUD operations with contracts and their sub elements such as subjects, rules, clauses, action refs and classifier refs. The last tab is for CRUD operations with EndpointGroups and their CNS and PNS. Created structure of data is possible to send by clicking on Submit button.

Using the GBP API

Please see:

- *Using the GBP OpenFlow Overlay (OfOverlay) renderer*
- *Policy Resolution*
- *Forwarding Model*
- *the **GBP** demo and development environments for tips*

It is recommended to use either:

- *Neutron mapper <gbp-neutron>*
- *the UX*

If the REST API must be used, and the above resources are not sufficient:

- feature:install odl-dlux-yangui
- browse to: `http://<odl-controller>:8181/index.html` and select YangUI from the left menu.

to explore the various **GBP** REST options

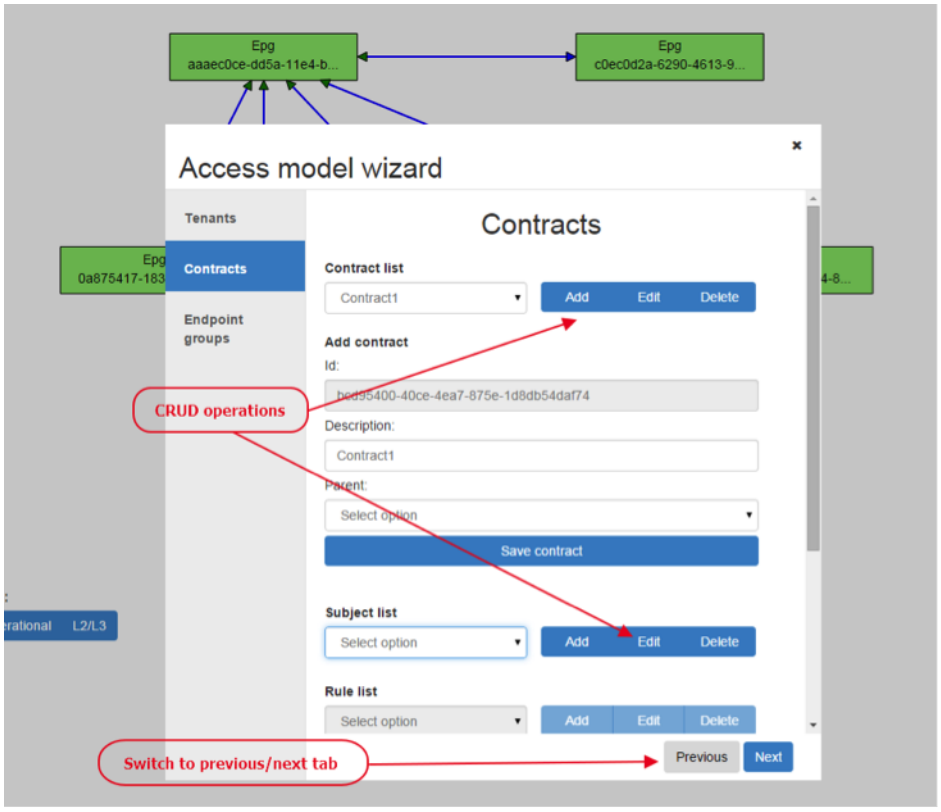

Fig. 48: Wizard

Using OpenStack with GBP

Overview

This section is for Application Developers and Network Administrators who are looking to integrate Group Based Policy with OpenStack.

To enable the **GBP** Neutron Mapper feature, at the Karaf console:

```
feature:install odl-groupbasedpolicy-neutronmapper
```

Neutron Mapper has the following dependencies that are automatically loaded:

```
odl-neutron-service
```

Neutron Northbound implementing REST API used by OpenStack

```
odl-groupbasedpolicy-base
```

Base **GBP** feature set, such as policy resolution, data model etc.

```
odl-groupbasedpolicy-ofoverlay
```

REST calls from OpenStack Neutron are by the Neutron NorthBound project.

GBP provides the implementation of the Neutron V2.0 API.

Features

List of supported Neutron entities:

- Port
- Network
  - Standard Internal
  - External provider L2/L3 network
- Subnet
- Security-groups
- Routers
  - Distributed functionality with local routing per compute
  - External gateway access per compute node (dedicated port required)
  - Multiple routers per tenant
- FloatingIP NAT
- IPv4/IPv6 support

The mapping of Neutron entities to **GBP** entities is as follows:

Neutron Port

The Neutron port is mapped to an endpoint.

The current implementation supports one IP address per Neutron port.

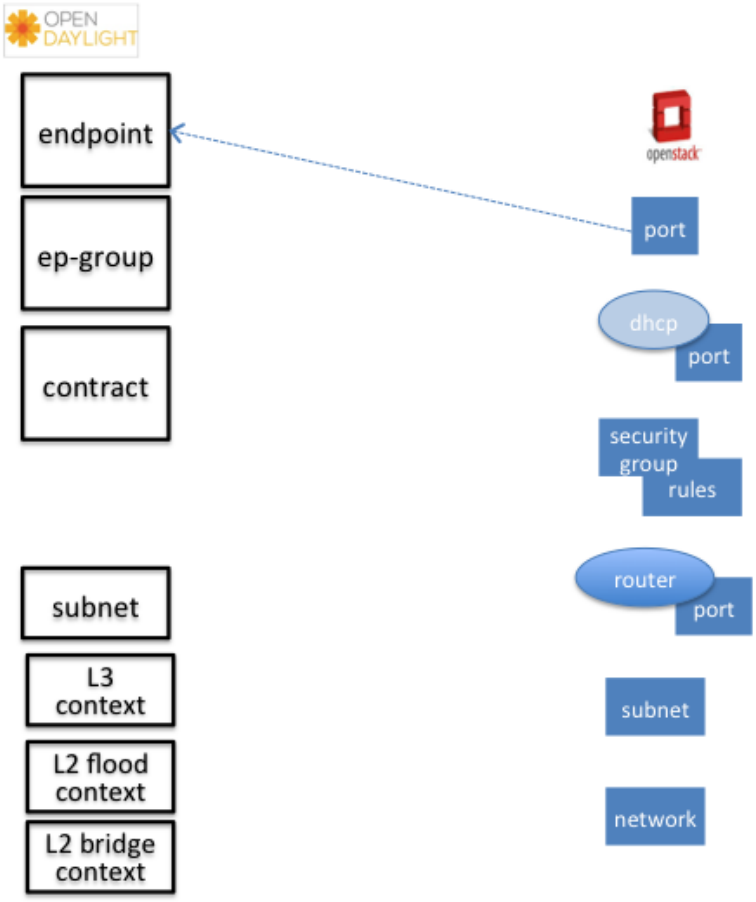

Fig. 49: Neutron Port

An endpoint and L3-endpoint belong to multiple EndpointGroups if the Neutron port is in multiple Neutron Security Groups.

The key for endpoint is L2-bridge-domain obtained as the parent of L2-flood-domain representing Neutron network. The MAC address is from the Neutron port. An L3-endpoint is created based on L3-context (the parent of the L2-bridge-domain) and IP address of Neutron Port.

Neutron Network

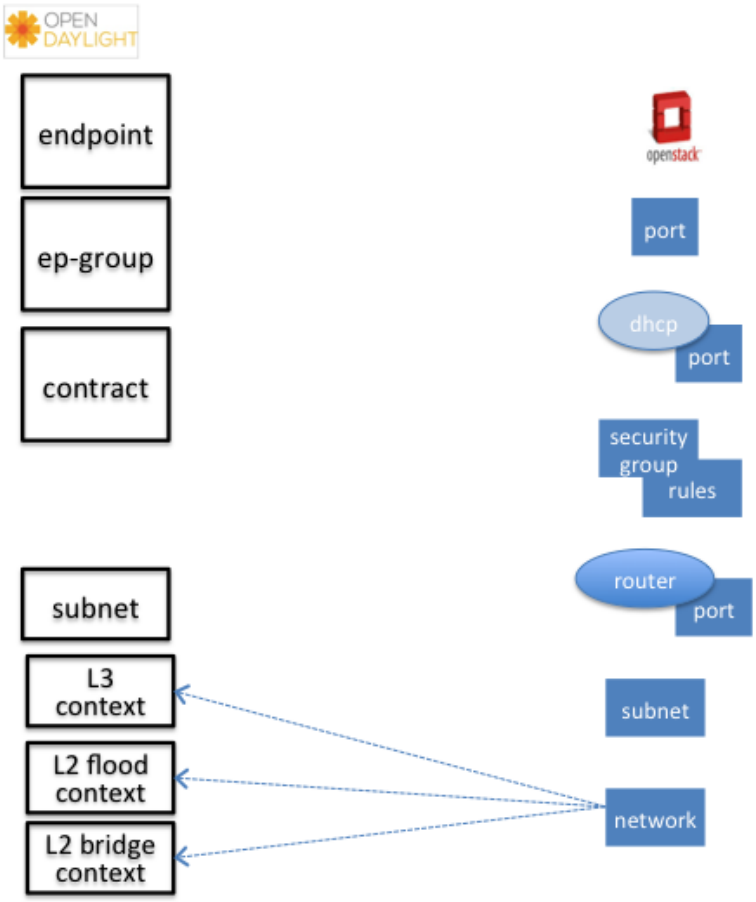

Fig. 50: Neutron Network

A Neutron network has the following characteristics:

- defines a broadcast domain
- defines a L2 transmission domain
- defines a L2 name space.

To represent this, a Neutron Network is mapped to multiple **GBP** entities. The first mapping is to an L2 flood-domain to reflect that the Neutron network is one flooding or broadcast domain. An L2-bridge-domain is then associated as the parent of L2 flood-domain. This reflects both the L2 transmission domain as well as the L2 addressing namespace.

The third mapping is to L3-context, which represents the distinct L3 address space. The L3-context is the parent of L2-bridge-domain.

Neutron Subnet

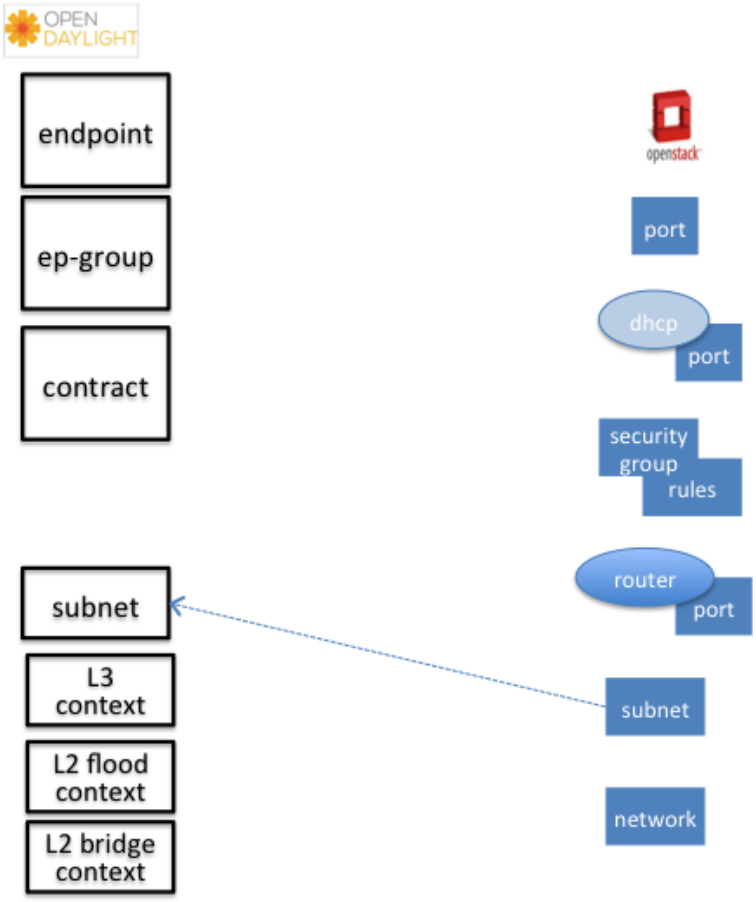

Fig. 51: Neutron Subnet

Neutron subnet is associated with a Neutron network. The Neutron subnet is mapped to a **GBP** subnet where the parent of the subnet is L2-flood-domain representing the Neutron network.

Neutron Security Group

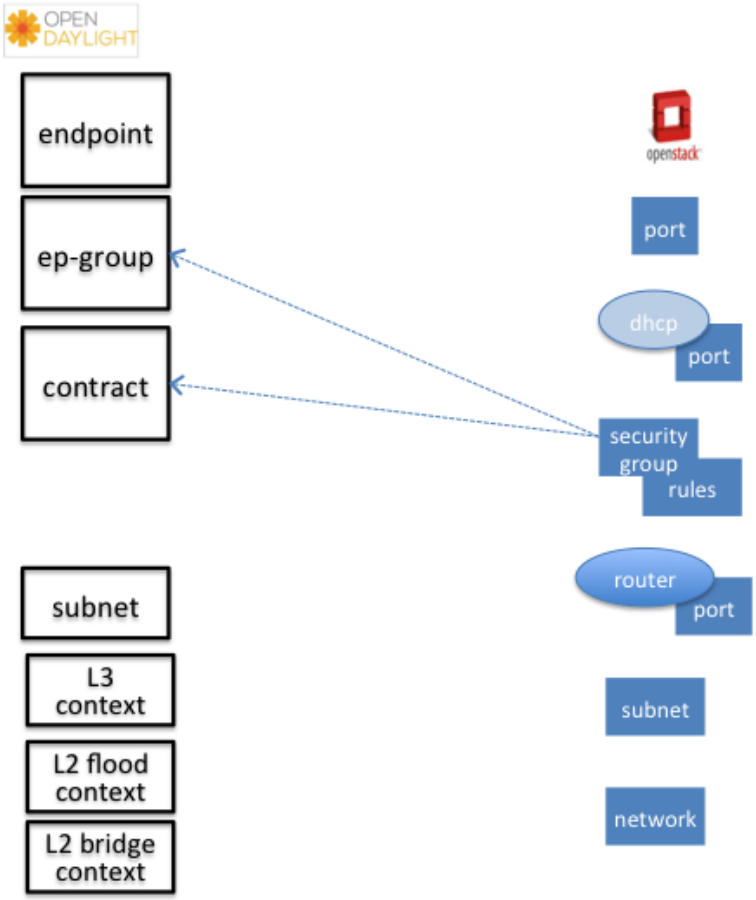

Fig. 52: Neutron Security Group and Rules

GBP entity representing Neutron security-group is EndpointGroup.

Infrastructure EndpointGroups

Neutron-mapper automatically creates EndpointGroups to manage key infrastructure items such as:

- DHCP EndpointGroup - contains endpoints representing Neutron DHCP ports
- Router EndpointGroup - contains endpoints representing Neutron router interfaces
- External EndpointGroup - holds L3-endpoints representing Neutron router gateway ports, also associated with FloatingIP ports.

Neutron Security Group Rules

This is the most involved amongst all the mappings because Neutron security-group-rules are mapped to contracts with clauses, subjects, rules, action-refs, classifier-refs, etc. Contracts are used between EndpointGroups representing Neutron Security Groups. For simplification it is important to note that Neutron security-group-rules are similar to a **GBP** rule containing:

- classifier with direction

- action of **allow**.

Neutron Routers

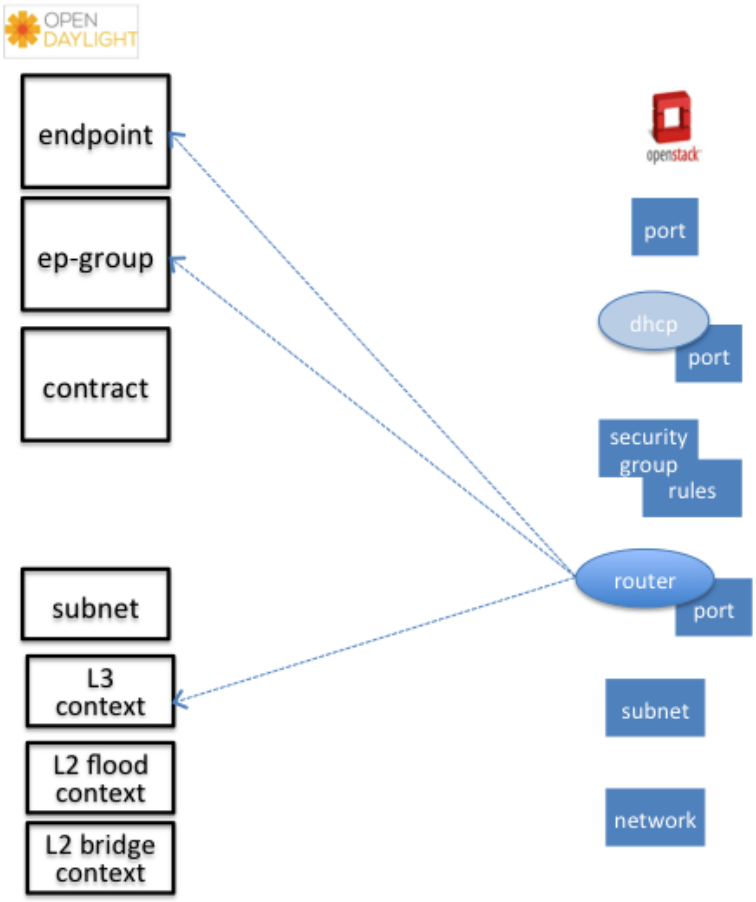

Fig. 53: Neutron Router

Neutron router is represented as a L3-context. This treats a router as a Layer3 namespace, and hence every network attached to it a part of that Layer3 namespace.

This allows for multiple routers per tenant with complete isolation.

The mapping of the router to an endpoint represents the router's interface or gateway port.

The mapping to an EndpointGroup represents the internal infrastructure EndpointGroups created by the **GBP** Neutron Mapper

When a Neutron router interface is attached to a network/subnet, that network/subnet and its associated endpoints or Neutron Ports are seamlessly added to the namespace.

Neutron FloatingIP

When associated with a Neutron Port, this leverages the *OfOverlay* renderer's NAT capabilities.

A dedicated *external* interface on each Nova compute host allows for disitributed external access. Each Nova instance associated with a FloatingIP address can access the external network directly without having to route via the Neutron controller, or having to enable any form of Neutron distributed routing functionality.

Assuming the gateway provisioned in the Neutron Subnet command for the external network is reachable, the combination of **GBP** Neutron Mapper and *OfOverlay renderer* will automatically ARP for this default gateway, requiring no user intervention.

Troubleshooting within GBP

Logging level for the mapping functionality can be set for package org.opendaylight.groupbasedpolicy.neutron.mapper. An example of enabling TRACE logging level on Karaf console:

```
log:set TRACE org.opendaylight.groupbasedpolicy.neutron.mapper
```

Neutron mapping example

As an example for mapping can be used creation of Neutron network, subnet and port. When a Neutron network is created 3 **GBP** entities are created: l2-flood-domain, l2-bridge-domain, l3-context.

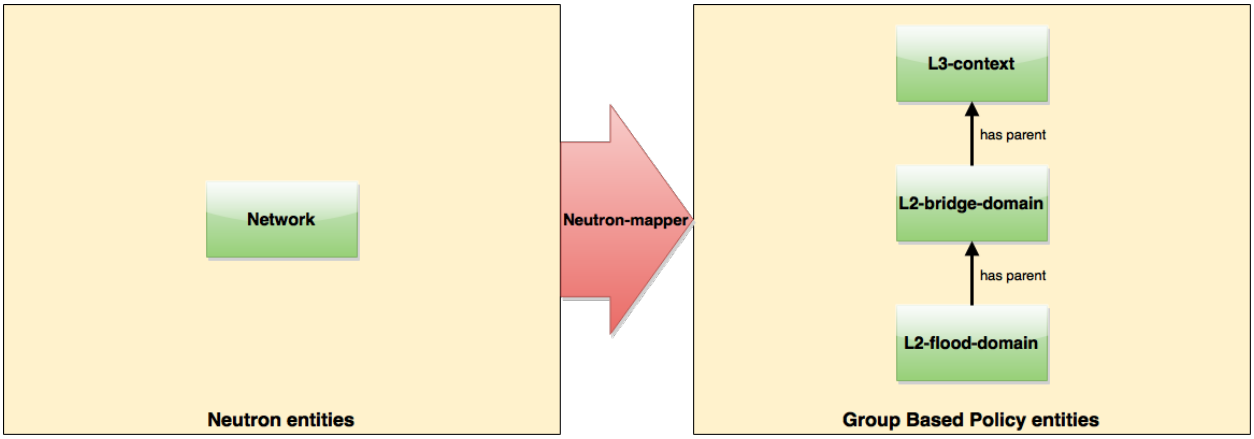

Fig. 54: Neutron network mapping

After an subnet is created in the network mapping looks like this.

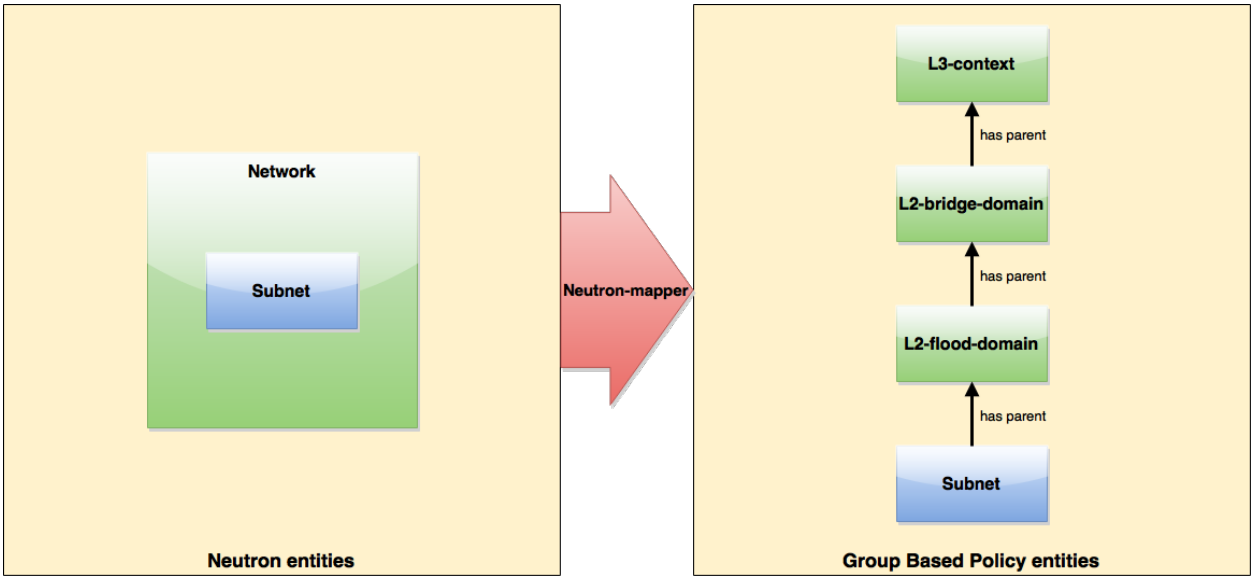

Fig. 55: Neutron subnet mapping

If an Neutron port is created in the subnet an endpoint and l3-endpoint are created. The endpoint has key composed from l2-bridge-domain and MAC address from Neutron port. A key of l3-endpoint is compesed from l3-context and IP address. The network containment of endpoint and l3-endpoint points to the subnet.

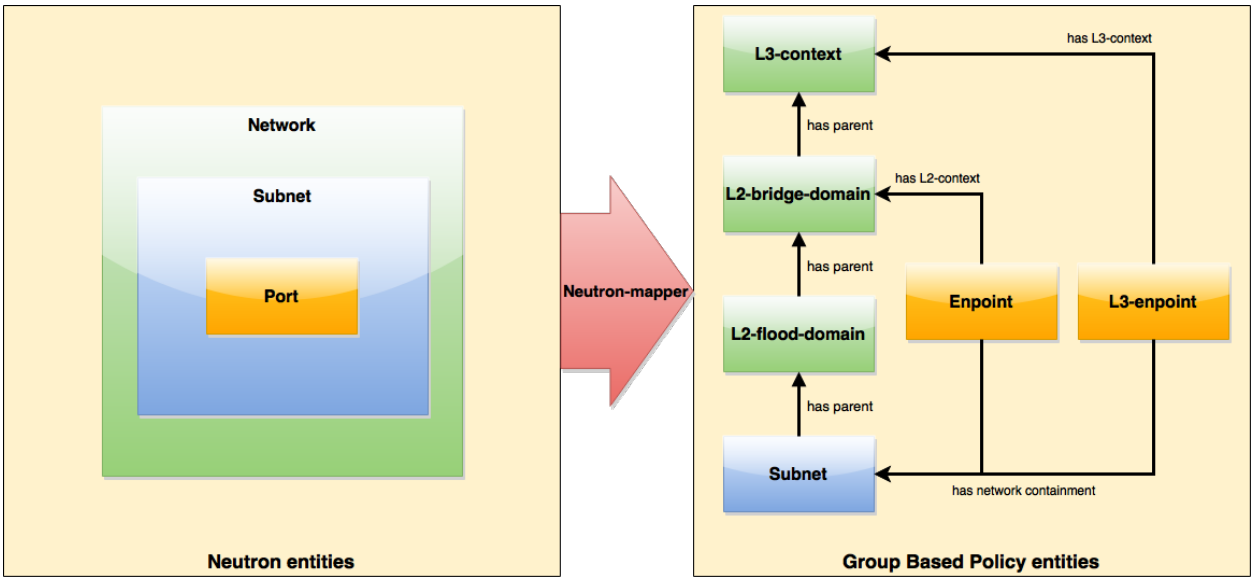

Fig. 56: Neutron port mapping

Configuring GBP Neutron

No intervention passed initial OpenStack setup is required by the user.

More information about configuration can be found in our DevStack demo environment on the GBP wiki.

Administering or Managing GBP Neutron

For consistencies sake, all provisioning should be performed via the Neutron API. (CLI or Horizon).

The mapped policies can be augmented via the **GBP** *UX*, to:

- Enable *Service Function Chaining*
- Add endpoints from outside of Neutron i.e. VMs/containers not provisioned in OpenStack
- Augment policies/contracts derived from Security Group Rules
- Overlay additional contracts or groupings

Tutorials

A DevStack demo environment can be found on the GBP wiki.

GBP Renderer manager

Overview

The GBP Renderer manager is an integral part of **GBP** base module. It dispatches information about endpoints' policy configuration to specific device renderer by writing a renderer policy configuration into the registered renderer's policy store.

Installing and Pre-requisites

Renderer manager is integrated into GBP base module, so no additional installation is required.

Architecture

Renderer manager gets data notifications about:

- Endoints (base-endpoint.yang)
- EndpointLocations (base-endpoint.yang)
- ResolvedPolicies (resolved-policy.yang)
- Forwarding (forwarding.yang)

Based on data from notifications it creates a configuration task for specific renderers by writing a renderer policy configuration into the registered renderer's policy store. Configuration is stored to CONF data store as Renderers (renderer.yang).

Configuration is signed with version number which is incremented by every change. All renderers are supposed to be on the same version. Renderer manager waits for all renderers to respond with version update in OPER data store. After a version of every renderer in OPER data store has the same value as the one in CONF data store, renderer manager moves to the next configuration with incremented version.

GBP Location manager

Overview

Location manager monitors information about Endpoint Location providers (see endpoint-location-provider.yang) and manages Endpoint locations in OPER data store accordingly.

Installing and Pre-requisites

Location manager is integrated into GBP base module, so no additional installation is required.

Architecture

The endpoint-locations container in OPER data store (see base-endpoint.yang) contains two lists for two types of EP location, namely address-endpoint-location and containment-endpoint-location. LocationResolver is a class that processes Location providers in CONF data store and puts location information to OPER data store.

When a new Location provider is created in CONF data store, its Address EP locations are being processed first, and their info is stored locally in accordance with processed Location provider's priority. Then a location of type

"absolute" with the highest priority is selected for an EP, and is put in OPER data store. If Address EP locations contain locations of type "relative", those are put to OPER data store.

If current Location provider contains Containment EP locations of type "relative", then those are put to OPER data store.

Similarly, when a Location provider is deleted, information of its locations is removed from the OPER data store.

Using the GBP OpenFlow Overlay (OfOverlay) renderer

Overview

The OpenFlow Overlay (OfOverlay) feature enables the OpenFlow Overlay renderer, which creates a network virtualization solution across nodes that host Open vSwitch software switches.

Installing and Pre-requisites

From the Karaf console in OpenDaylight:

```
feature:install odl-groupbasedpolicy-ofoverlay
```

This renderer is designed to work with OpenVSwitch (OVS) 2.1+ (although 2.3 is strongly recommended) and OpenFlow 1.3.

When used in conjunction with the *Neutron Mapper feature* no extra OfOverlay specific setup is required.

When this feature is loaded "standalone", the user is required to configure infrastructure, such as

- instantiating OVS bridges,
- attaching hosts to the bridges,
- and creating the VXLAN/VXLAN-GPE tunnel ports on the bridges.

The **GBP** OfOverlay renderer also supports a table offset option, to offset the pipeline post-table 0. The value of table offset is stored in the config datastore and it may be rewritten at runtime.

```
PUT http://{{controllerIp}}:8181/restconf/config/ofoverlay:of-overlay-config
{
    "of-overlay-config": {
        "gbp-ofoverlay-table-offset": 6
    }
}
```

The default value is set by changing: <gbp-ofoverlay-table-offset>0</gbp-ofoverlay-table-offset>

in file: distribution-karaf/target/assembly/etc/opendaylight/karaf/15-groupbasedpolicy-ofoverlay.xml

To avoid overwriting runtime changes, the default value is used only when the OfOverlay renderer starts and no other value has been written before.

OpenFlow Overlay Architecture

These are the primary components of **GBP**. The OfOverlay components are highlighted in red.

In terms of the inner components of the **GBP** OfOverlay renderer:

OfOverlay Renderer

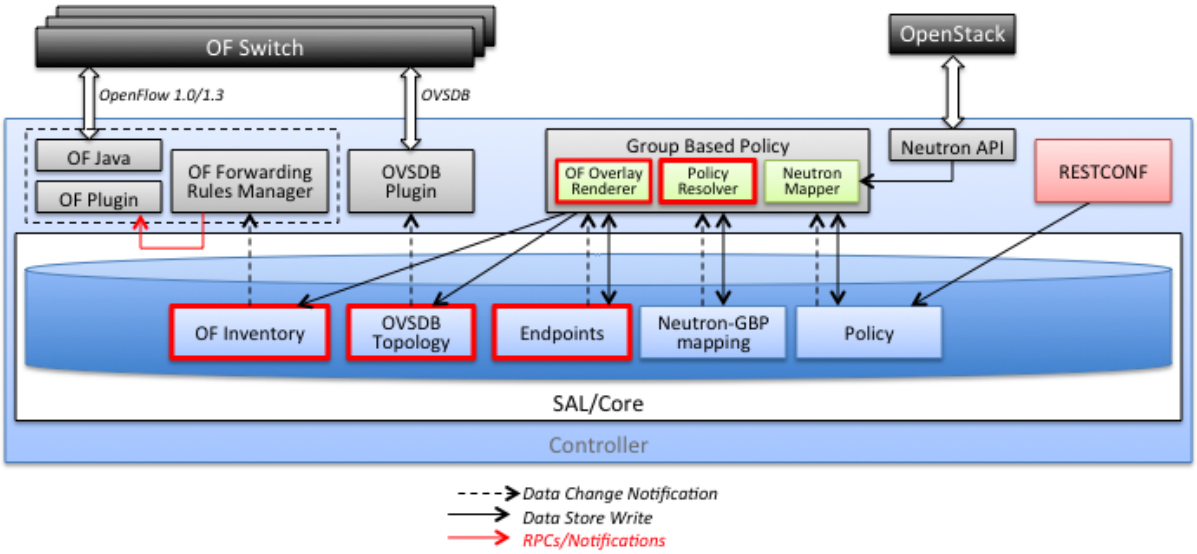

Fig. 57: OfOverlay within **GBP**

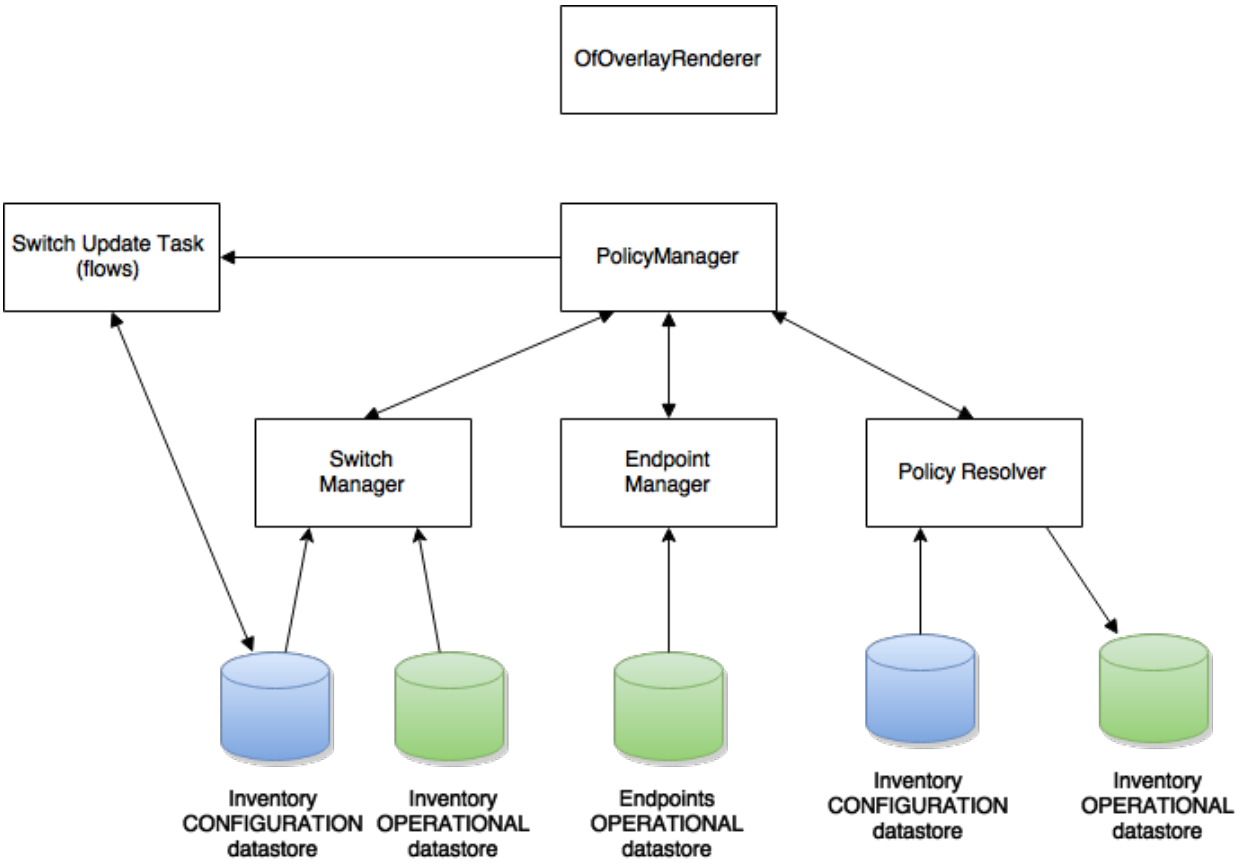

Fig. 58: OfOverlay expanded view:

Launches components below:

Policy Resolver

Policy resolution is completely domain independent, and the OfOverlay leverages process policy information internally. See Policy Resolution process.

It listens to inputs to the *Tenants* configuration datastore, validates tenant input, then writes this to the Tenants operational datastore.

From there an internal notification is generated to the PolicyManager.

In the next release, this will be moving to a non-renderer specific location.

Endpoint Manager

The endpoint repository operates in **orchestrated** mode. This means the user is responsible for the provisioning of endpoints via:

- *UX/GUI*
- REST API

Note: When using the *Neutron mapper* feature, everything is managed transparently via Neutron.

The Endpoint Manager is responsible for listening to Endpoint repository updates and notifying the Switch Manager when a valid Endpoint has been registered.

It also supplies utility functions to the flow pipeline process.

Switch Manager

The Switch Manager is purely a state manager.

Switches are in one of 3 states:

- DISCONNECTED
- PREPARING
- READY

Ready is denoted by a connected switch:

- having a tunnel interface
- having at least one endpoint connected.

In this way **GBP** is not writing to switches it has no business to.

Preparing simply means the switch has a controller connection but is missing one of the above *complete and necessary* conditions

Disconnected means a previously connected switch is no longer present in the Inventory operational datastore.

The OfOverlay leverages Nicira registers as follows:

- REG0 = Source EndpointGroup + Tenant ordinal
- REG1 = Source Conditions + Tenant ordinal
- REG2 = Destination EndpointGroup + Tenant ordinal
- REG3 = Destination Conditions + Tenant ordinal
- REG4 = Bridge Domain + Tenant ordinal

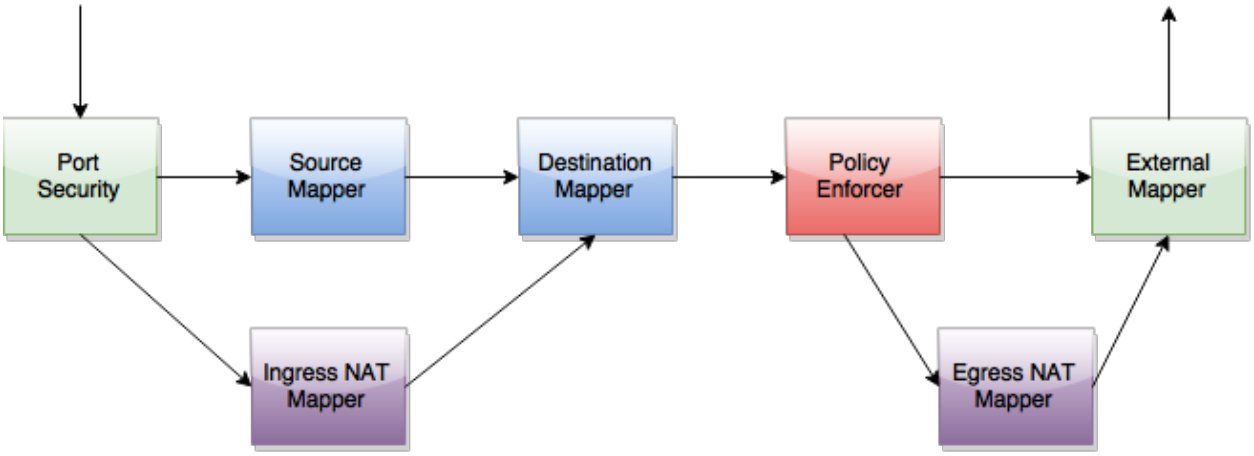

Fig. 59: OfOverlay Flow Pipeline

- REG5 = Flood Domain + Tenant ordinal
- REG6 = Layer 3 Context + Tenant ordinal

Port Security

Table 0 of the OpenFlow pipeline. Responsible for ensuring that only valid connections can send packets into the pipeline:

```
cookie=0x0, <snip> , priority=200,in_port=3 actions=goto_table:2
cookie=0x0, <snip> , priority=200,in_port=1 actions=goto_table:1
cookie=0x0, <snip> , priority=121,arp,in_port=5,dl_src=fa:16:3e:d5:b9:8d,arp_spa=10.1.
↪1.3 actions=goto_table:2
cookie=0x0, <snip> , priority=120,ip,in_port=5,dl_src=fa:16:3e:d5:b9:8d,nw_src=10.1.1.
↪3 actions=goto_table:2
cookie=0x0, <snip> , priority=115,ip,in_port=5,dl_src=fa:16:3e:d5:b9:8d,nw_dst=255.
↪255.255.255 actions=goto_table:2
cookie=0x0, <snip> , priority=112,ipv6 actions=drop
cookie=0x0, <snip> , priority=111, ip actions=drop
cookie=0x0, <snip> , priority=110,arp actions=drop
cookie=0x0, <snip> ,in_port=5,dl_src=fa:16:3e:d5:b9:8d actions=goto_table:2
cookie=0x0, <snip> , priority=1 actions=drop
```

Ingress from tunnel interface, go to Table *Source Mapper*:

```
cookie=0x0, <snip> , priority=200,in_port=3 actions=goto_table:2
```

Ingress from outside, goto Table *Ingress NAT Mapper*:

```
cookie=0x0, <snip> , priority=200,in_port=1 actions=goto_table:1
```

ARP from Endpoint, go to Table *Source Mapper*:

```
cookie=0x0, <snip> , priority=121,arp,in_port=5,dl_src=fa:16:3e:d5:b9:8d,arp_spa=10.1.
↪1.3 actions=goto_table:2
```

IPv4 from Endpoint, go to Table *Source Mapper*:

```
cookie=0x0, <snip> , priority=120,ip,in_port=5,dl_src=fa:16:3e:d5:b9:8d,nw_src=10.1.1.
↪3 actions=goto_table:2
```

DHCP DORA from Endpoint, go to Table *Source Mapper*:

```
cookie=0x0, <snip> , priority=115,ip,in_port=5,dl_src=fa:16:3e:d5:b9:8d,nw_dst=255.
↪255.255.255 actions=goto_table:2
```

Series of DROP tables with priority set to capture any non-specific traffic that should have matched above:

```
cookie=0x0, <snip> , priority=112,ipv6 actions=drop
cookie=0x0, <snip> , priority=111, ip actions=drop
cookie=0x0, <snip> , priority=110,arp actions=drop
```

"L2" catch all traffic not identified above:

```
cookie=0x0, <snip> ,in_port=5,dl_src=fa:16:3e:d5:b9:8d actions=goto_table:2
```

Drop Flow:

```
cookie=0x0, <snip> , priority=1 actions=drop
```

Ingress NAT Mapper

Table *offset* +1.

ARP responder for external NAT address:

```
cookie=0x0, <snip> , priority=150,arp,arp_tpa=192.168.111.51,arp_op=1␣
↪actions=move:NXM_OF_ETH_SRC[]->NXM_OF_ETH_DST[],set_field:fa:16:3e:58:c3:dd->eth_
↪src,load:0x2->NXM_OF_ARP_OP[],move:NXM_NX_ARP_SHA[]->NXM_NX_ARP_THA[],
↪load:0xfa163e58c3dd->NXM_NX_ARP_SHA[],move:NXM_OF_ARP_SPA[]->NXM_OF_ARP_TPA[],
↪load:0xc0a86f33->NXM_OF_ARP_SPA[],IN_PORT
```

Translate from Outside to Inside and perform same functions as SourceMapper.

```
cookie=0x0, <snip> , priority=100,ip,nw_dst=192.168.111.51 actions=set_field:10.1.1.2-
↪>ip_dst,set_field:fa:16:3e:58:c3:dd->eth_dst,load:0x2->NXM_NX_REG0[],load:0x1->NXM_
↪NX_REG1[],load:0x4->NXM_NX_REG4[],load:0x5->NXM_NX_REG5[],load:0x7->NXM_NX_REG6[],
↪load:0x3->NXM_NX_TUN_ID[0..31],goto_table:3
```

Source Mapper

Table *offset* +2.

Determines based on characteristics from the ingress port, which:

- EndpointGroup(s) it belongs to
- Forwarding context
- Tunnel VNID ordinal

Establishes tunnels at valid destination switches for ingress.

Ingress Tunnel established at remote node with VNID Ordinal that maps to Source EPG, Forwarding Context etc:

```
cookie=0x0, <snip>, priority=150,tun_id=0xd,in_port=3 actions=load:0xc->NXM_NX_REG0[],
↪load:0xffffff->NXM_NX_REG1[],load:0x4->NXM_NX_REG4[],load:0x5->NXM_NX_REG5[],
↪load:0x7->NXM_NX_REG6[],goto_table:3
```

Maps endpoint to Source EPG, Forwarding Context based on ingress port, and MAC:

```
cookie=0x0, <snip> , priority=100,in_port=5,dl_src=fa:16:3e:b4:b4:b1 actions=load:0xc-
↪>NXM_NX_REG0[],load:0x1->NXM_NX_REG1[],load:0x4->NXM_NX_REG4[],load:0x5->NXM_NX_
↪REG5[],load:0x7->NXM_NX_REG6[],load:0xd->NXM_NX_TUN_ID[0..31],goto_table:3
```

Generic drop:

```
cookie=0x0, duration=197.622s, table=2, n_packets=0, n_bytes=0, priority=1␣
↪actions=drop
```

Destination Mapper

Table *offset* +3.

Determines based on characteristics of the endpoint:

- EndpointGroup(s) it belongs to
- Forwarding context
- Tunnel Destination value

Manages routing based on valid ingress nodes ARP'ing for their default gateway, and matches on either gateway MAC or destination endpoint MAC.

ARP for default gateway for the 10.1.1.0/24 subnet:

```
cookie=0x0, <snip> , priority=150,arp,reg6=0x7,arp_tpa=10.1.1.1,arp_op=1␣
↪actions=move:NXM_OF_ETH_SRC[]->NXM_OF_ETH_DST[],set_field:fa:16:3e:28:4c:82->eth_
↪src,load:0x2->NXM_OF_ARP_OP[],move:NXM_NX_ARP_SHA[]->NXM_NX_ARP_THA[],
↪load:0xfa163e284c82->NXM_NX_ARP_SHA[],move:NXM_OF_ARP_SPA[]->NXM_OF_ARP_TPA[],
↪load:0xa010101->NXM_OF_ARP_SPA[],IN_PORT
```

Broadcast traffic destined for GroupTable:

```
cookie=0x0, <snip> , priority=140,reg5=0x5,dl_dst=01:00:00:00:00:00/01:00:00:00:00:00␣
↪actions=load:0x5->NXM_NX_TUN_ID[0..31],group:5
```

Layer3 destination matching flows, where priority=100+masklength. Since **GBP** now support L3Prefix endpoint, we can set default routes etc:

```
cookie=0x0, <snip>, priority=132,ip,reg6=0x7,dl_dst=fa:16:3e:b4:b4:b1,nw_dst=10.1.1.3␣
↪actions=load:0xc->NXM_NX_REG2[],load:0x1->NXM_NX_REG3[],load:0x5->NXM_NX_REG7[],set_
↪field:fa:16:3e:b4:b4:b1->eth_dst,dec_ttl,goto_table:4
```

Layer2 destination matching flows, designed to be caught only after last IP flow (lowest priority IP flow is 100):

```
cookie=0x0, duration=323.203s, table=3, n_packets=4, n_bytes=168, priority=50,
↪reg4=0x4,dl_dst=fa:16:3e:58:c3:dd actions=load:0x2->NXM_NX_REG2[],load:0x1->NXM_NX_
↪REG3[],load:0x2->NXM_NX_REG7[],goto_table:4
```

General drop flow: cookie=0x0, duration=323.207s, table=3, n_packets=6, n_bytes=588, priority=1 actions=drop

Policy Enforcer

Table *offset* +4.

Once the Source and Destination EndpointGroups are assigned, policy is enforced based on resolved rules.

In the case of *Service Function Chaining*, the encapsulation and destination for traffic destined to a chain, is discovered and enforced.

Policy flow, allowing IP traffic between EndpointGroups:

```
cookie=0x0, <snip> , priority=64998,ip,reg0=0x8,reg1=0x1,reg2=0xc,reg3=0x1␣
↪actions=goto_table:5
```

Egress NAT Mapper

Table *offset* +5.

Performs NAT function before Egressing OVS instance to the underlay network.

Inside to Outside NAT translation before sending to underlay:

```
cookie=0x0, <snip> , priority=100,ip,reg6=0x7,nw_src=10.1.1.2 actions=set_field:192.
↪168.111.51->ip_src,goto_table:6
```

External Mapper

Table *offset* +6.

Manages post-policy enforcement for endpoint specific destination effects. Specifically for *Service Function Chaining*, which is why we can support both symmetric and asymmetric chains and distributed ingress/egress classification.

Generic allow:

```
cookie=0x0, <snip>, priority=100 actions=output:NXM_NX_REG7[]
```

Configuring OpenFlow Overlay via REST

Note: Please see the *UX* section on how to configure **GBP** via the GUI.

Endpoint

```
POST http://{{controllerIp}}:8181/restconf/operations/endpoint:register-endpoint
{
    "input": {
        "endpoint-group": "<epg0>",
        "endpoint-groups" : ["<epg1>","<epg2>"],
        "network-containment" : "<fowarding-model-context1>",
        "l2-context": "<bridge-domain1>",
        "mac-address": "<mac1>",
        "l3-address": [
            {
                "ip-address": "<ipaddress1>",
                "l3-context": "<l3_context1>"
            }
        ],
        "*ofoverlay:port-name*": "<ovs port name>",
        "tenant": "<tenant1>"
    }
}
```

Note: The usage of "port-name" preceded by "ofoverlay". In OpenDaylight, base datastore objects can be *augmented*. In **GBP**, the base endpoint model has no renderer specifics, hence can be leveraged across multiple renderers.

OVS Augmentations to Inventory

```
PUT http://{{controllerIp}}:8181/restconf/config/opendaylight-inventory:nodes/
{
    "opendaylight-inventory:nodes": {
        "node": [
            {
                "id": "openflow:123456",
                "ofoverlay:tunnel": [
                    {
                        "tunnel-type": "overlay:tunnel-type-vxlan",
                        "ip": "<ip_address_of_ovs>",
                        "port": 4789,
                        "node-connector-id": "openflow:123456:1"
                    }
                ]
            },
            {
                "id": "openflow:654321",
                "ofoverlay:tunnel": [
                    {
                        "tunnel-type": "overlay:tunnel-type-vxlan",
                        "ip": "<ip_address_of_ovs>",
                        "port": 4789,
                        "node-connector-id": "openflow:654321:1"
                    }
                ]
            }
        ]
    }
}
```

Tenants see *Policy Resolution* and *Forwarding Model* for details:

```
{
  "policy:tenant": {
    "contract": [
      {
        "clause": [
          {
            "name": "allow-http-clause",
            "subject-refs": [
              "allow-http-subject",
              "allow-icmp-subject"
            ]
          }
        ],
        "id": "<id>",
        "subject": [
          {
            "name": "allow-http-subject",
            "rule": [
              {
                "classifier-ref": [
                  {
                    "direction": "in",
                    "name": "http-dest"
                  },
                  {
```

(continues on next page)

(continued from previous page)

```
                    "direction": "out",
                    "name": "http-src"
                  }
                ],
                "action-ref": [
                  {
                    "name": "allow1",
                    "order": 0
                  }
                ],
                "name": "allow-http-rule"
              }
            ]
          },
          {
            "name": "allow-icmp-subject",
            "rule": [
              {
                "classifier-ref": [
                  {
                    "name": "icmp"
                  }
                ],
                "action-ref": [
                  {
                    "name": "allow1",
                    "order": 0
                  }
                ],
                "name": "allow-icmp-rule"
              }
            ]
          }
        ]
      }
    ],
    "endpoint-group": [
      {
        "consumer-named-selector": [
          {
            "contract": [
              "<id>"
            ],
            "name": "<name>"
          }
        ],
        "id": "<id>",
        "provider-named-selector": []
      },
      {
        "consumer-named-selector": [],
        "id": "<id>",
        "provider-named-selector": [
          {
            "contract": [
              "<id>"
            ],
```

(continues on next page)

(continued from previous page)

```
            "name": "<name>"
          }
        ]
      }
    ],
    "id": "<id>",
    "l2-bridge-domain": [
      {
        "id": "<id>",
        "parent": "<id>"
      }
    ],
    "l2-flood-domain": [
      {
        "id": "<id>",
        "parent": "<id>"
      },
      {
        "id": "<id>",
        "parent": "<id>"
      }
    ],
    "l3-context": [
      {
        "id": "<id>"
      }
    ],
    "name": "GBPPOC",
    "subject-feature-instances": {
      "classifier-instance": [
        {
          "classifier-definition-id": "<id>",
          "name": "http-dest",
          "parameter-value": [
            {
              "int-value": "6",
              "name": "proto"
            },
            {
              "int-value": "80",
              "name": "destport"
            }
          ]
        },
        {
          "classifier-definition-id": "<id>",
          "name": "http-src",
          "parameter-value": [
            {
              "int-value": "6",
              "name": "proto"
            },
            {
              "int-value": "80",
              "name": "sourceport"
            }
          ]
```

(continues on next page)

(continued from previous page)

```
        },
        {
          "classifier-definition-id": "<id>",
          "name": "icmp",
          "parameter-value": [
            {
              "int-value": "1",
              "name": "proto"
            }
          ]
        }
      ],
      "action-instance": [
        {
          "name": "allow1",
          "action-definition-id": "<id>"
        }
      ]
    },
    "subnet": [
      {
        "id": "<id>",
        "ip-prefix": "<ip_prefix>",
        "parent": "<id>",
        "virtual-router-ip": "<ip address>"
      },
      {
        "id": "<id>",
        "ip-prefix": "<ip prefix>",
        "parent": "<id>",
        "virtual-router-ip": "<ip address>"
      }
    ]
  }
}
```

Tutorials

Comprehensive tutorials, along with a demonstration environment leveraging Vagrant can be found on the GBP wiki

Using the GBP eBPF IO Visor Agent renderer

Overview

The IO Visor renderer feature enables container endpoints (e.g. Docker, LXC) to leverage GBP policies.

The renderer interacts with a IO Visor module from the Linux Foundation IO Visor project.

Installing and Pre-requisites

From the Karaf console in OpenDaylight:

```
feature:install odl-groupbasedpolicy-iovisor odl-restconf
```

Installation details, usage, and other information for the IO Visor GBP module can be found here: IO Visor github repo for IO Modules

Using the GBP FaaS renderer

Overview

The FaaS renderer feature enables leveraging the FaaS project as a GBP renderer.

Installing and Pre-requisites

From the Karaf console in OpenDaylight:

```
feature:install odl-groupbasedpolicy-faas
```

More information about FaaS can be found here: https://wiki.opendaylight.org/view/FaaS:GBPIntegration

Using Service Function Chaining (SFC) with GBP Neutron Mapper and OfOverlay

Overview

Please refer to the Service Function Chaining project for specifics on SFC provisioning and theory.

GBP allows for the use of a chain, by name, in policy.

This takes the form of an *action* in **GBP**.

Using the *GBP demo and development environment* as an example:

In the topology above, a symmetrical chain between H35_2 and H36_3 could take path:

H35_2 to sw1 to sff1 to sf1 to sff1 to sff2 to sf2 to sff2 to sw6 to H36_3

If symmetric chaining was desired, the return path is:

If asymmetric chaining was desired, the return path could be direct, or an **entirely different chain**.

All these scenarios are supported by the integration.

In the **Subject Feature Instance** section of the tenant config, we define the instances of the classifier definitions for ICMP and HTTP:

```
"subject-feature-instances": {
  "classifier-instance": [
    {
      "name": "icmp",
      "parameter-value": [
        {
          "name": "proto",
          "int-value": 1
        }
      ]
    },
```

(continues on next page)

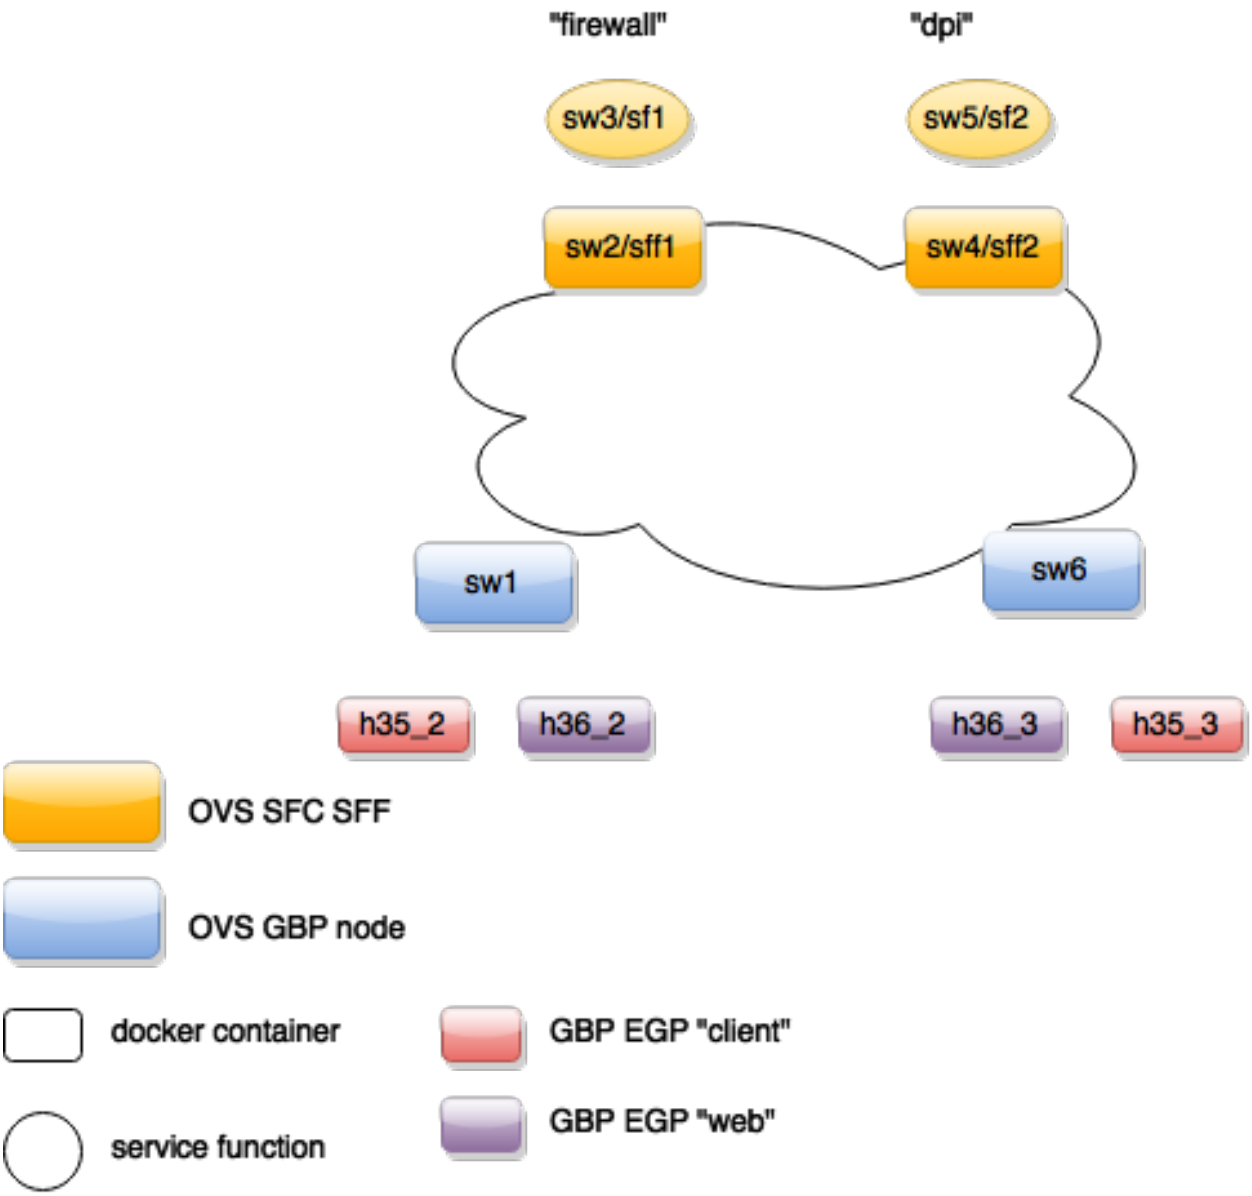

Fig. 60: GBP and SFC integration environment

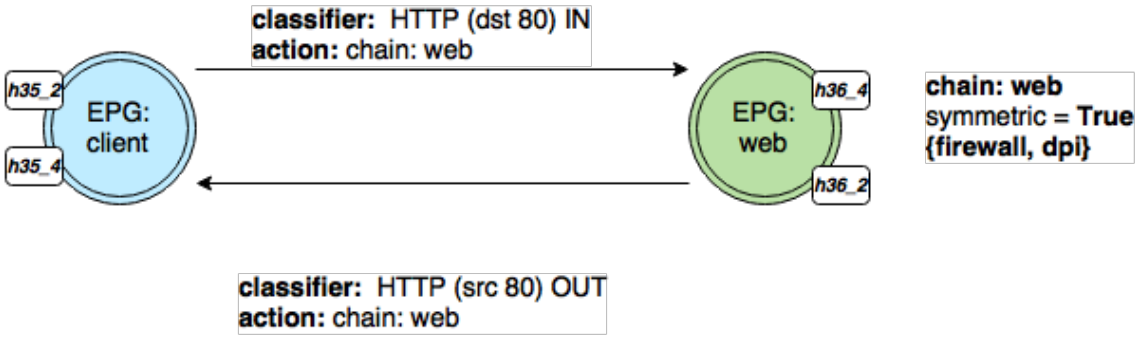

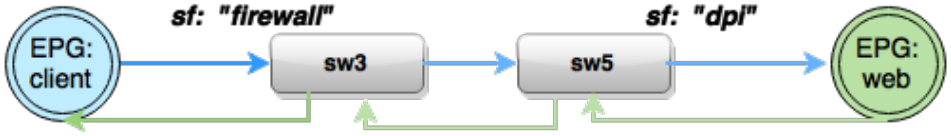

Fig. 61: GBP and SFC symmetric chain environment

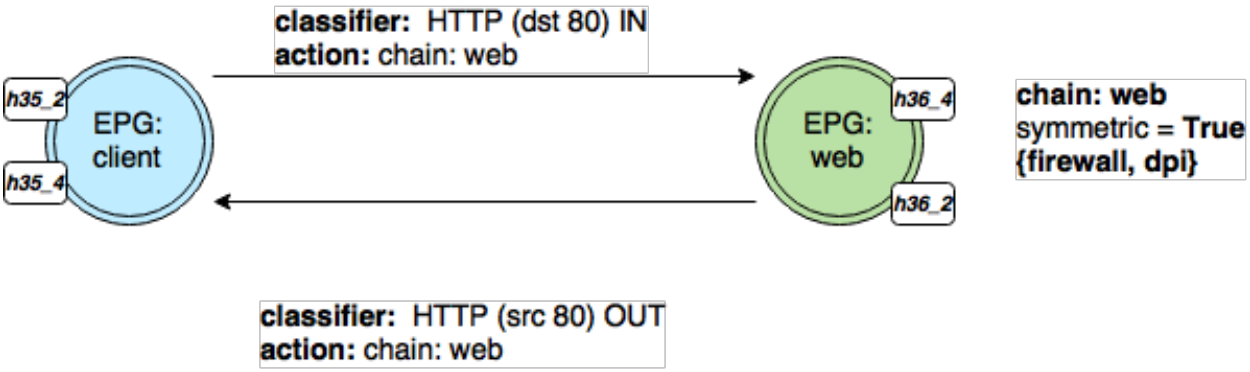

Fig. 62: GBP and SFC assymmetric chain environment

(continued from previous page)

```
    {
      "name": "http-dest",
      "parameter-value": [
        {
          "int-value": "6",
          "name": "proto"
        },
        {
          "int-value": "80",
          "name": "destport"
        }
      ]
    },
    {
      "name": "http-src",
      "parameter-value": [
        {
          "int-value": "6",
          "name": "proto"
        },
        {
          "int-value": "80",
          "name": "sourceport"
        }
      ]
    }
  ],
```

Then the action instances to associate to traffic that matches classifiers are defined.

Note the *SFC chain name* must exist in SFC, and is validated against the datastore once the tenant configuration is entered, before entering a valid tenant configuration into the operational datastore (which triggers policy resolution).

```
  "action-instance": [
    {
      "name": "chain1",
      "parameter-value": [
        {
          "name": "sfc-chain-name",
          "string-value": "SFCGBP"
        }
      ]
    },
    {
      "name": "allow1",
    }
  ]
},
```

When ICMP is matched, allow the traffic:

```
"contract": [
  {
    "subject": [
      {
        "name": "icmp-subject",
        "rule": [
```

(continues on next page)

(continued from previous page)

```
          {
            "name": "allow-icmp-rule",
            "order" : 0,
            "classifier-ref": [
              {
                "name": "icmp"
              }
            ],
            "action-ref": [
              {
                "name": "allow1",
                "order": 0
              }
            ]
          }

        ]
      },
```

When HTTP is matched, **in** to the provider of the contract with a TCP destination port of 80 (HTTP) or the HTTP request. The chain action is triggered, and similarly **out** from the provider for traffic with TCP source port of 80 (HTTP), or the HTTP response.

```
{
  "name": "http-subject",
  "rule": [
    {
      "name": "http-chain-rule-in",
      "classifier-ref": [
        {
          "name": "http-dest",
          "direction": "in"
        }
      ],
      "action-ref": [
        {
          "name": "chain1",
          "order": 0
        }
      ]
    },
    {
      "name": "http-chain-rule-out",
      "classifier-ref": [
        {
          "name": "http-src",
          "direction": "out"
        }
      ],
      "action-ref": [
        {
          "name": "chain1",
          "order": 0
        }
      ]
    }
```

(continues on next page)

(continued from previous page)

```
  ]
}
```

To enable asymmetrical chaining, for instance, the user desires that HTTP requests traverse the chain, but the HTTP response does not, the HTTP response is set to *allow* instead of chain:

```
{
  "name": "http-chain-rule-out",
  "classifier-ref": [
    {
      "name": "http-src",
      "direction": "out"
    }
  ],
  "action-ref": [
    {
      "name": "allow1",
      "order": 0
    }
  ]
}
```

Demo/Development environment

The **GBP** project for this release has two demo/development environments.

- Docker based GBP and GBP+SFC integration Vagrant environment
- DevStack based GBP+Neutron integration Vagrant environment

Demo @ GBP wiki

L2 Switch User Guide

Overview

The L2 Switch project provides Layer2 switch functionality.

L2 Switch Architecture

- Packet Handler
  - Decodes the packets coming to the controller and dispatches them appropriately
- Loop Remover
  - Removes loops in the network
- Arp Handler
  - Handles the decoded ARP packets
- Address Tracker
  - Learns the Addresses (MAC and IP) of entities in the network

- Host Tracker
  - Tracks the locations of hosts in the network
- L2 Switch Main
  - Installs flows on each switch based on network traffic

Configurable parameters in L2 Switch

The sections below give details about the configuration settings for the components that can be configured.

The process to change the configuration has been changed with the introduction of Blueprint in the Boron release. Please refer to *Change configuration in L2 Switch* for an example illustrating how to change the configurations.

Configurable parameters in Loop Remover

- l2switch/loopremover/implementation/src/main/yang/loop-remover-config.yang
  - is-install-lldp-flow
    - "true" means a flow that sends all LLDP packets to the controller will be installed on each switch
    - "false" means this flow will not be installed
    - default value is true
  - lldp-flow-table-id
    - The LLDP flow will be installed on the specified flow table of each switch
    - This field is only relevant when "is-install-lldp-flow" is set to "true"
    - default value is 0
  - lldp-flow-priority
    - The LLDP flow will be installed with the specified priority
    - This field is only relevant when "is-install-lldp-flow" is set to "true"
    - default value is 100
  - lldp-flow-idle-timeout
    - The LLDP flow will timeout (removed from the switch) if the flow doesn't forward a packet for *x* seconds
    - This field is only relevant when "is-install-lldp-flow" is set to "true"
    - default value is 0
  - lldp-flow-hard-timeout
    - The LLDP flow will timeout (removed from the switch) after *x* seconds, regardless of how many packets it is forwarding
    - This field is only relevant when "is-install-lldp-flow" is set to "true"
    - default value is 0
  - graph-refresh-delay
    - A graph of the network is maintained and gets updated as network elements go up/down (i.e. links go up/down and switches go up/down)

        * After a network element going up/down, it waits *graph-refresh-delay* seconds before recomputing the graph
        * A higher value has the advantage of doing less graph updates, at the potential cost of losing some packets because the graph didn't update immediately.
        * A lower value has the advantage of handling network topology changes quicker, at the cost of doing more computation.
        * default value is 1000

Configurable parameters in Arp Handler

- l2switch/arphandler/src/main/yang/arp-handler-config.yang
  - is-proactive-flood-mode
    * "true" means that flood flows will be installed on each switch. With this flood flow, each switch will flood a packet that doesn't match any other flows.
      · Advantage: Fewer packets are sent to the controller because those packets are flooded to the network.
      · Disadvantage: A lot of network traffic is generated.
    * "false" means the previously mentioned flood flows will not be installed. Instead an ARP flow will be installed on each switch that sends all ARP packets to the controller.
      · Advantage: Less network traffic is generated.
      · Disadvantage: The controller handles more packets (ARP requests & replies) and the ARP process takes longer than if there were flood flows.
    * default value is true
  - flood-flow-table-id
    * The flood flow will be installed on the specified flow table of each switch
    * This field is only relevant when "is-proactive-flood-mode" is set to "true"
    * default value is 0
  - flood-flow-priority
    * The flood flow will be installed with the specified priority
    * This field is only relevant when "is-proactive-flood-mode" is set to "true"
    * default value is 2
  - flood-flow-idle-timeout
    * The flood flow will timeout (removed from the switch) if the flow doesn't forward a packet for *x* seconds
    * This field is only relevant when "is-proactive-flood-mode" is set to "true"
    * default value is 0
  - flood-flow-hard-timeout
    * The flood flow will timeout (removed from the switch) after *x* seconds, regardless of how many packets it is forwarding
    * This field is only relevant when "is-proactive-flood-mode" is set to "true"

        - default value is 0
    - arp-flow-table-id
        - The ARP flow will be installed on the specified flow table of each switch
        - This field is only relevant when "is-proactive-flood-mode" is set to "false"
        - default value is 0
    - arp-flow-priority
        - The ARP flow will be installed with the specified priority
        - This field is only relevant when "is-proactive-flood-mode" is set to "false"
        - default value is 1
    - arp-flow-idle-timeout
        - The ARP flow will timeout (removed from the switch) if the flow doesn't forward a packet for *x* seconds
        - This field is only relevant when "is-proactive-flood-mode" is set to "false"
        - default value is 0
    - arp-flow-hard-timeout
        - The ARP flow will timeout (removed from the switch) after *arp-flow-hard-timeout* seconds, regardless of how many packets it is forwarding
        - This field is only relevant when "is-proactive-flood-mode" is set to "false"
        - default value is 0

Configurable parameters in Address Tracker

- l2switch/addresstracker/implementation/src/main/yang/address-tracker-config.yang
    - timestamp-update-interval
        - A last-seen timestamp is associated with each address. This last-seen timestamp will only be updated after *timestamp-update-interval* milliseconds.
        - A higher value has the advantage of performing less writes to the database.
        - A lower value has the advantage of knowing how fresh an address is.
        - default value is 600000
    - observe-addresses-from
        - IP and MAC addresses can be observed/learned from ARP, IPv4, and IPv6 packets. Set which packets to make these observations from.
        - default value is arp

Configurable parameters in L2 Switch Main

- l2switch/l2switch-main/src/main/yang/l2switch-config.yang
    - is-install-dropall-flow

    - "true" means a drop-all flow will be installed on each switch, so the default action will be to drop a packet instead of sending it to the controller
    - "false" means this flow will not be installed
    - default value is true
- dropall-flow-table-id
    - The dropall flow will be installed on the specified flow table of each switch
    - This field is only relevant when "is-install-dropall-flow" is set to "true"
    - default value is 0
- dropall-flow-priority
    - The dropall flow will be installed with the specified priority
    - This field is only relevant when "is-install-dropall-flow" is set to "true"
    - default value is 0
- dropall-flow-idle-timeout
    - The dropall flow will timeout (removed from the switch) if the flow doesn't forward a packet for *x* seconds
    - This field is only relevant when "is-install-dropall-flow" is set to "true"
    - default value is 0
- dropall-flow-hard-timeout
    - The dropall flow will timeout (removed from the switch) after *x* seconds, regardless of how many packets it is forwarding
    - This field is only relevant when "is-install-dropall-flow" is set to "true"
    - default value is 0
- is-learning-only-mode
    - "true" means that the L2 Switch will only be learning addresses. No additional flows to optimize network traffic will be installed.
    - "false" means that the L2 Switch will react to network traffic and install flows on the switches to optimize traffic. Currently, MAC-to-MAC flows are installed.
    - default value is false
- reactive-flow-table-id
    - The reactive flow will be installed on the specified flow table of each switch
    - This field is only relevant when "is-learning-only-mode" is set to "false"
    - default value is 0
- reactive-flow-priority
    - The reactive flow will be installed with the specified priority
    - This field is only relevant when "is-learning-only-mode" is set to "false"
    - default value is 10
- reactive-flow-idle-timeout

        * The reactive flow will timeout (removed from the switch) if the flow doesn't forward a packet for *x* seconds
        * This field is only relevant when "is-learning-only-mode" is set to "false"
        * default value is 600
    - reactive-flow-hard-timeout
        * The reactive flow will timeout (removed from the switch) after *x* seconds, regardless of how many packets it is forwarding
        * This field is only relevant when "is-learning-only-mode" is set to "false"
        * default value is 300

Change configuration in L2 Switch

Note: For more information on Blueprint in OpenDaylight, see this wiki page.

The following is an example on how to change the configurations of the L2 Switch components.

Use Case: Change the L2 switch from proactive flood mode to reactive mode.

Option 1: (external xml file)

1. Navigate to etc folder under download distribution
2. Create following directory structure:

```
mkdir - p opendaylight/datastore/initial/config
```

3. Create a new xml file corresponding to `<yang module name>_<container name>.xml`:

```
vi arp-handler-config_arp-handler-config.xml
```

4. Add following contents to the created file:

```
<?xml version="1.0" encoding="UTF-8"?>
  <arp-handler-config xmlns="urn:opendaylight:packet:arp-handler-config">
  <is-proactive-flood-mode>false</is-proactive-flood-mode>
</arp-handler-config>
```

5. Restart the controller which injects the configurations.

Option 2: (REST URL)

1. Make the following REST call
    - *URL:* `http://{{LOCALIP}}:8181/restconf/config/arp-handler-config:arp-handler-config/`
    - *Content-Type:* application/json
    - *Body*:

```
{
  "arp-handler-config":
  {
    "is-proactive-flood-mode":false
```

(continues on next page)

(continued from previous page)

```
    }
}
```

- *Expected Result:* 201 Created

2. Restart the controller to see updated configurations. With out a restart new configurations will be merged with old configurations which is not desirable.

Running the L2 Switch

To run the L2 Switch inside the OpenDaylight distribution simply install the `odl-l2switch-switch-ui` feature;

```
feature:install odl-l2switch-switch-ui
```

Create a network using mininet

```
sudo mn --controller=remote,ip=<Controller IP> --topo=linear,3 --switch ovsk,
↪protocols=OpenFlow13
sudo mn --controller=remote,ip=127.0.0.1 --topo=linear,3 --switch ovsk,
↪protocols=OpenFlow13
```

The above command will create a virtual network consisting of 3 switches. Each switch will connect to the controller located at the specified IP, i.e. 127.0.0.1

```
sudo mn --controller=remote,ip=127.0.0.1 --mac --topo=linear,3 --switch ovsk,
↪protocols=OpenFlow13
```

The above command has the "mac" option, which makes it easier to distinguish between Host MAC addresses and Switch MAC addresses.

Generating network traffic using mininet

```
h1 ping h2
```

The above command will cause host1 (h1) to ping host2 (h2)

```
pingall
```

pingall will cause each host to ping every other host.

Checking Address Observations

Address Observations are added to the Inventory data tree.

The Address Observations on a Node Connector can be checked through a browser or a REST Client.

```
http://10.194.126.91:8080/restconf/operational/opendaylight-inventory:nodes/node/
↪openflow:1/node-connector/openflow:1:1
```

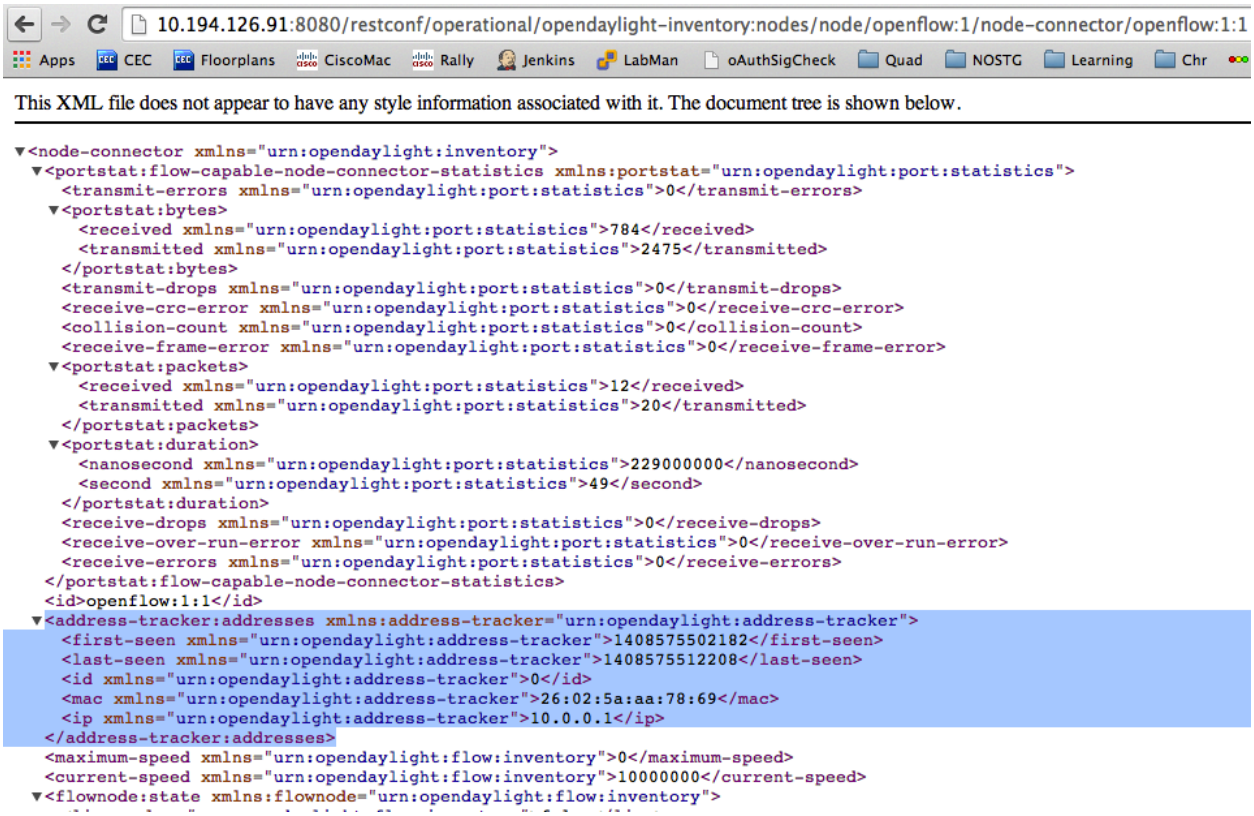

Fig. 63: Address Observations

Checking Hosts

Host information is added to the Topology data tree.

- Host address
- Attachment point (link) to a node/switch

This host information and attachment point information can be checked through a browser or a REST Client.

```
http://10.194.126.91:8080/restconf/operational/network-topology:network-topology/
↪topology/flow:1/
```

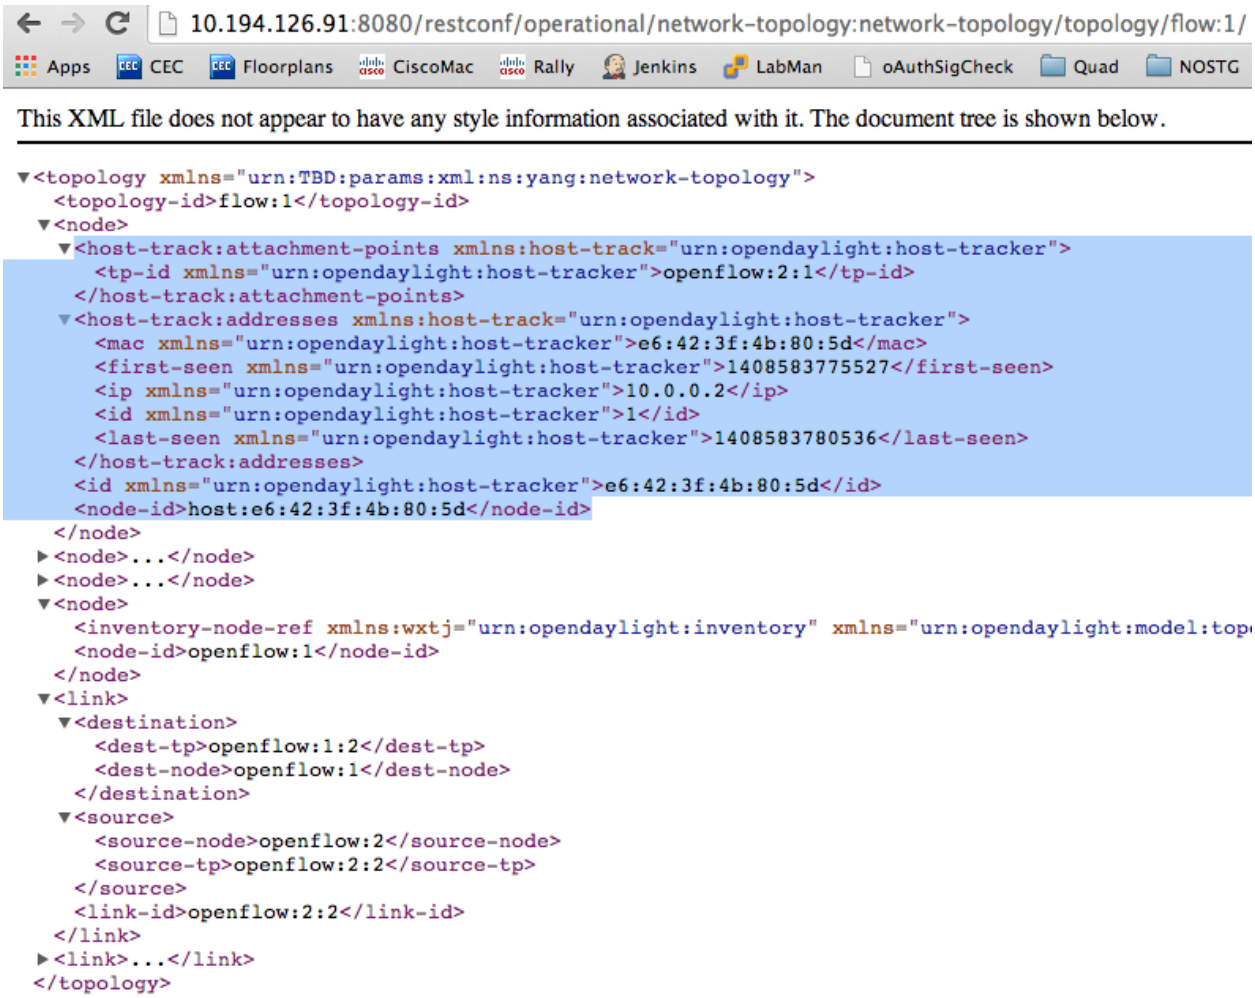

Fig. 64: Hosts

Checking STP status of each link

STP Status information is added to the Inventory data tree.

- A status of “forwarding” means the link is active and packets are flowing on it.
- A status of “discarding” means the link is inactive and packets are not sent over it.

The STP status of a link can be checked through a browser or a REST Client.

```
http://10.194.126.91:8080/restconf/operational/opendaylight-inventory:nodes/node/
↪openflow:1/node-connector/openflow:1:2
```

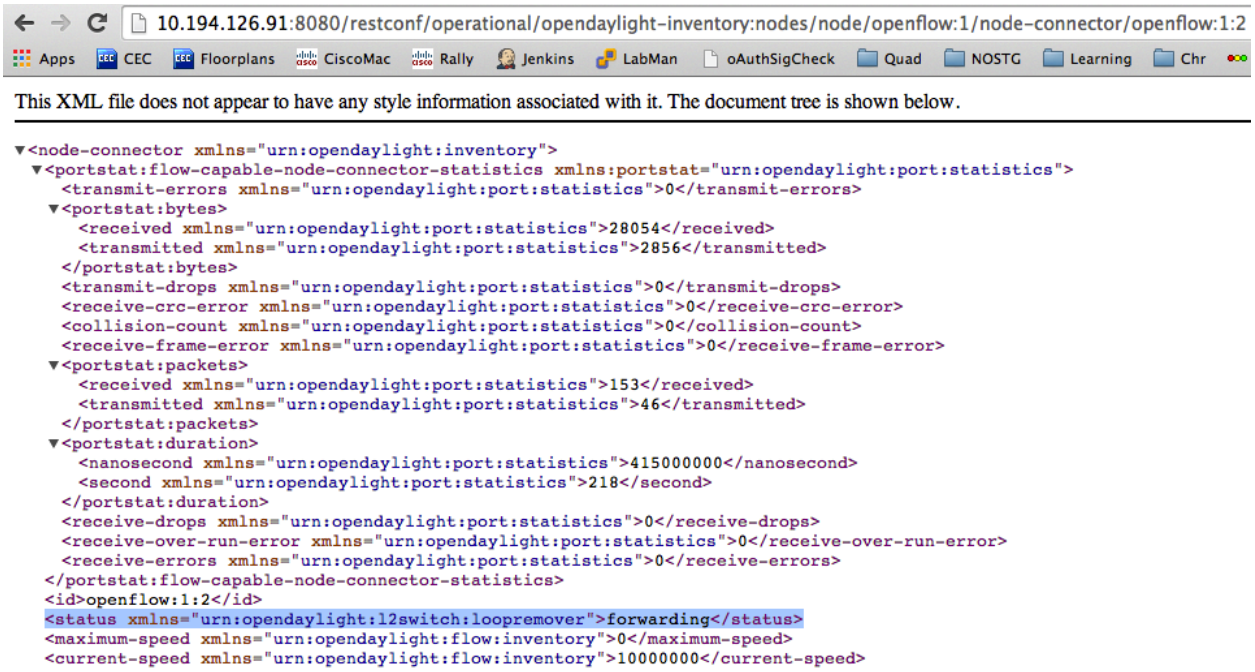

Fig. 65: STP status

Miscellaneous mininet commands

```
link s1 s2 down
```

This will bring the link between switch1 (s1) and switch2 (s2) down

```
link s1 s2 up
```

This will bring the link between switch1 (s1) and switch2 (s2) up

```
link s1 h1 down
```

This will bring the link between switch1 (s1) and host1 (h1) down

Link Aggregation Control Protocol User Guide

Overview

This section contains information about how to use the LACP plugin project with OpenDaylight, including configurations.

Link Aggregation Control Protocol Architecture

The LACP Project within OpenDaylight implements Link Aggregation Control Protocol (LACP) as an MD-SAL service module and will be used to auto-discover and aggregate multiple links between an OpenDaylight controlled network and LACP-enabled endpoints or switches. The result is the creation of a logical channel, which represents the aggregation of the links. Link aggregation provides link resiliency and bandwidth aggregation. This implementation adheres to IEEE Ethernet specification 802.3ad.

Configuring Link Aggregation Control Protocol

This feature can be enabled in the Karaf console of the OpenDaylight Karaf distribution by issuing the following command:

```
feature:install odl-lacp-ui
```

Note:

1. Ensure that legacy (non-OpenFlow) switches are configured with LACP mode active with a long timeout to allow for the LACP plugin in OpenDaylight to respond to its messages.
2. Flows that want to take advantage of LACP-configured Link Aggregation Groups (LAGs) must explicitly use a OpenFlow group table entry created by the LACP plugin. The plugin only creates group table entries, it does not program any flows on its own.

Administering or Managing Link Aggregation Control Protocol

LACP-discovered network inventory and network statistics can be viewed using the following REST APIs.

1. List of aggregators available for a node:

   ```
   http://<ControllerIP>:8181/restconf/operational/opendaylight-inventory:nodes/node/
   ↪<node-id>
   ```

   Aggregator information will appear within the `<lacp-aggregators>` XML tag.

2. To view only the information of an aggregator:

   ```
   http://<ControllerIP>:8181/restconf/operational/opendaylight-inventory:nodes/node/
   ↪<node-id>/lacp-aggregators/<agg-id>
   ```

   The group ID associated with the aggregator can be found inside the `<lag-groupid>` XML tag.

   The group table entry information for the `<lag-groupid>` added for the aggregator is also available in the `opendaylight-inventory` node database.

3. To view physical port information.

   ```
   http://<ControllerIP>:8181/restconf/operational/opendaylight-inventory:nodes/node/
   ↪<node-id>/node-connector/<node-connector-id>
   ```

   Ports that are associated with an aggregator will have the tag `<lacp-agg-ref>` updated with valid aggregator information.

Tutorials

The below tutorial demonstrates LACP LAG creation for a sample mininet topology.

Sample LACP Topology creation on Mininet

```
sudo mn --controller=remote,ip=<Controller IP> --topo=linear,1 --switch ovsk,
↪protocols=OpenFlow13
```

The above command will create a virtual network consisting of a switch and a host. The switch will be connected to the controller.

Once the topology is discovered, verify the presence of a flow entry with "dl_type" set to "0x8809" to handle LACP packets using the below ovs-ofctl command:

```
ovs-ofctl -O OpenFlow13 dump-flows s1
 OFPST_FLOW reply (OF1.3) (xid=0x2):
 cookie=0x300000000000001e, duration=60.067s, table=0, n_packets=0, n_bytes=0,␣
↪priority=5,dl_dst=01:80:c2:00:00:02,dl_type=0x8809 actions=CONTROLLER:65535
```

Configure an additional link between the switch (s1) and host (h1) using the below command on mininet shell to aggregate 2 links:

```
mininet> py net.addLink(s1, net.get('h1'))
mininet> py s1.attach('s1-eth2')
```

The LACP module will listen for LACP control packets that are generated from legacy switch (non-OpenFlow enabled). In our example, host (h1) will act as a LACP packet generator. In order to generate the LACP control packets, a bond interface has to be created on the host (h1) with mode type set to LACP with long-timeout. To configure bond interface, create a new file bonding.conf under the /etc/modprobe.d/ directory and insert the below lines in this new file:

```
alias bond0 bonding
options bonding mode=4
```

Here mode=4 is referred to LACP and the default timeout is set to long.

Enable bond interface and associate both physical interface h1-eth0 & h1-eth1 as members of the bond interface on host (h1) using the below commands on the mininet shell:

```
mininet> py net.get('h1').cmd('modprobe bonding')
mininet> py net.get('h1').cmd('ip link add bond0 type bond')
mininet> py net.get('h1').cmd('ip link set bond0 address <bond-mac-address>')
mininet> py net.get('h1').cmd('ip link set h1-eth0 down')
mininet> py net.get('h1').cmd('ip link set h1-eth0 master bond0')
mininet> py net.get('h1').cmd('ip link set h1-eth1 down')
mininet> py net.get('h1').cmd('ip link set h1-eth1 master bond0')
mininet> py net.get('h1').cmd('ip link set bond0 up')
```

Once the bond0 interface is up, the host (h1) will send LACP packets to the switch (s1). The LACP Module will then create a LAG through exchange of LACP packets between the host (h1) and switch (s1). To view the bond interface output on the host (h1) side:

```
mininet> py net.get('h1').cmd('cat /proc/net/bonding/bond0')
Ethernet Channel Bonding Driver: v3.7.1 (April 27, 2011)
```

(continues on next page)

(continued from previous page)

```
Bonding Mode: IEEE 802.3ad Dynamic link aggregation
Transmit Hash Policy: layer2 (0)
MII Status: up
MII Polling Interval (ms): 100
Up Delay (ms): 0
Down Delay (ms): 0
802.3ad info
LACP rate: slow
Min links: 0
Aggregator selection policy (ad_select): stable
Active Aggregator Info:
        Aggregator ID: 1
        Number of ports: 2
        Actor Key: 33
        Partner Key: 27
        Partner Mac Address: 00:00:00:00:01:01
```

```
Slave Interface: h1-eth0
MII Status: up
Speed: 10000 Mbps
Duplex: full
Link Failure Count: 0
Permanent HW addr: 00:00:00:00:00:11
Aggregator ID: 1
Slave queue ID: 0
```

```
Slave Interface: h1-eth1
MII Status: up
Speed: 10000 Mbps
Duplex: full
Link Failure Count: 0
Permanent HW addr: 00:00:00:00:00:12
Aggregator ID: 1
Slave queue ID: 0
```

A corresponding group table entry would be created on the OpenFlow switch (s1) with "type" set to "select" to perform the LAG functionality. To view the group entries:

```
mininet>ovs-ofctl -O Openflow13 dump-groups s1
OFPST_GROUP_DESC reply (OF1.3) (xid=0x2):
 group_id=60169,type=select,bucket=weight:0,actions=output:1,output:2
```

To apply the LAG functionality on the switches, the flows should be configured with action set to GroupId instead of output port. A sample add-flow configuration with output action set to GroupId:

```
sudo ovs-ofctl -O Openflow13 add-flow s1 dl_type=0x0806,dl_src=SRC_MAC,dl_dst=DST_MAC,
↪actions=group:60169
```

LISP Flow Mapping User Guide

Overview

Locator/ID Separation Protocol

Locator/ID Separation Protocol (LISP) is a technology that provides a flexible map-and-encap framework that can be used for overlay network applications such as data center network virtualization and Network Function Virtualization (NFV).

LISP provides the following name spaces:

- Endpoint Identifiers (EIDs)
- Routing Locators (RLOCs)

In a virtualization environment EIDs can be viewed as virtual address space and RLOCs can be viewed as physical network address space.

The LISP framework decouples network control plane from the forwarding plane by providing:

- A data plane that specifies how the virtualized network addresses are encapsulated in addresses from the underlying physical network.
- A control plane that stores the mapping of the virtual-to-physical address spaces, the associated forwarding policies and serves this information to the data plane on demand.

Network programmability is achieved by programming forwarding policies such as transparent mobility, service chaining, and traffic engineering in the mapping system; where the data plane elements can fetch these policies on demand as new flows arrive. This chapter describes the LISP Flow Mapping project in OpenDaylight and how it can be used to enable advanced SDN and NFV use cases.

LISP data plane Tunnel Routers are available at OpenOverlayRouter.org in the open source community on the following platforms:

- Linux
- Android
- OpenWRT

For more details and support for LISP data plane software please visit the OOR web site.

LISP Flow Mapping Service

The LISP Flow Mapping service provides LISP Mapping System services. This includes LISP Map-Server and LISP Map-Resolver services to store and serve mapping data to data plane nodes as well as to OpenDaylight applications. Mapping data can include mapping of virtual addresses to physical network address where the virtual nodes are reachable or hosted at. Mapping data can also include a variety of routing policies including traffic engineering and load balancing. To leverage this service, OpenDaylight applications and services can use the northbound REST API to define the mappings and policies in the LISP Mapping Service. Data plane devices capable of LISP control protocol can leverage this service through a southbound LISP plugin. LISP-enabled devices must be configured to use this OpenDaylight service as their Map Server and/or Map Resolver.

The southbound LISP plugin supports the LISP control protocol (Map-Register, Map-Request, Map-Reply messages), and can also be used to register mappings in the OpenDaylight mapping service.

LISP Flow Mapping Architecture

The following figure shows the various LISP Flow Mapping modules.

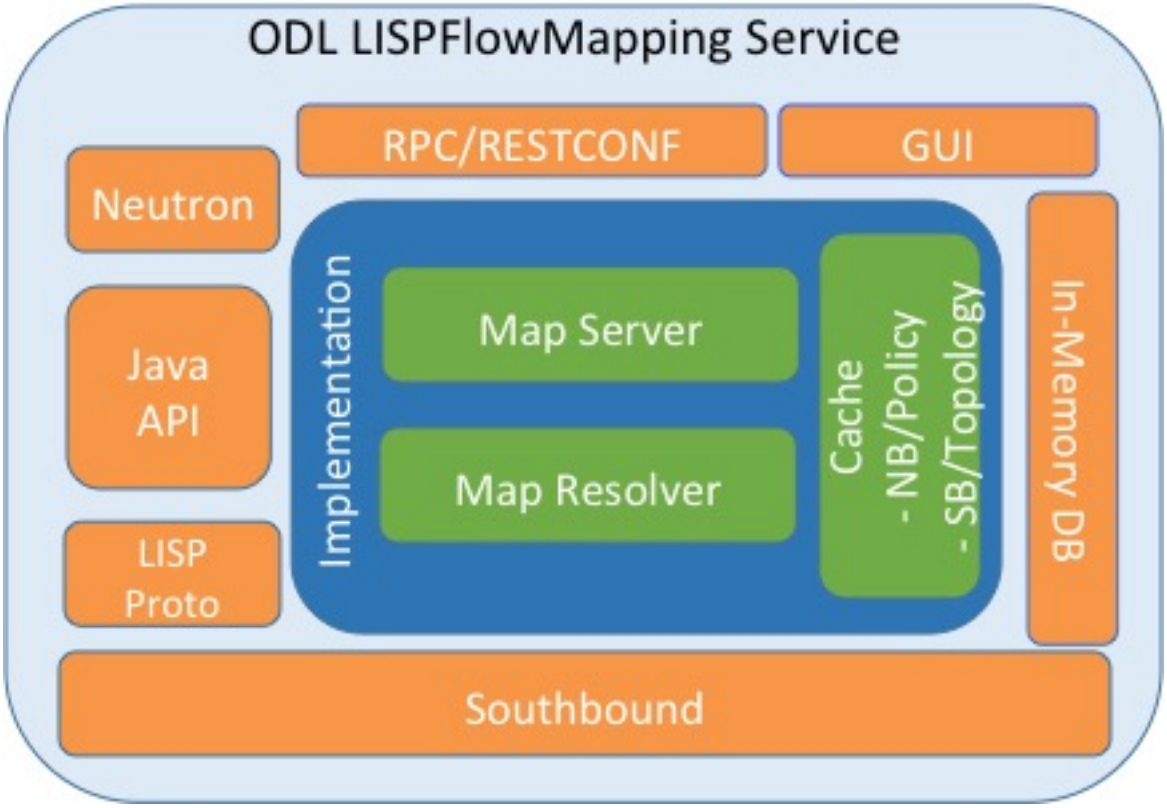

Fig. 66: LISP Mapping Service Internal Architecture

A brief description of each module is as follows:

- **DAO (Data Access Object):** This layer separates the LISP logic from the database, so that we can separate the map server and map resolver from the specific implementation of the mapping database. Currently we have an implementation of this layer with an in-memory HashMap, but it can be switched to any other key/value store and you only need to implement the ILispDAO interface.
- **Map Server:** This module processes the adding or registration of authentication tokens (keys) and mappings. For a detailed specification of LISP Map Server, see LISP.
- **Map Resolver:** This module receives and processes the mapping lookup queries and provides the mappings to requester. For a detailed specification of LISP Map Server, see LISP.
- **RPC/RESTCONF:** This is the auto-generated RESTCONF-based northbound API. This module enables defining key-EID associations as well as adding mapping information through the Map Server. Key-EID associations and mappings can also be queried via this API.
- **GUI:** This module enables adding and querying the mapping service through a GUI based on ODL DLUX.
- **Neutron:** This module implements the OpenDaylight Neutron Service APIs. It provides integration between the LISP service and the OpenDaylight Neutron service, and thus OpenStack.
- **Java API:** The API module exposes the Map Server and Map Resolver capabilities via a Java API.
- **LISP Proto:** This module includes LISP protocol dependent data types and associated processing.

- **In Memory DB:** This module includes the in memory database implementation of the mapping service.
- **LISP Southbound Plugin:** This plugin enables data plane devices that support LISP control plane protocol (see LISP) to register and query mappings to the LISP Flow Mapping via the LISP control plane protocol.

Configuring LISP Flow Mapping

In order to use the LISP mapping service for registering EID to RLOC mappings from northbound or southbound, keys have to be defined for the EID prefixes first. Once a key is defined for an EID prefix, it can be used to add mappings for that EID prefix multiple times. If the service is going to be used to process Map-Register messages from the southbound LISP plugin, the same key must be used by the data plane device to create the authentication data in the Map-Register messages for the associated EID prefix.

The `etc/custom.properties` file in the Karaf distribution allows configuration of several OpenDaylight parameters. The LISP service has the following properties that can be adjusted:

lisp.mappingOverwrite (default: *true*) Configures handling of mapping updates. When set to *true* (default) a mapping update (either through the southbound plugin via a Map-Register message or through a northbound API PUT REST call) the existing RLOC set associated to an EID prefix is overwritten. When set to *false*, the RLOCs of the update are merged to the existing set.

lisp.smr (default: *false*) Enables/disables the Solicit-Map-Request (SMR) functionality. SMR is a method to notify changes in an EID-to-RLOC mapping to "subscribers". The LISP service considers all Map-Request's source RLOC as a subscriber to the requested EID prefix, and will send an SMR control message to that RLOC if the mapping changes.

lisp.elpPolicy (default: *default*) Configures how to build a Map-Reply southbound message from a mapping containing an Explicit Locator Path (ELP) RLOC. It is used for compatibility with dataplane devices that don't understand the ELP LCAF format. The *default* setting doesn't alter the mapping, returning all RLOCs unmodified. The *both* setting adds a new RLOC to the mapping, with a lower priority than the ELP, that is the next hop in the service chain. To determine the next hop, it searches the source RLOC of the Map-Request in the ELP, and chooses the next hop, if it exists, otherwise it chooses the first hop. The *replace* setting adds a new RLOC using the same algorithm as the *both* setting, but using the origin priority of the ELP RLOC, which is removed from the mapping.

lisp.lookupPolicy (default: *northboundFirst*) Configures the mapping lookup algorithm. When set to *northboundFirst* mappings programmed through the northbound API will take precedence. If no northbound programmed mappings exist, then the mapping service will return mappings registered through the southbound plugin, if any exists. When set to *northboundAndSouthbound* the mapping programmed by the northbound is returned, updated by the up/down status of these mappings as reported by the southbound (if existing).

lisp.mappingMerge (default: *false*) Configures the merge policy on the southbound registrations through the LISP SB Plugin. When set to *false*, only the latest mapping registered through the SB plugin is valid in the southbound mapping database, independent of which device it came from. When set to *true*, mappings for the same EID registered by different devices are merged together and a union of the locators is maintained as the valid mapping for that EID.

Textual Conventions for LISP Address Formats

In addition to the more common IPv4, IPv6 and MAC address data types, the LISP control plane supports arbitrary Address Family Identifiers assigned by IANA, and in addition to those the LISP Canoncal Address Format (LCAF).

The LISP Flow Mapping project in OpenDaylight implements support for many of these different address formats, the full list being summarized in the following table. While some of the address formats have well defined and widely used textual representation, many don't. It became necessary to define a convention to use for text rendering of all implemented address types in logs, URLs, input fields, etc. The below table lists the supported formats, along with

their AFI number and LCAF type, including the prefix used for disambiguation of potential overlap, and examples output.

Name	AFI	LCAF	Prefix	Text Rendering
No Address	0	•	no:	No Address Present
IPv4 Prefix	1	•	ipv4:	192.0.2.0/24
IPv6 Prefix	2	•	ipv6:	2001:db8::/32
MAC Address	16389	•	mac:	00:00:5E:00:53:00
Distinguished Name	17	•	dn:	stringAsIs
AS Number	18	•	as:	AS64500
AFI List	16387	1	list:	{192.0.2.1,192.0.2.2,2001:db8::1 }
Instance ID	16387	2	•	[223] 192.0.2.0/24
Application Data	16387	4	appdata:	192.0.2.1!128!17!80-81!6667-7000
Explicit Locator Path	16387	10	elp:	{192.0.2.1→192.0.2.2|lps→192.0.2.3}
Source/Destination Key	16387	12	srcdst:	192.0.2.1/32|192.0.2.2/32
Key/Value Address Pair	16387	15	kv:	192.0.2.1192.0.2.2
Service Path	16387	N/A	sp:	42(3)

Table: LISP Address Formats

Please note that the forward slash character `/` typically separating IPv4 and IPv6 addresses from the mask length is transformed into `%2f` when used in a URL.

Karaf commands

In this section we will discuss two types of Karaf commands: built-in, and LISP specific. Some built-in commands are quite useful, and are needed for the tutorial, so they will be discussed here. A reference of all LISP specific commands, added by the LISP Flow Mapping project is also included. They are useful mostly for debugging.

Useful built-in commands

`help` Lists all available command, with a short description of each.

`help <command_name>` Show detailed help about a specific command.

`feature:list [-i]` Show all locally available features in the Karaf container. The `-i` option lists only features that are currently installed. It is possible to use `| grep` to filter the output (for all commands, not just this one).

feature:install <feature_name> Install feature `feature_name`.

log:set <level> <class> Set the log level for `class` to `level`. The default log level for all classes is INFO. For debugging, or learning about LISP internals it is useful to run `log:set TRACE org.opendaylight.lispflowmapping` right after Karaf starts up.

log:display Outputs the log file to the console, and returns control to the user.

log:tail Continuously shows log output, requires `Ctrl+C` to return to the console.

LISP specific commands

The available lisp commands can always be obtained by `help mappingservice`. Currently they are:

mappingservice:addkey Add the default password `password` for the IPv4 EID prefix 0.0.0.0/0 (all addresses). This is useful when experimenting with southbound devices, and using the REST interface would be combersome for whatever reason.

mappingservice:mappings Show the list of all mappings stored in the internal non-persistent data store (the DAO), listing the full data structure. The output is not human friendly, but can be used for debugging.

LISP Flow Mapping Karaf Features

LISP Flow Mapping has the following Karaf features that can be installed from the Karaf console:

odl-lispflowmapping-msmr This includes the core features required to use the LISP Flow Mapping Service such as mapping service and the LISP southbound plugin.

odl-lispflowmapping-ui This includes the GUI module for the LISP Mapping Service.

odl-lispflowmapping-neutron This is the experimental Neutron provider module for LISP mapping service.

Tutorials

This section provides a tutorial demonstrating various features in this service. We have included tutorials using two forwarding platforms:

1. Using Open Overlay Router (OOR)
2. Using FD.io

Both have different approaches to create the overlay but ultimately do the same job. Details of both approaches have been explained below.

Creating a LISP overlay with OOR

This section provides instructions to set up a LISP network of three nodes (one "client" node and two "server" nodes) using OOR as data plane LISP nodes and the LISP Flow Mapping project from OpenDaylight as the LISP programmable mapping system for the LISP network.

Overview

The steps shown below will demonstrate setting up a LISP network between a client and two servers, then performing a failover between the two "server" nodes.

Prerequisites

- **OpenDaylight Boron**
- **The Postman Chrome App**: the most convenient way to follow along this tutorial is to use the Postman App to edit and send the requests. The project git repository hosts a collection of the requests that are used in this tutorial in the `resources/tutorial/OOR/Beryllium_Tutorial.json.postman_collection` file. You can import this file to Postman by clicking *Import* at the top, choosing *Download from link* and then entering the following URL: https://git.opendaylight.org/gerrit/gitweb?p=lispflowmapping.git;a=blob_plain;f=resources/tutorial/OOR/Beryllium_Tutorial.json.postman_collection;hb=refs/heads/stable/boron. Alternatively, you can save the file on your machine, or if you have the repository checked out, you can import from there. You will need to create a new Postman Environment and define some variables within: `controllerHost` set to the hostname or IP address of the machine running the OpenDaylight instance, and `restconfPort` to 8181, if you didn't modify the default controller settings.
- **OOR version 1.0 or later** The README.md lists the dependencies needed to build it from source.
- **A virtualization platform**

Target Environment

The three LISP data plane nodes and the LISP mapping system are assumed to be running in Linux virtual machines, which have the `eth0` interface in NAT mode to allow outside internet access and `eth1` connected to a host-only network, with the following IP addresses (please adjust configuration files, JSON examples, etc. accordingly if you're using another addressing scheme):

Node	Node Type	IP Address
controller	OpenDaylight	192.168.16.11
client	OOR	192.168.16.30
server1	OOR	192.168.16.31
server2	OOR	192.168.16.32
service-node	OOR	192.168.16.33

Table: Nodes in the tutorial

The figure below gives a sketch of network topology that will be used in the tutorial.

In LISP terminology **client**, **server1** and **server2** are mobile nodes (MN in OOR), **controller** is a MS/MR and **service-node** is a RTR.

Note: While the tutorial uses OOR as the data plane, it could be any LISP-enabled hardware or software router (commercial/open source).

Instructions

The below steps use the command line tool cURL to talk to the LISP Flow Mapping RPC REST API. This is so that you can see the actual request URLs and body content on the page.

1. Install and run OpenDaylight Boron release on the controller VM. Please follow the general OpenDaylight Boron Installation Guide for this step. Once the OpenDaylight controller is running install the *odl-lispflowmapping-msmr* feature from the Karaf CLI:

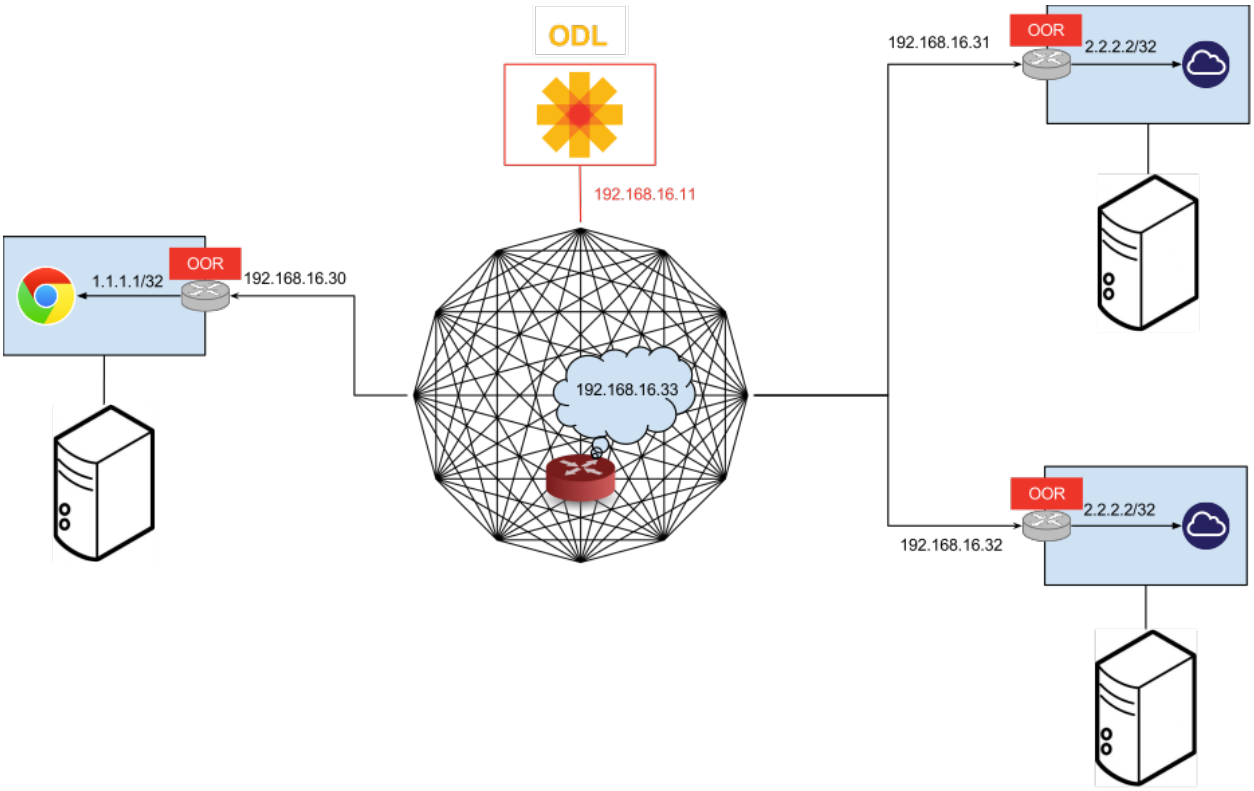


```
feature:install odl-lispflowmapping-msmr
```

It takes quite a while to load and initialize all features and their dependencies. It's worth running the command `log:tail` in the Karaf console to see when the log output is winding down, and continue with the tutorial after that.

2. Install OOR on the **client**, **server1**, **server2**, and **service-node** VMs following the installation instructions from the OOR README file.

3. Configure the OOR installations from the previous step. Take a look at the `oor.conf.example` to get a general idea of the structure of the conf file. First, check if the file `/etc/oor.conf` exists. If the file doesn't exist, create the file `/etc/oor.conf`. Set the EID in `/etc/oor.conf` file from the IP address space selected for your virtual/LISP network. In this tutorial the EID of the **client** is set to 1.1.1.1/32, and that of **server1** and **server2** to 2.2.2.2/32.

4. Set the RLOC interface to `eth1` in each `oor.conf` file. LISP will determine the RLOC (IP address of the corresponding VM) based on this interface.

5. Set the Map-Resolver address to the IP address of the **controller**, and on the **client** the Map-Server too. On **server1** and **server2** remove the Map-Server configuration, so that it doesn't interfere with the mappings on the controller, since we're going to program them manually.

6. Modify the "key" parameter in each `oor.conf` file to a key/password of your choice (*password* in this tutorial).

Note: The `resources/tutorial/OOR` directory in the *stable/boron* branch of the project git repository has the files used in the tutorial checked in, so you can just copy the files to `/etc/oor.conf` on the respective VMs. You will also find the JSON files referenced below in the same directory.

---

7. Define a key and EID prefix association in OpenDaylight using the RPC REST API for the **client** EID

(1.1.1.1/32) to allow registration from the southbound. Since the mappings for the server EID will be configured from the REST API, no such association is necessary. Run the below command on the **controller** (or any machine that can reach **controller**, by replacing *localhost* with the IP address of **controller**).

```
curl -u "admin":"admin" -H "Content-type: application/json" -X PUT \
    http://localhost:8181/restconf/config/odl-mappingservice:mapping-database/
↪virtual-network-identifier/0/authentication-key/ipv4:1.1.1.1%2f32/ \
    --data @add-key.json
```

where the content of the *add-key.json* file is the following:

```
{
    "authentication-key": {
        "eid-uri": "ipv4:1.1.1.1/32",
        "eid": {
            "address-type": "ietf-lisp-address-types:ipv4-prefix-afi",
            "ipv4-prefix": "1.1.1.1/32"
        },
        "mapping-authkey": {
            "key-string": "password",
            "key-type": 1
        }
    }
}
```

8. Verify that the key is added properly by requesting the following URL:

```
curl -u "admin":"admin" -H "Content-type: application/json" -X GET \
    http://localhost:8181/restconf/config/odl-mappingservice:mapping-database/
↪virtual-network-identifier/0/authentication-key/ipv4:1.1.1.1%2f32/
```

The output the above invocation should look like this:

```
{
    "authentication-key":[
        {
            "eid-uri":"ipv4:1.1.1.1/32",
            "eid":{
                "ipv4-prefix":"1.1.1.1/32",
                "address-type":"ietf-lisp-address-types:ipv4-prefix-afi"
            },
            "mapping-authkey":{
                "key-string":"password"
                ,"key-type":1
            }
        }
    ]
}
```

9. Run the `oor` OOR daemon on all VMs:

```
oor -f /etc/oor.conf
```

For more information on accessing OOR logs, take a look at OOR README

10. The **client** OOR node should now register its EID-to-RLOC mapping in OpenDaylight. To verify you can lookup the corresponding EIDs via the REST API

```
curl -u "admin":"admin" -H "Content-type: application/json" -X GET \
    http://localhost:8181/restconf/operational/odl-mappingservice:mapping-
↪database/virtual-network-identifier/0/mapping/ipv4:1.1.1.1%2f32/southbound/
```

An alternative way for retrieving mappings from OpenDaylight using the southbound interface is using the `lig <https://github.com/davidmeyer/lig>`__ open source tool.

11. Register the EID-to-RLOC mapping of the server EID 2.2.2.2/32 to the controller, pointing to **server1** and **server2** with a higher priority for **server1**

```
curl -u "admin":"admin" -H "Content-type: application/json" -X PUT \
    http://localhost:8181/restconf/config/odl-mappingservice:mapping-database/
↪virtual-network-identifier/0/mapping/ipv4:2.2.2.2%2f32/northbound/ \
    --data @mapping.json
```

where the *mapping.json* file looks like this:

```
{
    "mapping": {
        "eid-uri": "ipv4:2.2.2.2/32",
        "origin": "northbound",
        "mapping-record": {
            "recordTtl": 1440,
            "action": "NoAction",
            "authoritative": true,
            "eid": {
                "address-type": "ietf-lisp-address-types:ipv4-prefix-afi",
                "ipv4-prefix": "2.2.2.2/32"
            },
            "LocatorRecord": [
                {
                    "locator-id": "server1",
                    "priority": 1,
                    "weight": 1,
                    "multicastPriority": 255,
                    "multicastWeight": 0,
                    "localLocator": true,
                    "rlocProbed": false,
                    "routed": true,
                    "rloc": {
                        "address-type": "ietf-lisp-address-types:ipv4-afi",
                        "ipv4": "192.168.16.31"
                    }
                },
                {
                    "locator-id": "server2",
                    "priority": 2,
                    "weight": 1,
                    "multicastPriority": 255,
                    "multicastWeight": 0,
                    "localLocator": true,
                    "rlocProbed": false,
                    "routed": true,
                    "rloc": {
                        "address-type": "ietf-lisp-address-types:ipv4-afi",
                        "ipv4": "192.168.16.32"
                    }
```

(continues on next page)

(continued from previous page)

```
                }
            ]
        }
    }
}
```

Here the priority of the second RLOC (192.168.16.32 - **server2**) is 2, a higher numeric value than the priority of 192.168.16.31, which is 1. This policy is saying that **server1** is preferred to **server2** for reaching EID 2.2.2.2/32. Note that lower priority value has higher preference in LISP.

12. Verify the correct registration of the 2.2.2.2/32 EID:

```
curl -u "admin":"admin" -H "Content-type: application/json" -X GET \
    http://localhost:8181/restconf/config/odl-mappingservice:mapping-database/
↪virtual-network-identifier/0/mapping/ipv4:2.2.2.2%2f32/northbound/
```

13. Now the LISP network is up. To verify, log into the **client** VM and ping the server EID:

```
ping 2.2.2.2
```

14. Let's test fail-over now. Suppose you had a service on **server1** which became unavailable, but **server1** itself is still reachable. LISP will not automatically fail over, even if the mapping for 2.2.2.2/32 has two locators, since both locators are still reachable and uses the one with the higher priority (lowest priority value). To force a failover, we need to set the priority of **server2** to a lower value. Using the file mapping.json above, swap the priority values between the two locators (lines 14 and 28 in *mapping.json*) and repeat the request from step 11. You can also repeat step 12 to see if the mapping is correctly registered. If you leave the ping on, and monitor the traffic using wireshark, you can see that the ping traffic to 2.2.2.2 will be diverted from the **server1** RLOC to the **server2** RLOC.

    With the default OpenDaylight configuration the failover should be near instantaneous (we observed 3 lost pings in the worst case), because of the LISP Solicit-Map-Request (SMR) mechanism that can ask a LISP data plane element to update its mapping for a certain EID (enabled by default). It is controlled by the `lisp.smr` variable in `etc/custom.porperties`. When enabled, any mapping change from the RPC interface will trigger an SMR packet to all data plane elements that have requested the mapping in the last 24 hours (this value was chosen because it's the default TTL of Cisco IOS xTR mapping registrations). If disabled, ITRs keep their mappings until the TTL specified in the Map-Reply expires.

15. To add a service chain into the path from the client to the server, we can use an Explicit Locator Path, specifying the **service-node** as the first hop and **server1** (or **server2**) as the second hop. The following will achieve that:

```
curl -u "admin":"admin" -H "Content-type: application/json" -X PUT \
    http://localhost:8181/restconf/config/odl-mappingservice:mapping-database/
↪virtual-network-identifier/0/mapping/ipv4:2.2.2.2%2f32/northbound/ \
    --data @elp.json
```

where the *elp.json* file is as follows:

```
{
    "mapping": {
        "eid-uri": "ipv4:2.2.2.2/32",
        "origin": "northbound",
        "mapping-record": {
            "recordTtl": 1440,
            "action": "NoAction",
            "authoritative": true,
            "eid": {
```

(continues on next page)

(continued from previous page)

```
                "address-type": "ietf-lisp-address-types:ipv4-prefix-afi",
                "ipv4-prefix": "2.2.2.2/32"
            },
            "LocatorRecord": [
                {
                    "locator-id": "ELP",
                    "priority": 1,
                    "weight": 1,
                    "multicastPriority": 255,
                    "multicastWeight": 0,
                    "localLocator": true,
                    "rlocProbed": false,
                    "routed": true,
                    "rloc": {
                        "address-type": "ietf-lisp-address-types:explicit-locator-
↪path-lcaf",
                        "explicit-locator-path": {
                            "hop": [
                                {
                                    "hop-id": "service-node",
                                    "address": "192.168.16.33",
                                    "lrs-bits": "strict"
                                },
                                {
                                    "hop-id": "server1",
                                    "address": "192.168.16.31",
                                    "lrs-bits": "strict"
                                }
                            ]
                        }
                    }
                }
            ]
        }
    }
}
```

After the mapping for 2.2.2.2/32 is updated with the above, the ICMP traffic from **client** to **server1** will flow through the **service-node**. You can confirm this in the OOR logs, or by sniffing the traffic on either the **service-node** or **server1**. Note that service chains are unidirectional, so unless another ELP mapping is added for the return traffic, packets will go from **server1** to **client** directly.

16. Suppose the **service-node** is actually a firewall, and traffic is diverted there to support access control lists (ACLs). In this tutorial that can be emulated by using `iptables` firewall rules in the **service-node** VM. To deny traffic on the service chain defined above, the following rule can be added:

```
iptables -A OUTPUT --dst 192.168.16.31 -j DROP
```

The ping from the **client** should now have stopped.

In this case the ACL is done on the destination RLOC. There is an effort underway in the OOR community to allow filtering on EIDs, which is the more logical place to apply ACLs.

17. To delete the rule and restore connectivity on the service chain, delete the ACL by issuing the following command:

```
iptables -D OUTPUT --dst 192.168.16.31 -j DROP
```

which should restore connectivity.

Creating a simple LISP overlay with FD.io

In this section, we use the Overlay Network Engine (ONE) project in FD.io to facilitate fully scripted setup and testing of a LISP/VXLAN-GPE network. Overlay Network Engine (ONE) is a FD.io project that enables programmable dynamic software defined overlays. Details about this project can be found in ONE wiki.

The steps shown below will demonstrate setting up a LISP network between a client and a server using VPP. We demonstrate how to use VPP lite to build a IP4 LISP overlay on an Ubuntu host using namespaces and af_packet interfaces. All configuration files used in the tutorials can be found here.

Prerequisites

- **OpenDaylight Boron**
- **The Postman Chrome App**: Please follow the *instructions* and import postman collection from the following URL: https://git.opendaylight.org/gerrit/gitweb?p=lispflowmapping.git;a=blob;f=resources/tutorial/FD_io/lfm_vpp.postman_collection.json;hb=HEAD.
- **Vagrant** (optional): Download it from Vagrant website and follow the setup instructions.

Target Environment

Unlike the case with OOR, we use network namespace functionality of Linux to create the overlay in this case. The following table contains ip addresses of nodes in the overlay topology used in the tutorial. Our objective will be to create this topology and be able to ping from client to server through an intermediary hop, **service node**, which is a `rtr node` providing the service of re-encapsulation. So, all the packets from client to server will be through this **service node**.

Node	Node Type	IP Address
controller	OpenDaylight	6.0.3.100
client	VPP	6.0.2.2
server	VPP	6.0.4.4
service node	VPP	6.0.3.3

Table: Nodes in the tutorial

The figure below gives a sketch of network topology that will be used in the tutorial.

Instructions

Follow the instructions below sequentially.

1. Pull the VPP code anonymously using:

```
git clone https://gerrit.fd.io/r/vpp
```

2. Then, use the vagrant file from repository to build virtual machine with proper environment.

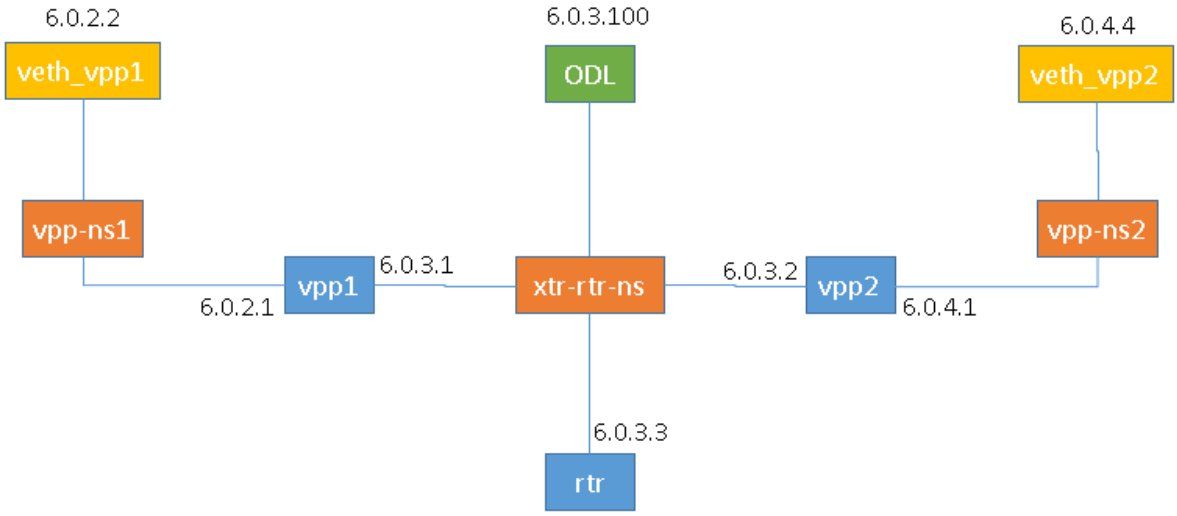


```
cd vpp/build-root/vagrant/
vagrant up
vagrant ssh
```

3. In case there is any error from `vagrant up`, try `vargant ssh`. if it works, no worries. If it still doesn't work, you can try any Ubuntu virtual machine. Or sometimes there is an issue with the Vagrant properly copying the VPP repo code from the host VM after the first installation. In that case `/vpp` doesn't exist. In both cases, follow the instructions from below.

   (a) Clone the code in `/` directory. So, the codes will be in `/vpp`.

   (b) **Run the following commands:**

   ```
   cd /vpp/build-root
   make distclean
   ./bootstrap.sh
   make V=0 PLATFORM=vpp TAG=vpp install-deb
   sudo dpkg -i /vpp/build-root/*.deb
   ```

   Alternative and more detailed build instructions can be found in VPP's wiki

4. By now, you should have a Ubuntu VM with VPP repository in `/vpp` with `sudo` access. Now, we need VPP Lite build. The following commands builds VPP Lite.

   ```
   cd /vpp
   export PLATFORM=vpp_lite
   make build
   ```

   Successful build create the binary in `/vpp/build-root/install-vpp_lite_debug-native/vpp/bin`

5. Install bridge-utils and ethtool if needed by using following commands:

   ```
   sudo apt-get install bridge-utils ethtool
   ```

6. Now, install and run OpenDaylight Boron release on the VM. Please follow the general OpenDaylight Boron Installation Guide for this step from *Installing OpenDaylight*. Before running OpenDaylight, we need to change

the configuration for RTR to work. Update `etc/custom.properties` with the `lisp.elpPolicy` to be replace.

```
lisp.elpPolicy = replace
```

Then, run OpenDaylight. For details regarding configuring LISP Flow Mapping, please take a look at *Configuring LISP Flow Mapping*. Once the OpenDaylight controller is running install the *odl-lispflowmapping-msmr* feature from the Karaf CLI:

```
feature:install odl-lispflowmapping-msmr
```

It may take quite a while to load and initialize all features and their dependencies. It's worth running the command `log:tail` in the Karaf console to see when the log output is winding down, and continue with the tutorial after that.

7. For setting up VPP, get the files from `resources/tutorial/FD_io` folder of the lispflowmapping repo. The files can also be found here. Copy the `vpp1.config`, `vpp2.config` and `rtr.config` files in `/etc/vpp/lite/`.

8. In this example, VPP doesn't make any southbound map registers to OpenDaylight. So, we add the mappings directly from northbound. For that, we need to add the mappings to OpenDaylight via RESTCONF API.

   Register EID-to-RLOC mapping of the Client EID 6.0.2.0/24.

```
curl -u "admin":"admin" -H "Content-type: application/json" -X PUT \
    http://localhost:8181/restconf/config/odl-mappingservice:mapping-database/
↪virtual-network-identifier/0/mapping/ipv4:6.0.2.0%2f24/northbound/ \
    --data @epl1.json
```

   Content of epl1.json:

```
{
    "mapping": {
        "eid-uri": "ipv4:6.0.2.0/24",
        "origin": "northbound",
        "mapping-record": {
            "recordTtl": 1440,
            "action": "NoAction",
            "authoritative": true,
            "eid": {
                    "address-type": "ietf-lisp-address-types:ipv4-prefix-afi",
                    "ipv4-prefix": "6.0.2.0/24"
            },
            "LocatorRecord": [
                {
                    "locator-id": "ELP",
                    "priority": 1,
                    "weight": 1,
                    "multicastPriority": 255,
                    "multicastWeight": 0,
                    "localLocator": true,
                    "rlocProbed": false,
                    "routed": false,
                    "rloc": {
                        "address-type": "ietf-lisp-address-types:explicit-locator-
↪path-lcaf",
                        "explicit-locator-path": {
                            "hop": [
```

(continues on next page)

(continued from previous page)

```
                                    {
                                        "hop-id": "Hop 1",
                                        "address": "6.0.3.3",
                                        "lrs-bits": "lookup rloc-probe strict"
                                    },
                                    {
                                        "hop-id": "Hop 2",
                                        "address": "6.0.3.1",
                                        "lrs-bits": "lookup strict"
                                    }
                                ]
                            }
                        }
                    }
                ]
            }
        }
}
```

Similarly add EID-to-RLOC mapping of the Server EID 6.0.4.0/24.

```
curl -u "admin":"admin" -H "Content-type: application/json" -X PUT \
    http://localhost:8181/restconf/config/odl-mappingservice:mapping-database/
↪virtual-network-identifier/0/mapping/ipv4:6.0.4.0%2f24/northbound/ \
    --data @epl2.json
```

Content of elp2.json:

```
{
    "mapping": {
        "eid-uri": "ipv4:6.0.4.0/24",
        "origin": "northbound",
        "mapping-record": {
            "recordTtl": 1440,
            "action": "NoAction",
            "authoritative": true,
            "eid": {
                    "address-type": "ietf-lisp-address-types:ipv4-prefix-afi",
                    "ipv4-prefix": "6.0.4.0/24"
            },
            "LocatorRecord": [
                {
                    "locator-id": "ELP",
                    "priority": 1,
                    "weight": 1,
                    "multicastPriority": 255,
                    "multicastWeight": 0,
                    "localLocator": true,
                    "rlocProbed": false,
                    "routed": false,
                    "rloc": {
                        "address-type": "ietf-lisp-address-types:explicit-locator-
↪path-lcaf",
                        "explicit-locator-path": {
                            "hop": [
                                {
                                    "hop-id": "Hop 1",
```

(continues on next page)

(continued from previous page)

```
                                    "address": "6.0.3.3",
                                    "lrs-bits": "lookup rloc-probe strict"
                                },
                                {
                                    "hop-id": "Hop 2",
                                    "address": "6.0.3.2",
                                    "lrs-bits": "lookup strict"
                                }
                            ]
                        }
                    }
                }
            ]
        }
    }
}
```

The JSON files regarding these can be found in here. Even though there is no southbound registration for mapping to OpenDaylight, using northbound policy we can specify mappings, when Client requests for the Server eid, Client gets a reply from OpenDaylight.

9. Assuming all files have been created and OpenDaylight has been configured as explained above, execute the host script you've created or the `topology_setup.sh` script from here.
10. If all goes well, you can now test connectivity between the namespaces with:

```
sudo ip netns exec vpp-ns1 ping 6.0.4.4
```

11. Traffic and control plane message exchanges can be checked with a wireshark listening on the odl interface.
12. **Important:** Delete the topology by running the `topology_setup.sh` with `clean` argument.

```
sudo ./topology_setup.sh clean
```

LISP Flow Mapping Support

For support the lispflowmapping project can be reached by emailing the developer mailing list: lispflowmapping-dev@lists.opendaylight.org or on the #opendaylight-lispflowmapping IRC channel on irc.freenode.net.

Additional information is also available on the Lisp Flow Mapping wiki

Clustering in LISP Flow Mapping

Documentation regarding setting up a 3-node OpenDaylight cluster is described at following odl wiki page.

To turn on clustering in LISP Flow Mapping it is necessary:

- run script **deploy.py** script. This script is in integration-test project placed at *tools/clustering/cluster-deployer/deploy.py*. A whole deploy.py command can looks like:

{path_to_integration_test_project}/tools/clustering/cluster-deployer/**deploy.py**
–**distribution** {path_to_distribution_in_zip_format}

–**rootdir** {dir_at_remote_host_where_copy_odl_distribution}
–**hosts** {ip1},{ip2},{ip3}
–**clean**
–**template** lispflowmapping
–**rf** 3
–**user** {user_name_of_remote_hosts}
–**password** {password_to_remote_hosts}

Running this script will cause that specified **distribution** to be deployed to remote **hosts** specified through their IP adresses with using credentials (**user** and **password**). The distribution will be copied to specified **rootdir**. As part of the deployment, a **template** which contains a set of controller files which are different from standard ones. In this case it is specified in
{path_to_integration_test_project}/tools/clustering/cluster-deployer/lispflowmapping directory.
Lispflowmapping templates are part of integration-test project. There are 5 template files:

- akka.conf.template
- jolokia.xml.template
- module-shards.conf.template
- modules.conf.template
- org.apache.karaf.features.cfg.template

After copying the distribution, it is unzipped and started on all of specified **hosts** in cluster aware manner.

Remarks

It is necessary to have:

- **unzip** program installed on all of the host
- set all remote hosts /etc/sudoers files to not **requiretty** (should only matter on debian hosts)

NATApp User Guide

The NATApp User Guide contains information about configuration, administration, management, using and troubleshooting the feature.

Overview

NATApp provides network different types of address translation functionality for OpenDaylight. After installing this feature, network administrators can select the type of NAT functionality they want to enable by sending a REST API command. Subsequently, the user may enter the gloabl IP addresses to the YANG Data Store through REST APIs. When an OpenDaylight managed enterprise network with local IPs tries to connect to external networks such as Internet, NATApp comes into play and installs appropriate flow rules at the OpenFlow switch for bidirectional NAT translation.

NATApp Architecture

NATApp listens on OpenFlow southbound interface for Packet_In messages. The application parses the message for header information. If the received message has a local IP address the application installs rules on the OpenFlow switch for network address translation from local to global IP addresses. NATApp has NATPacketHandler class that implements the PacketProcessing interface to override the OnPacketReceived notification by which the application is notified of Packet_In messages.

Configuring NATApp

REST APIs are available at the following URI: http://localhost:8181/apidoc/explorer/index.html#!/natapp(2016-01-25)

Mininet Topology

```
sudo mn --mac --topo=single,10 --controller=remote,ip=127.0.0.1,port=6653
```

Install a flow to flood the ARP packets.

```
sh ovs-ofctl add-flow s1 dl_type=0x0806,actions=FLOOD
```

Check the flow for ARP Flooding

```
sh ovs-ofctl dump-flows s1
```

Administering or Managing NATApp

Static NAT and Dynamic NAT

First user has to select the type of NAT he wants by using the following URI:

POST URI http://localhost:8181/restconf/operations/natapp:nat-type

Sample Input {“natapp:input”: { “type:static”:”“}}

Sample Input {“natapp:input”: { “type:dynamic”:”“}}

Then user can inject the Global IPs using the following URI

PUT URI http://localhost:8181/restconf/config/natapp:staticNat/

Sample Input {“natapp:staticNat”: {“globalIP”:[“172.0.0.1/32”,”172.0.0.2/32”, “172.0.0.3/32”, “172.0.0.4/32”, “172.0.0.5/32”, “172.0.0.6/32”, “172.0.0.7/32”, “172.0.0.8/32”, “172.0.0.9/32”, “172.0.0.10/32”] }}

From mininet verify any pair of hosts can ping each other. The NATApp modifies the destination IP address of the ICMP Echo request with the global IP address. Check the mininet flows for this modification.

```
sh ovs-ofctl dump-flows s1
```

Port Address Translation (PAT)

User can select PAT by using the following URI.

POST URI http://localhost:8181/restconf/operations/natapp:nat-type

Sample Input {“natapp:input”: { “type:pat”:”“}}

Then user can inject the Global IPs using the following URI

PUT URI http://localhost:8181/restconf/config/natapp:patNat/

Sample Input {“natapp:patNat”: {“globalIP”:”172.0.0.1/32”}}

From Mininet use the command as xterm h1 h5. At h5 give the following commands

```
$ ip r add 172.0.0.1/32 dev h5-eth0
$ arp -s 172.0.0.1 00:00:00:00:00:01
$ nc -l 5000
```

At h1, Give the following command

```
$ echo "TCS" | nc -p 8000 10.0.0.5 5000
```

```
mininet> sh ovs-ofctl dump-flows s1
NXST_FLOW reply (xid=0x4):
 cookie=0x0, duration=811.272s, table=0, n_packets=5, n_bytes=342, idle_age=13,␣
↪priority=210,tcp,in_port=1,tp_src=8000 actions=mod_nw_src:172.0.0.1,mod_tp_src:2000,
↪output:5
 cookie=0x0, duration=499.843s, table=0, n_packets=2, n_bytes=84, idle_age=13, arp␣
↪actions=FLOOD
 cookie=0x0, duration=811.203s, table=0, n_packets=3, n_bytes=206, idle_age=13,␣
↪priority=209,tcp,in_port=5,tp_dst=2000 actions=mod_nw_dst:10.0.0.1,mod_tp_dst:8000,
↪output:1
```

NEtwork MOdeling (NEMO)

This section describes how to use the NEMO feature in OpenDaylight and contains contains configuration, administration, and management sections for the feature.

Overview

With the network becoming more complicated, users and applications must handle more complex configurations to deploy new services. NEMO project aims to simplify the usage of network by providing a new intent northbound interface (NBI). Instead of tons of APIs, users/applications just need to describe their intent without caring about complex physical devices and implementation means. The intent will be translated into detailed configurations on the devices in the NEMO engine. A typical scenario is user just need to assign which nodes to implement an VPN, without considering which technique is used.

NEMO Engine Architecture

- NEMO API * The NEMO API provide users the NEMO model, which guides users how to construct the instance of intent, and how to construct the instance of predefined types.
- NEMO REST * The NEMO REST provides users REST APIs to access NEMO engine, that is, user could transmit the intent instance to NEMO engine through basic REST methods.
- NEMO UI * The NEMO UI provides users a visual interface to deploy service with NEMO model, and display the state in DLUX UI.

Installing NEMO engine

To install NEMO engine, download OpenDaylight and use the Karaf console to install the following feature:

odl-nemo-engine-ui

Administering or Managing NEMO Engine

After install features NEMO engine used, user could use NEMO to express his intent with NEMO UI or REST APIs in apidoc.

Go to `http://{controller-ip}:8181/index.html`. In this interface, user could go to NEMO UI, and use the tabs and input box to input intent, and see the state of intent deployment with the image.

Go to `http://{controller-ip}:8181/apidoc/explorer/index.html`. In this interface, user could REST methods "POST", "PUT","GET" and "DELETE" to deploy intent or query the state of deployment.

Tutorials

Below are tutorials for NEMO Engine.

Using NEMO Engine

The purpose of the tutorial is to describe how to use use UI to deploy intent.

Overview

This tutorial will describe how to use the NEMO UI to check the operated resources, the steps to deploy service, and the ultimate state.

Prerequisites

To understand the tutorial well, we hope there are a physical or virtual network exist, and OpenDaylight with NEMO engine must be deployed in one host.

Target Environment

The intent expressed by NEMO model is depended on network resources, so user need to have enough resources to use, or else, the deployment of intent will fail.

Instructions

- Run the OpenDaylight distribution and install odl-nemo-engine-ui from the Karaf console.
- Go to `http://{controller-ip}:8181/index.html`, and sign in.
- Go the NEMO UI interface. And Register a new user with user name, password, and tenant.
- Check the existing resources to see if it is consistent with yours.

- Deploy service with NEMO model by the create intent menu.

NETCONF User Guide

Overview

NETCONF is an XML-based protocol used for configuration and monitoring devices in the network. The base NETCONF protocol is described in RFC-6241.

NETCONF in OpenDaylight:.

OpenDaylight supports the NETCONF protocol as a northbound server as well as a southbound plugin. It also includes a set of test tools for simulating NETCONF devices and clients.

Southbound (netconf-connector)

The NETCONF southbound plugin is capable of connecting to remote NETCONF devices and exposing their configuration/operational datastores, RPCs and notifications as MD-SAL mount points. These mount points allow applications and remote users (over RESTCONF) to interact with the mounted devices.

In terms of RFCs, the connector supports:

- RFC-6241
- RFC-5277
- RFC-6022
- draft-ietf-netconf-yang-library-06

Netconf-connector is fully model-driven (utilizing the YANG modeling language) so in addition to the above RFCs, it supports any data/RPC/notifications described by a YANG model that is implemented by the device.

Tip: NETCONF southbound can be activated by installing `odl-netconf-connector-all` Karaf feature.

Netconf-connector configuration

There are 2 ways for configuring netconf-connector: NETCONF or RESTCONF. This guide focuses on using RESTCONF.

Default configuration

The default configuration contains all the necessary dependencies (file: 01-netconf.xml) and a single instance of netconf-connector (file: 99-netconf-connector.xml) called **controller-config** which connects itself to the NETCONF northbound in OpenDaylight in a loopback fashion. The connector mounts the NETCONF server for config-subsystem in order to enable RESTCONF protocol for config-subsystem. This RESTCONF still goes via NETCONF, but using RESTCONF is much more user friendly than using NETCONF.

Spawning additional netconf-connectors while the controller is running

Preconditions:

1. OpenDaylight is running
2. In Karaf, you must have the netconf-connector installed (at the Karaf prompt, type: `feature:install odl-netconf-connector-all`); the loopback NETCONF mountpoint will be automatically configured and activated
3. Wait until log displays following entry: RemoteDevice{controller-config}: NETCONF connector initialized successfully

To configure a new netconf-connector you need to send following request to RESTCONF:

POST http://localhost:8181/restconf/config/network-topology:network-topology/topology/topology-netconf/node/controller-config/yang-ext:mount/config:modules

Headers:

- Accept application/xml
- Content-Type application/xml

```
<module xmlns="urn:opendaylight:params:xml:ns:yang:controller:config">
  <type xmlns:prefix=
↪"urn:opendaylight:params:xml:ns:yang:controller:md:sal:connector:netconf">
↪prefix:sal-netconf-connector</type>
  <name>new-netconf-device</name>
  <address xmlns=
↪"urn:opendaylight:params:xml:ns:yang:controller:md:sal:connector:netconf">127.0.0.1
↪</address>
  <port xmlns="urn:opendaylight:params:xml:ns:yang:controller:md:sal:connector:netconf
↪">830</port>
  <username xmlns=
↪"urn:opendaylight:params:xml:ns:yang:controller:md:sal:connector:netconf">admin</
↪username>
  <password xmlns=
↪"urn:opendaylight:params:xml:ns:yang:controller:md:sal:connector:netconf">admin</
↪password>
  <tcp-only xmlns=
↪"urn:opendaylight:params:xml:ns:yang:controller:md:sal:connector:netconf">false</
↪tcp-only>
  <event-executor xmlns=
↪"urn:opendaylight:params:xml:ns:yang:controller:md:sal:connector:netconf">
    <type xmlns:prefix="urn:opendaylight:params:xml:ns:yang:controller:netty">
↪prefix:netty-event-executor</type>
    <name>global-event-executor</name>
  </event-executor>
  <binding-registry xmlns=
↪"urn:opendaylight:params:xml:ns:yang:controller:md:sal:connector:netconf">
    <type xmlns:prefix="urn:opendaylight:params:xml:ns:yang:controller:md:sal:binding
↪">prefix:binding-broker-osgi-registry</type>
    <name>binding-osgi-broker</name>
  </binding-registry>
  <dom-registry xmlns=
↪"urn:opendaylight:params:xml:ns:yang:controller:md:sal:connector:netconf">
    <type xmlns:prefix="urn:opendaylight:params:xml:ns:yang:controller:md:sal:dom">
↪prefix:dom-broker-osgi-registry</type>
    <name>dom-broker</name>
```

(continues on next page)

(continued from previous page)

```
  </dom-registry>
  <client-dispatcher xmlns=
↪"urn:opendaylight:params:xml:ns:yang:controller:md:sal:connector:netconf">
    <type xmlns:prefix="urn:opendaylight:params:xml:ns:yang:controller:config:netconf
↪">prefix:netconf-client-dispatcher</type>
    <name>global-netconf-dispatcher</name>
  </client-dispatcher>
  <processing-executor xmlns=
↪"urn:opendaylight:params:xml:ns:yang:controller:md:sal:connector:netconf">
    <type xmlns:prefix="urn:opendaylight:params:xml:ns:yang:controller:threadpool">
↪prefix:threadpool</type>
    <name>global-netconf-processing-executor</name>
  </processing-executor>
  <keepalive-executor xmlns=
↪"urn:opendaylight:params:xml:ns:yang:controller:md:sal:connector:netconf">
    <type xmlns:prefix="urn:opendaylight:params:xml:ns:yang:controller:threadpool">
↪prefix:scheduled-threadpool</type>
    <name>global-netconf-ssh-scheduled-executor</name>
  </keepalive-executor>
</module>
```

This spawns a new netconf-connector which tries to connect to (or mount) a NETCONF device at 127.0.0.1 and port 830. You can check the configuration of config-subsystem's configuration datastore. The new netconf-connector will now be present there. Just invoke:

GET http://localhost:8181/restconf/config/network-topology:network-topology/topology/topology-netconf/node/controller-config/yang-ext:mount/config:modules

The response will contain the module for new-netconf-device.

Right after the new netconf-connector is created, it writes some useful metadata into the datastore of MD-SAL under the network-topology subtree. This metadata can be found at:

GET http://localhost:8181/restconf/operational/network-topology:network-topology/

Information about connection status, device capabilities, etc. can be found there.

Connecting to a device not supporting NETCONF monitoring

The netconf-connector in OpenDaylight relies on ietf-netconf-monitoring support when connecting to remote NETCONF device. The ietf-netconf-monitoring support allows netconf-connector to list and download all YANG schemas that are used by the device. NETCONF connector can only communicate with a device if it knows the set of used schemas (or at least a subset). However, some devices use YANG models internally but do not support NETCONF monitoring. Netconf-connector can also communicate with these devices, but you have to side load the necessary yang models into OpenDaylight's YANG model cache for netconf-connector. In general there are 2 situations you might encounter:

1. NETCONF device does not support ietf-netconf-monitoring but it does list all its YANG models as capabilities in HELLO message

This could be a device that internally uses only ietf-inet-types YANG model with revision 2010-09-24. In the HELLO message that is sent from this device there is this capability reported:

```
urn:ietf:params:xml:ns:yang:ietf-inet-types?module=ietf-inet-types&revision=2010-09-24
```

For such devices you only need to put the schema into folder cache/schema inside your Karaf distribution.

Important: The file with YANG schema for ietf-inet-types has to be called ietf-inet-types@2010-09-24.yang. It is the required naming format of the cache.

2. NETCONF device does not support ietf-netconf-monitoring and it does NOT list its YANG models as capabilities in HELLO message

Compared to device that lists its YANG models in HELLO message, in this case there would be no capability with ietf-inet-types in the HELLO message. This type of device basically provides no information about the YANG schemas it uses so its up to the user of OpenDaylight to properly configure netconf-connector for this device.

Netconf-connector has an optional configuration attribute called yang-module-capabilities and this attribute can contain a list of "YANG module based" capabilities. So by setting this configuration attribute, it is possible to override the "yang-module-based" capabilities reported in HELLO message of the device. To do this, we need to modify the configuration of netconf-connector by adding this XML (It needs to be added next to the address, port, username etc. configuration elements):

```
<yang-module-capabilities xmlns=
↪"urn:opendaylight:params:xml:ns:yang:controller:md:sal:connector:netconf">
  <capability xmlns=
↪"urn:opendaylight:params:xml:ns:yang:controller:md:sal:connector:netconf">
    urn:ietf:params:xml:ns:yang:ietf-inet-types?module=ietf-inet-types&amp;
↪revision=2010-09-24
  </capability>
</yang-module-capabilities>
```

Remember to also put the YANG schemas into the cache folder.

Note: For putting multiple capabilities, you just need to replicate the capability xml element inside yang-module-capability element. Capability element is modeled as a leaf-list. With this configuration, we would make the remote device report usage of ietf-inet-types in the eyes of netconf-connector.

Reconfiguring Netconf-Connector While the Controller is Running

It is possible to change the configuration of a running module while the whole controller is running. This example will continue where the last left off and will change the configuration for the brand new netconf-connector after it was spawned. Using one RESTCONF request, we will change both username and password for the netconf-connector.

To update an existing netconf-connector you need to send following request to RESTCONF:

PUT http://localhost:8181/restconf/config/network-topology:network-topology/topology/topology-netconf/node/controller-config/yang-ext:mount/config:modules/module/odl-sal-netconf-connector-cfg:sal-netconf-connector/new-netconf-device

```
<module xmlns="urn:opendaylight:params:xml:ns:yang:controller:config">
  <type xmlns:prefix=
↪"urn:opendaylight:params:xml:ns:yang:controller:md:sal:connector:netconf">
↪prefix:sal-netconf-connector</type>
  <name>new-netconf-device</name>
  <username xmlns=
↪"urn:opendaylight:params:xml:ns:yang:controller:md:sal:connector:netconf">bob</
↪username>
  <password xmlns=
↪"urn:opendaylight:params:xml:ns:yang:controller:md:sal:connector:netconf">passwd</
↪password>
```

(continues on next page)

(continued from previous page)

```
  <tcp-only xmlns=
↪"urn:opendaylight:params:xml:ns:yang:controller:md:sal:connector:netconf">false</
↪tcp-only>
  <event-executor xmlns=
↪"urn:opendaylight:params:xml:ns:yang:controller:md:sal:connector:netconf">
    <type xmlns:prefix="urn:opendaylight:params:xml:ns:yang:controller:netty">
↪prefix:netty-event-executor</type>
    <name>global-event-executor</name>
  </event-executor>
  <binding-registry xmlns=
↪"urn:opendaylight:params:xml:ns:yang:controller:md:sal:connector:netconf">
    <type xmlns:prefix="urn:opendaylight:params:xml:ns:yang:controller:md:sal:binding
↪">prefix:binding-broker-osgi-registry</type>
    <name>binding-osgi-broker</name>
  </binding-registry>
  <dom-registry xmlns=
↪"urn:opendaylight:params:xml:ns:yang:controller:md:sal:connector:netconf">
    <type xmlns:prefix="urn:opendaylight:params:xml:ns:yang:controller:md:sal:dom">
↪prefix:dom-broker-osgi-registry</type>
    <name>dom-broker</name>
  </dom-registry>
  <client-dispatcher xmlns=
↪"urn:opendaylight:params:xml:ns:yang:controller:md:sal:connector:netconf">
    <type xmlns:prefix="urn:opendaylight:params:xml:ns:yang:controller:config:netconf
↪">prefix:netconf-client-dispatcher</type>
    <name>global-netconf-dispatcher</name>
  </client-dispatcher>
  <processing-executor xmlns=
↪"urn:opendaylight:params:xml:ns:yang:controller:md:sal:connector:netconf">
    <type xmlns:prefix="urn:opendaylight:params:xml:ns:yang:controller:threadpool">
↪prefix:threadpool</type>
    <name>global-netconf-processing-executor</name>
  </processing-executor>
  <keepalive-executor xmlns=
↪"urn:opendaylight:params:xml:ns:yang:controller:md:sal:connector:netconf">
    <type xmlns:prefix="urn:opendaylight:params:xml:ns:yang:controller:threadpool">
↪prefix:scheduled-threadpool</type>
    <name>global-netconf-ssh-scheduled-executor</name>
  </keepalive-executor>
</module>
```

Since a PUT is a replace operation, the whole configuration must be specified along with the new values for username and password. This should result in a 2xx response and the instance of netconf-connector called new-netconf-device will be reconfigured to use username bob and password passwd. New configuration can be verified by executing:

GET http://localhost:8181/restconf/config/network-topology:network-topology/topology/topology-netconf/node/controller-config/yang-ext:mount/config:modules/module/odl-sal-netconf-connector-cfg:sal-netconf-connector/new-netconf-device

With new configuration, the old connection will be closed and a new one established.

Destroying Netconf-Connector While the Controller is Running

Using RESTCONF one can also destroy an instance of a module. In case of netconf-connector, the module will be destroyed, NETCONF connection dropped and all resources will be cleaned. To do this, simply issue a request to following URL:

DELETE http://localhost:8181/restconf/config/network-topology:network-topology/topology/topology-netconf/node/controller-config/yang-ext:mount/config:modules/module/odl-sal-netconf-connector-cfg:sal-netconf-connector/new-netconf-device

The last element of the URL is the name of the instance and its predecessor is the type of that module (In our case the type is **sal-netconf-connector** and name **new-netconf-device**). The type and name are actually the keys of the module list.

Netconf-connector configuration with MD-SAL

It is also possible to configure new NETCONF connectors directly through MD-SAL with the usage of the network-topology model. You can configure new NETCONF connectors both through the NETCONF server for MD-SAL (port 2830) or RESTCONF. This guide focuses on RESTCONF.

Tip: To enable NETCONF connector configuration through MD-SAL install either the `odl-netconf-topology` or `odl-netconf-clustered-topology` feature. We will explain the difference between these features later.

Preconditions

1. OpenDaylight is running
2. In Karaf, you must have the `odl-netconf-topology` or `odl-netconf-clustered-topology` feature installed.
3. Feature `odl-restconf` must be installed
4. Wait until log displays following entry:

   ```
   Successfully pushed configuration snapshot 02-netconf-topology.xml(odl-netconf-
   ↪topology,odl-netconf-topology)
   ```

   or until

   ```
   GET http://localhost:8181/restconf/operational/network-topology:network-topology/
   ↪topology/topology-netconf/
   ```

   returns a non-empty response, for example:

   ```
   <topology xmlns="urn:TBD:params:xml:ns:yang:network-topology">
     <topology-id>topology-netconf</topology-id>
   </topology>
   ```

Spawning new NETCONF connectors

To create a new NETCONF connector you need to send the following request to RESTCONF:

```
PUT http://localhost:8181/restconf/config/network-topology:network-topology/topology/
↪topology-netconf/node/new-netconf-device
```

Headers:

- Accept: application/xml
- Content-Type: application/xml

Payload:

```
<node xmlns="urn:TBD:params:xml:ns:yang:network-topology">
  <node-id>new-netconf-device</node-id>
  <host xmlns="urn:opendaylight:netconf-node-topology">127.0.0.1</host>
  <port xmlns="urn:opendaylight:netconf-node-topology">17830</port>
  <username xmlns="urn:opendaylight:netconf-node-topology">admin</username>
  <password xmlns="urn:opendaylight:netconf-node-topology">admin</password>
  <tcp-only xmlns="urn:opendaylight:netconf-node-topology">false</tcp-only>
  <!-- non-mandatory fields with default values, you can safely remove these if you␣
↪do not wish to override any of these values-->
  <reconnect-on-changed-schema xmlns="urn:opendaylight:netconf-node-topology">false</
↪reconnect-on-changed-schema>
  <connection-timeout-millis xmlns="urn:opendaylight:netconf-node-topology">20000</
↪connection-timeout-millis>
  <max-connection-attempts xmlns="urn:opendaylight:netconf-node-topology">0</max-
↪connection-attempts>
  <between-attempts-timeout-millis xmlns="urn:opendaylight:netconf-node-topology">2000
↪</between-attempts-timeout-millis>
  <sleep-factor xmlns="urn:opendaylight:netconf-node-topology">1.5</sleep-factor>
  <!-- keepalive-delay set to 0 turns off keepalives-->
  <keepalive-delay xmlns="urn:opendaylight:netconf-node-topology">120</keepalive-
↪delay>
</node>
```

Note that the device name in <node-id> element must match the last element of the restconf URL.

Reconfiguring an existing connector

The steps to reconfigure an existing connector are exactly the same as when spawning a new connector. The old connection will be disconnected and a new connector with the new configuration will be created.

Deleting an existing connector

To remove an already configured NETCONF connector you need to send the following:

```
DELETE http://localhost:8181/restconf/config/network-topology:network-topology/
↪topology/topology-netconf/node/new-netconf-device
```

Connecting to a device supporting only NETCONF 1.0

OpenDaylight is schema-based distribution and heavily depends on YANG models. However some legacy NETCONF devices are not schema-based and implement just RFC 4741. This type of device does not utilize YANG models internally and OpenDaylight does not know how to communicate with such devices, how to validate data, or what the semantics of data are.

NETCONF connector can communicate also with these devices, but the trade-offs are worsened possibilities in utilization of NETCONF mountpoints. Using RESTCONF with such devices is not suported. Also communicating with schemaless devices from application code is slightly different.

To connect to schemaless device, there is a optional configuration option in netconf-node-topology model called schemaless. You have to set this option to true.

Clustered NETCONF connector

To spawn NETCONF connectors that are cluster-aware you need to install the `odl-netconf-clustered-topology` karaf feature.

> **Warning:** The `odl-netconf-topology` and `odl-netconf-clustered-topology` features are considered **INCOMPATIBLE**. They both manage the same space in the datastore and would issue conflicting writes if installed together.

Configuration of clustered NETCONF connectors works the same as the configuration through the topology model in the previous section.

When a new clustered connector is configured the configuration gets distributed among the member nodes and a NETCONF connector is spawned on each node. From these nodes a master is chosen which handles the schema download from the device and all the communication with the device. You will be able to read/write to/from the device from all slave nodes due to the proxy data brokers implemented.

You can use the `odl-netconf-clustered-topology` feature in a single node scenario as well but the code that uses akka will be used, so for a scenario where only a single node is used, `odl-netconf-topology` might be preferred.

Netconf-connector utilization

Once the connector is up and running, users can utilize the new mount point instance. By using RESTCONF or from their application code. This chapter deals with using RESTCONF and more information for app developers can be found in the developers guide or in the official tutorial application **ncmount** that can be found in the coretutorials project:

- https://github.com/opendaylight/coretutorials/tree/stable/beryllum/ncmount

Reading data from the device

Just invoke (no body needed):

GET http://localhost:8080/restconf/operational/network-topology:network-topology/topology/topology-netconf/node/new-netconf-device/yang-ext:mount/

This will return the entire content of operation datastore from the device. To view just the configuration datastore, change **operational** in this URL to **config**.

Writing configuration data to the device

In general, you cannot simply write any data you want to the device. The data have to conform to the YANG models implemented by the device. In this example we are adding a new interface-configuration to the mounted device (assuming the device supports Cisco-IOS-XR-ifmgr-cfg YANG model). In fact this request comes from the tutorial dedicated to the **ncmount** tutorial app.

POST http://localhost:8181/restconf/config/network-topology:network-topology/topology/topology-netconf/node/new-netconf-device/yang-ext:mount/Cisco-IOS-XR-ifmgr-cfg:interface-configurations

```
<interface-configuration xmlns="http://cisco.com/ns/yang/Cisco-IOS-XR-ifmgr-cfg">
    <active>act</active>
    <interface-name>mpls</interface-name>
    <description>Interface description</description>
    <bandwidth>32</bandwidth>
    <link-status></link-status>
</interface-configuration>
```

Should return 200 response code with no body.

Tip: This call is transformed into a couple of NETCONF RPCs. Resulting NETCONF RPCs that go directly to the device can be found in the OpenDaylight logs after invoking `log:set TRACE org.opendaylight.controller.sal.connect.netconf` in the Karaf shell. Seeing the NETCONF RPCs might help with debugging.

This request is very similar to the one where we spawned a new netconf device. That's because we used the loopback netconf-connector to write configuration data into config-subsystem datastore and config-subsystem picked it up from there.

Invoking custom RPC

Devices can implement any additional RPC and as long as it provides YANG models for it, it can be invoked from OpenDaylight. Following example shows how to invoke the get-schema RPC (get-schema is quite common among netconf devices). Invoke:

POST http://localhost:8181/restconf/operations/network-topology:network-topology/topology/topology-netconf/node/new-netconf-device/yang-ext:mount/ietf-netconf-monitoring:get-schema

```
<input xmlns="urn:ietf:params:xml:ns:yang:ietf-netconf-monitoring">
  <identifier>ietf-yang-types</identifier>
  <version>2013-07-15</version>
</input>
```

This call should fetch the source for ietf-yang-types YANG model from the mounted device.

Netconf-connector + Netopeer

Netopeer (an open-source NETCONF server) can be used for testing/exploring NETCONF southbound in OpenDaylight.

Netopeer installation

A Docker container with netopeer will be used in this guide. To install Docker and start the netopeer image perform following steps:

1. Install docker http://docs.docker.com/linux/step_one/
2. Start the netopeer image:

   ```
   docker run -rm -t -p 1831:830 dockeruser/netopeer
   ```

3. Verify netopeer is running by invoking (netopeer should send its HELLO message right away:

```
ssh root@localhost -p 1831 -s netconf
(password root)
```

Mounting netopeer NETCONF server

Preconditions:

- OpenDaylight is started with features `odl-restconf-all` and `odl-netconf-connector-all`.
- Netopeer is up and running in docker

Now just follow the chapter: *Spawning netconf-connector*. In the payload change the:

- name, e.g., to netopeer
- username/password to your system credentials
- ip to localhost
- port to 1831.

After netopeer is mounted successfully, its configuration can be read using RESTCONF by invoking:

GET http://localhost:8181/restconf/config/network-topology:network-topology/topology/topology-netconf/node/netopeer/yang-ext:mount/

Northbound (NETCONF servers)

OpenDaylight provides 2 types of NETCONF servers:

- **NETCONF server for config-subsystem (listening by default on port 1830)**
  - Serves as a default interface for config-subsystem and allows users to spawn/reconfigure/destroy modules (or applications) in OpenDaylight
- **NETCONF server for MD-SAL (listening by default on port 2830)**
  - Serves as an alternative interface for MD-SAL (besides RESTCONF) and allows users to read/write data from MD-SAL's datastore and to invoke its rpcs (NETCONF notifications are not available in the Boron release of OpenDaylight)

Note: The reason for having 2 NETCONF servers is that config-subsystem and MD-SAL are 2 different components of OpenDaylight and require different approach for NETCONF message handling and data translation. These 2 components will probably merge in the future.

NETCONF server for config-subsystem

This NETCONF server is the primary interface for config-subsystem. It allows the users to interact with config-subsystem in a standardized NETCONF manner.

In terms of RFCs, these are supported:

- RFC-6241
- RFC-5277
- RFC-6470

  - (partially, only the schema-change notification is available in Boron release)
- RFC-6022

For regular users it is recommended to use RESTCONF + the controller-config loopback mountpoint instead of using pure NETCONF. How to do that is spesific for each component/module/application in OpenDaylight and can be found in their dedicated user guides.

NETCONF server for MD-SAL

This NETCONF server is just a generic interface to MD-SAL in OpenDaylight. It uses the stadard MD-SAL APIs and serves as an alternative to RESTCONF. It is fully model driven and supports any data and rpcs that are supported by MD-SAL.

In terms of RFCs, these are supported:

- RFC-6241
- RFC-6022
- draft-ietf-netconf-yang-library-06

Notifications over NETCONF are not supported in the Boron release.

Tip: Install NETCONF northbound for MD-SAL by installing feature: `odl-netconf-mdsal` in karaf. Default binding port is **2830**.

Configuration

The default configuration can be found in file: *08-netconf-mdsal.xml*. The file contains the configuration for all necessary dependencies and a single SSH endpoint starting on port 2830. There is also a (by default disabled) TCP endpoint. It is possible to start multiple endpoints at the same time either in the initial configuration file or while OpenDaylight is running.

The credentials for SSH endpoint can also be configured here, the defaults are admin/admin. Credentials in the SSH endpoint are not yet managed by the centralized AAA component and have to be configured separately.

Verifying MD-SAL's NETCONF server

After the NETCONF server is available it can be examined by a command line ssh tool:

```
ssh admin@localhost -p 2830 -s netconf
```

The server will respond by sending its HELLO message and can be used as a regular NETCONF server from then on.

Mounting the MD-SAL's NETCONF server

To perform this operation, just spawn a new netconf-connector as described in *Spawning netconf-connector*. Just change the ip to "127.0.0.1" port to "2830" and its name to "controller-mdsal".

Now the MD-SAL's datastore can be read over RESTCONF via NETCONF by invoking:

GET http://localhost:8181/restconf/operational/network-topology:network-topology/topology/topology-netconf/node/controller-mdsal/yang-ext:mount

Note: This might not seem very useful, since MD-SAL can be accessed directly from RESTCONF or from Application code, but the same method can be used to mount and control other OpenDaylight instances by the "master OpenDaylight".

NETCONF testtool

NETCONF testtool is a set of standalone runnable jars that can:

- Simulate NETCONF devices (suitable for scale testing)
- Stress/Performance test NETCONF devices
- Stress/Performance test RESTCONF devices

These jars are part of OpenDaylight's controller project and are built from the NETCONF codebase in OpenDaylight.

Tip: Download testtool from OpenDaylight Nexus at: https://nexus.opendaylight.org/content/repositories/public/org/opendaylight/netconf/netconf-testtool/1.1.0-Boron/

Nexus contains 3 executable tools:

- executable.jar - device simulator
- stress.client.tar.gz - NETCONF stress/performance measuring tool
- perf-client.jar - RESTCONF stress/performance measuring tool

Tip: Each executable tool provides help. Just invoke `java -jar <name-of-the-tool.jar> --help`

NETCONF device simulator

NETCONF testtool (or NETCONF device simulator) is a tool that

- Simulates 1 or more NETCONF devices
- Is suitable for scale, performance or crud testing
- Uses core implementation of NETCONF server from OpenDaylight
- Generates configuration files for controller so that the OpenDaylight distribution (Karaf) can easily connect to all simulated devices
- Provides broad configuration options
- Can start a fully fledged MD-SAL datastore
- Supports notifications

Building testtool

1. Check out latest NETCONF repository from git
2. Move into the `opendaylight/netconf/tools/netconf-testtool/` folder

3. Build testtool using the `mvn clean install` command

Downloading testtool

Netconf-testtool is now part of default maven build profile for controller and can be also downloaded from nexus. The executable jar for testtool can be found at: nexus-artifacts

Running testtool

1. After successfully building or downloading, move into the `opendaylight/netconf/tools/netconf-testtool/target/` folder and there is file `netconf-testtool-1.1.0-SNAPSHOT-executable.jar` (or if downloaded from nexus just take that jar file)
2. Execute this file using, e.g.:

```
java -jar netconf-testtool-1.1.0-SNAPSHOT-executable.jar
```

This execution runs the testtool with default for all parameters and you should see this log output from the testtool :

```
10:31:08.206 [main] INFO  o.o.c.n.t.t.NetconfDeviceSimulator - Starting 1, SSH␣
↪simulated devices starting on port 17830
10:31:08.675 [main] INFO  o.o.c.n.t.t.NetconfDeviceSimulator - All simulated␣
↪devices started successfully from port 17830 to 17830
```

Default Parameters

The default parameters for testtool are:

- Use SSH
- Run 1 simulated device
- Device port is 17830
- YANG modules used by device are only: ietf-netconf-monitoring, ietf-yang-types, ietf-inet-types (these modules are required for device in order to support NETCONF monitoring and are included in the netconf-testtool)
- Connection timeout is set to 30 minutes (quite high, but when testing with 10000 devices it might take some time for all of them to fully establish a connection)
- Debug level is set to false
- No distribution is modified to connect automatically to the NETCONF testtool

Verifying testtool

To verify that the simulated device is up and running, we can try to connect to it using command line ssh tool. Execute this command to connect to the device:

```
ssh admin@localhost -p 17830 -s netconf
```

Just accept the server with yes (if required) and provide any password (testtool accepts all users with all passwords). You should see the hello message sent by simulated device.

Testtool help

```
usage: netconf testool [-h] [--device-count DEVICES-COUNT] [--devices-per-port␣
↪DEVICES-PER-PORT] [--schemas-dir SCHEMAS-DIR] [--notification-file NOTIFICATION-
↪FILE]
                       [--initial-config-xml-file INITIAL-CONFIG-XML-FILE] [--
↪starting-port STARTING-PORT] [--generate-config-connection-timeout GENERATE-CONFIG-
↪CONNECTION-TIMEOUT]
                       [--generate-config-address GENERATE-CONFIG-ADDRESS] [--
↪generate-configs-batch-size GENERATE-CONFIGS-BATCH-SIZE] [--distribution-folder␣
↪DISTRO-FOLDER] [--ssh SSH] [--exi EXI]
                       [--debug DEBUG] [--md-sal MD-SAL]

NETCONF device simulator. Detailed info can be found at https://wiki.opendaylight.org/
↪view/OpenDaylight_Controller:Netconf:Testtool#Building_testtool

optional arguments:
  -h, --help            show this help message and exit
  --device-count DEVICES-COUNT
                        Number of simulated netconf devices to spin. This is the␣
↪number of actual ports open for the devices.
  --devices-per-port DEVICES-PER-PORT
                        Amount of config files generated per port to spoof more␣
↪devices then are actually running
  --schemas-dir SCHEMAS-DIR
                        Directory containing yang schemas to describe simulated␣
↪devices. Some schemas e.g. netconf monitoring and inet types are included by default
  --notification-file NOTIFICATION-FILE
                        Xml file containing notifications that should be sent to␣
↪clients after create subscription is called
  --initial-config-xml-file INITIAL-CONFIG-XML-FILE
                        Xml file containing initial simulatted configuration to be␣
↪returned via get-config rpc
  --starting-port STARTING-PORT
                        First port for simulated device. Each other device will have␣
↪previous+1 port number
  --generate-config-connection-timeout GENERATE-CONFIG-CONNECTION-TIMEOUT
                        Timeout to be generated in initial config files
  --generate-config-address GENERATE-CONFIG-ADDRESS
                        Address to be placed in generated configs
  --generate-configs-batch-size GENERATE-CONFIGS-BATCH-SIZE
                        Number of connector configs per generated file
  --distribution-folder DISTRO-FOLDER
                        Directory where the karaf distribution for controller is␣
↪located
  --ssh SSH             Whether to use ssh for transport or just pure tcp
  --exi EXI             Whether to use exi to transport xml content
  --debug DEBUG         Whether to use debug log level instead of INFO
  --md-sal MD-SAL       Whether to use md-sal datastore instead of default simulated␣
↪datastore.
```

Supported operations

Testtool default simple datastore supported operations:

get-schema returns YANG schemas loaded from user specified directory,

edit-config always returns OK and stores the XML from the input in a local variable available for get-config and get RPC. Every edit-config replaces the previous data,

commit always returns OK, but does not actually commit the data,

get-config returns local XML stored by edit-config,

get returns local XML stored by edit-config with netconf-state subtree, but also supports filtering.

(un)lock returns always OK with no lock guarantee

create-subscription returns always OK and after the operation is triggered, provided NETCONF notifications (if any) are fed to the client. No filtering or stream recognition is supported.

Note: when operation="delete" is present in the payload for edit-config, it will wipe its local store to simulate the removal of data.

When using the MD-SAL datastore testtool behaves more like normal NETCONF server and is suitable for crud testing. create-subscription is not supported when testtool is running with the MD-SAL datastore.

Notification support

Testtool supports notifications via the –notification-file switch. To trigger the notification feed, create-subscription operation has to be invoked. The XML file provided should look like this example file:

```
<?xml version='1.0' encoding='UTF-8' standalone='yes'?>
<notifications>

<!-- Notifications are processed in the order they are defined in XML -->

<!-- Notification that is sent only once right after create-subscription is called -->
<notification>
    <!-- Content of each notification entry must contain the entire notification with␣
↪event time. Event time can be hardcoded, or generated by testtool if XXXX is set as␣
↪eventtime in this XML -->
    <content><![CDATA[
        <notification xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0">
            <eventTime>2011-01-04T12:30:46</eventTime>
            <random-notification xmlns="http://www.opendaylight.org/netconf/event:1.0
↪">
                <random-content>single no delay</random-content>
            </random-notification>
        </notification>
    ]]></content>
</notification>

<!-- Repeated Notification that is sent 5 times with 2 second delay inbetween -->
<notification>
    <!-- Delay in seconds from previous notification -->
    <delay>2</delay>
    <!-- Number of times this notification should be repeated -->
    <times>5</times>
    <content><![CDATA[
        <notification xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0">
            <eventTime>XXXX</eventTime>
            <random-notification xmlns="http://www.opendaylight.org/netconf/event:1.0
↪">
                <random-content>scheduled 5 times 10 seconds each</random-content>
            </random-notification>
```

(continues on next page)

(continued from previous page)

```
        </notification>
    ]]></content>
</notification>

<!-- Single notification that is sent only once right after the previous notification␣
↪-->
<notification>
    <delay>2</delay>
    <content><![CDATA[
        <notification xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0">
            <eventTime>XXXX</eventTime>
            <random-notification xmlns="http://www.opendaylight.org/netconf/event:1.0
↪">
                <random-content>single with delay</random-content>
            </random-notification>
        </notification>
    ]]></content>
</notification>

</notifications>
```

Connecting testtool with controller Karaf distribution

Auto connect to OpenDaylight

It is possible to make OpenDaylight auto connect to the simulated devices spawned by testtool (so user does not have to post a configuration for every NETCONF connector via RESTCONF). The testtool is able to modify the OpenDaylight distribution to auto connect to the simulated devices after feature `odl-netconf-connector-all` is installed. When running testtool, issue this command (just point the testool to the distribution:

```
java -jar netconf-testtool-1.1.0-SNAPSHOT-executable.jar --device-count 10 --
↪distribution-folder ~/distribution-karaf-0.4.0-SNAPSHOT/ --debug true
```

With the distribution-folder parameter, the testtool will modify the distribution to include configuration for netconf-connector to connect to all simulated devices. So there is no need to spawn netconf-connectors via RESTCONF.

Running testtool and OpenDaylight on different machines

The testtool binds by default to 0.0.0.0 so it should be accessible from remote machines. However you need to set the parameter "generate-config-address" (when using autoconnect) to the address of machine where testtool will be run so OpenDaylight can connect. The default value is localhost.

Executing operations via RESTCONF on a mounted simulated device

Simulated devices support basic RPCs for editing their config. This part shows how to edit data for simulated device via RESTCONF.

Test YANG schema

The controller and RESTCONF assume that the data that can be manipulated for mounted device is described by a YANG schema. For demonstration, we will define a simple YANG model:

```
module test {
    yang-version 1;
    namespace "urn:opendaylight:test";
    prefix "tt";

    revision "2014-10-17";


   container cont {

        leaf l {
            type string;
        }
   }
}
```

Save this schema in file called test@2014-10-17.yang and store it a directory called test-schemas/, e.g., your home folder.

Editing data for simulated device

- Start the device with following command:

```
java -jar netconf-testtool-1.1.0-SNAPSHOT-executable.jar --device-count 10 --
↪distribution-folder ~/distribution-karaf-0.4.0-SNAPSHOT/ --debug true --schemas-
↪dir ~/test-schemas/
```

- Start OpenDaylight
- Install odl-netconf-connector-all feature
- Install odl-restconf feature
- Check that you can see config data for simulated device by executing GET request to

```
http://localhost:8181/restconf/config/network-topology:network-topology/topology/
↪topology-netconf/node/17830-sim-device/yang-ext:mount/
```

- The data should be just and empty data container
- Now execute edit-config request by executing a POST request to:

```
http://localhost:8181/restconf/config/network-topology:network-topology/topology/
↪topology-netconf/node/17830-sim-device/yang-ext:mount
```

  with headers:

```
Accept application/xml
Content-Type application/xml
```

  and payload:

```
<cont xmlns="urn:opendaylight:test">
  <l>Content</l>
</cont>
```

- Check that you can see modified config data for simulated device by executing GET request to

```
http://localhost:8181/restconf/config/network-topology:network-topology/topology/
↪topology-netconf/node/17830-sim-device/yang-ext:mount/
```

- Check that you can see the same modified data in operational for simulated device by executing GET request to

```
http://localhost:8181/restconf/operational/network-topology:network-topology/
↪topology/topology-netconf/node/17830-sim-device/yang-ext:mount/
```

Warning: Data will be mirrored in operational datastore only when using the default simple datastore.

Known problems

Slow creation of devices on virtual machines

When testtool seems to take unusually long time to create the devices use this flag when running it:

```
-Dorg.apache.sshd.registerBouncyCastle=false
```

Too many files open

When testtool or OpenDaylight starts to fail with TooManyFilesOpen exception, you need to increase the limit of open files in your OS. To find out the limit in linux execute:

```
ulimit -a
```

Example sufficient configuration in linux:

```
core file size          (blocks, -c) 0
data seg size           (kbytes, -d) unlimited
scheduling priority             (-e) 0
file size               (blocks, -f) unlimited
pending signals                 (-i) 63338
max locked memory       (kbytes, -l) 64
max memory size         (kbytes, -m) unlimited
open files                      (-n) 500000
pipe size            (512 bytes, -p) 8
POSIX message queues     (bytes, -q) 819200
real-time priority              (-r) 0
stack size              (kbytes, -s) 8192
cpu time               (seconds, -t) unlimited
max user processes              (-u) 63338
virtual memory          (kbytes, -v) unlimited
file locks                      (-x) unlimited
```

To set these limits edit file: /etc/security/limits.conf, for example:

```
*          hard    nofile      500000
*          soft    nofile      500000
root       hard    nofile      500000
root       soft    nofile      500000
```

"Killed"

The testtool might end unexpectedly with a simple message: "Killed". This means that the OS killed the tool due to too much memory consumed or too many threads spawned. To find out the reason on linux you can use following command:

```
dmesg | egrep -i -B100 'killed process'
```

Also take a look at this file: /proc/sys/kernel/threads-max. It limits the number of threads spawned by a process. Sufficient (but probably much more than enough) value is, e.g., 126676

NETCONF stress/performance measuring tool

This is basically a NETCONF client that puts NETCONF servers under heavy load of NETCONF RPCs and measures the time until a configurable amount of them is processed.

RESTCONF stress-performance measuring tool

Very similar to NETCONF stress tool with the difference of using RESTCONF protocol instead of NETCONF.

YANGLIB remote repository

There are scenarios in NETCONF deployment, that require for a centralized YANG models repository. YANGLIB plugin provides such remote repository.

To start this plugin, you have to install odl-yanglib feature. Then you have to configure YANGLIB either through RESTCONF or NETCONF. We will show how to configure YANGLIB through RESTCONF.

YANGLIB configuration through RESTCONF

You have to specify what local YANG modules directory you want to provide. Then you have to specify address and port whre you want to provide YANG sources. For example, we want to serve yang sources from folder /sources on localhost:5000 adress. The configuration for this scenario will be as follows:

```
PUT  http://localhost:8181/restconf/config/network-topology:network-topology/topology/
↪topology-netconf/node/controller-config/yang-ext:mount/config:modules/module/
↪yanglib:yanglib/example
```

Headers:

- Accept: application/xml
- Content-Type: application/xml

Payload:

```
<module xmlns="urn:opendaylight:params:xml:ns:yang:controller:config">
  <name>example</name>
  <type xmlns:prefix="urn:opendaylight:params:xml:ns:yang:controller:yanglib:impl">
↪prefix:yanglib</type>
  <broker xmlns="urn:opendaylight:params:xml:ns:yang:controller:yanglib:impl">
    <type xmlns:prefix="urn:opendaylight:params:xml:ns:yang:controller:md:sal:binding
↪">prefix:binding-broker-osgi-registry</type>
    <name>binding-osgi-broker</name>
  </broker>
  <cache-folder xmlns="urn:opendaylight:params:xml:ns:yang:controller:yanglib:impl">/
↪sources</cache-folder>
  <binding-addr xmlns="urn:opendaylight:params:xml:ns:yang:controller:yanglib:impl">
↪localhost</binding-addr>
  <binding-port xmlns="urn:opendaylight:params:xml:ns:yang:controller:yanglib:impl">
↪5000</binding-port>
</module>
```

This should result in a 2xx response and new YANGLIB instance should be created. This YANGLIB takes all YANG sources from /sources folder and for each generates URL in form:

```
http://localhost:5000/schemas/{modelName}/{revision}
```

On this URL will be hosted YANG source for particular module.

YANGLIB instance also write this URL along with source identifier to ietf-netconf-yang-library/modules-state/module list.

Netconf-connector with YANG library as fallback

There is an optional configuration in netconf-connector called yang-library. You can specify YANG library to be plugged as additional source provider into the mount's schema repository. Since YANGLIB plugin is advertising provided modules through yang-library model, we can use it in mount point's configuration as YANG library. To do this, we need to modify the configuration of netconf-connector by adding this XML

```
<yang-library xmlns="urn:opendaylight:netconf-node-topology">
  <yang-library-url xmlns="urn:opendaylight:netconf-node-topology">http://
↪localhost:8181/restconf/operational/ietf-yang-library:modules-state</yang-library-
↪url>
  <username xmlns="urn:opendaylight:netconf-node-topology">admin</username>
  <password xmlns="urn:opendaylight:netconf-node-topology">admin</password>
</yang-library>
```

This will register YANGLIB provided sources as a fallback schemas for particular mount point.

NetIDE User Guide

Overview

OpenDaylight's NetIDE project allows users to run SDN applications written for different SDN controllers, e.g., Floodlight or Ryu, on top of OpenDaylight managed infrastructure. The NetIDE Network Engine integrates a client controller layer that executes the modules that compose a Network Application and interfaces with a server SDN controller layer that drives the underlying infrastructure. In addition, it provides a uniform interface to common tools that are intended to allow the inspection/debug of the control channel and the management of the network resources.

The Network Engine provides a compatibility layer capable of translating calls of the network applications running on top of the client controllers, into calls for the server controller framework. The communication between the client and the server layers is achieved through the NetIDE intermediate protocol, which is an application-layer protocol on top of TCP that transmits the network control/management messages from the client to the server controller and vice-versa. Between client and server controller sits the Core Layer which also speaks the intermediate protocol.

NetIDE API

Architecture and Design

The NetIDE engine follows the ONF's proposed Client/Server SDN Application architecture.

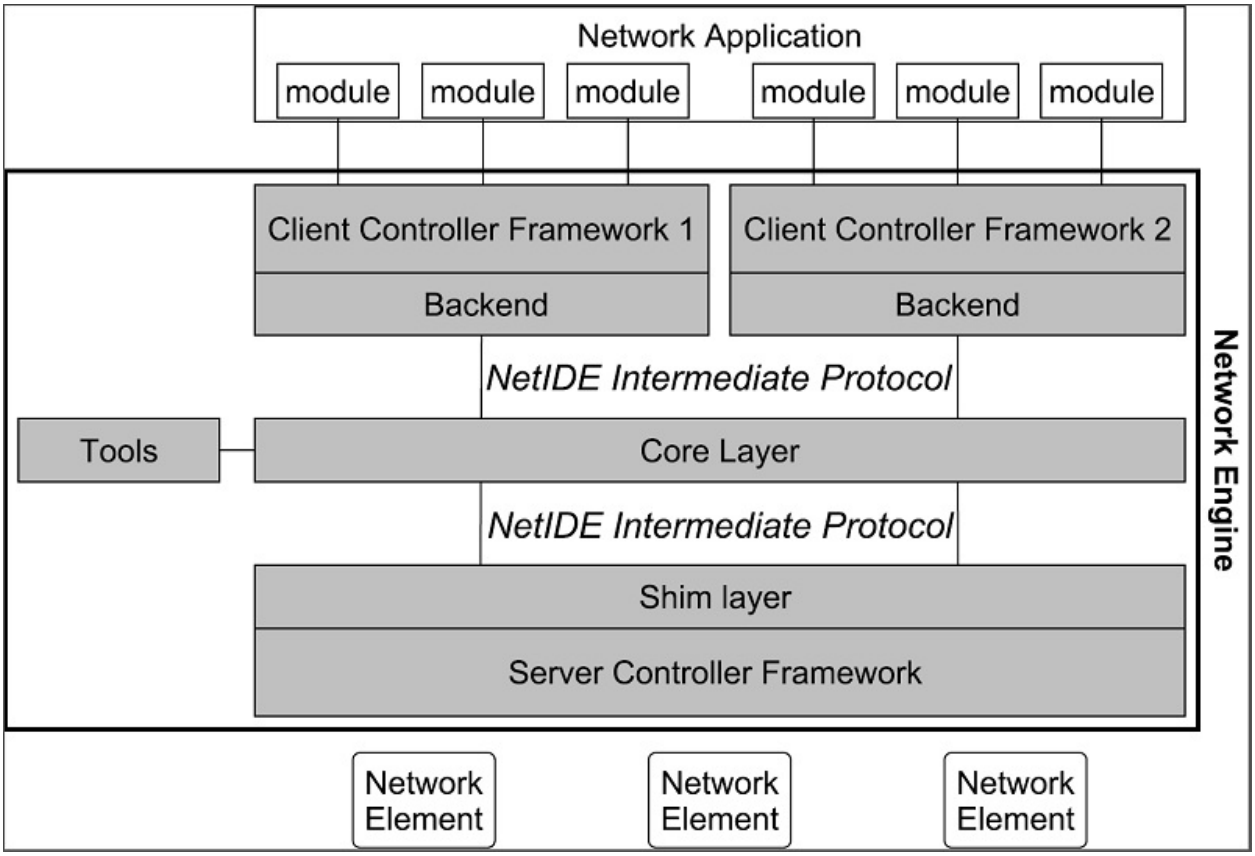

Fig. 67: NetIDE Network Engine Architecture

Core

The NetIDE Core is a message-based system that allows for the exchange of messages between OpenDaylight and subscribed Client SDN Controllers

Handling reply messages correctly

When an application module sends a request to the network (e.g. flow statistics, features, etc.), the Network Engine must be able to correctly drive the corresponding reply to such a module. This is not a trivial task, as many modules

may compose the network application running on top of the Network Engine, and there is no way for the Core to pair replies and requests. The transaction IDs (xid) in the OpenFlow header are unusable in this case, as it may happen that different modules use the same values.

In the proposed approach, represented in the figure below, the task of pairing replies with requests is performed by the Shim Layer which replaces the original xid of the OpenFlow requests coming from the core with new unique xid values. The Shim also saves the original OpenFlow xid value and the module id it finds in the NetIDE header. As the network elements must use the same xid values in the replies, the Shim layer can easily pair a reply with the correct request as it is using unique xid values.

The below figure shows how the Network Engine should handle the controller-to-switch OpenFlow messages. The diagram shows the case of a request message sent by an application module to a network element where the Backend inserts the module id of the module in the NetIDE header (X in the Figure). For other messages generated by the client controller platform (e.g. echo requests) or by the Backend, the module id of the Backend is used (Y in the Figure).

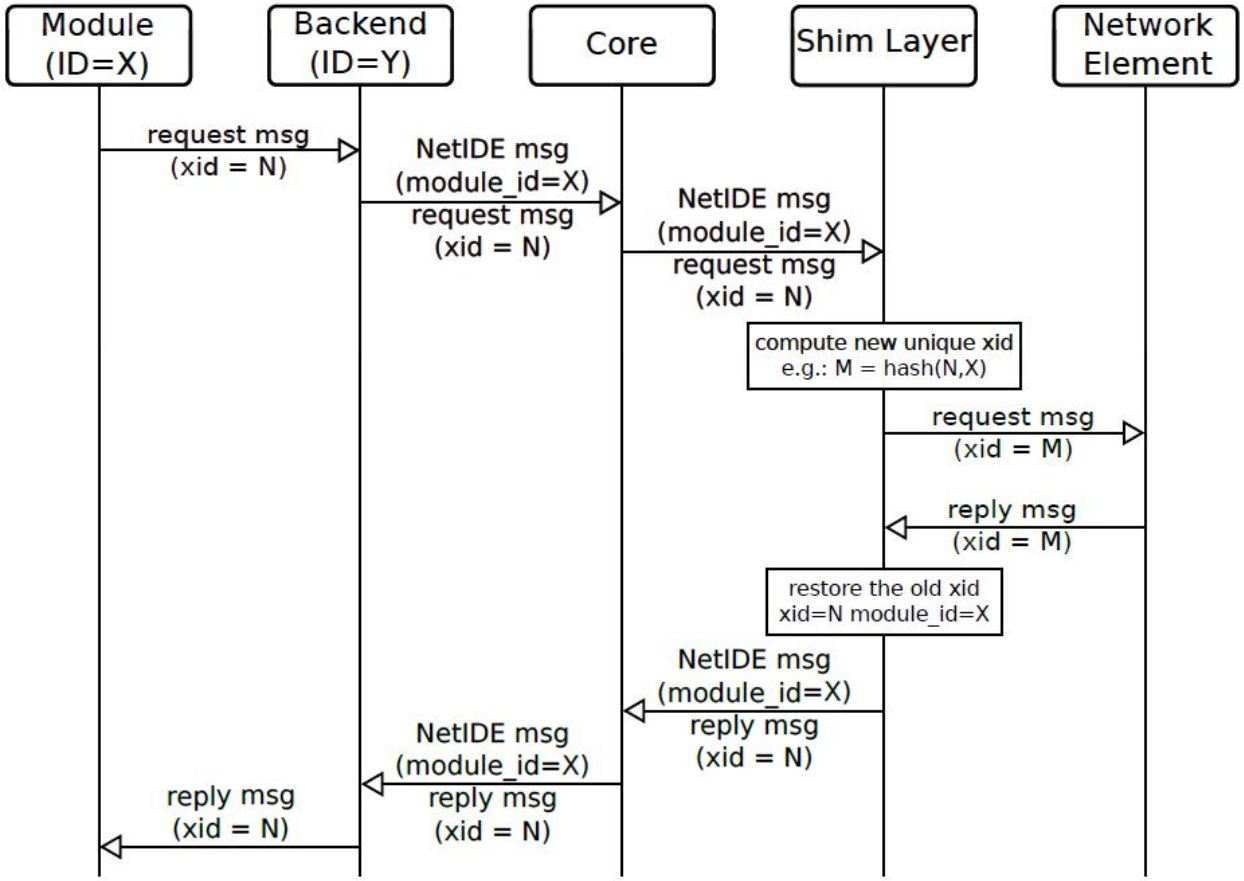

Fig. 68: NetIDE Communication Flow

Configuration

Below are the configuration items which can be edited, including their default values.

- core-address: This is the ip address of the NetIDE Core, default is 127.0.0.1
- core-port: The port of on which the NetIDE core is listening on
- address: IP address where the controller listens for switch connections, default is 127.0.0.1

- port: Port where controller listens for switch connections, default: 6644
- transport-protocol: default is TCP
- switch-idle-timeout: default is 15000ms

NetVirt User Guide

L3VPN Service: User Guide

Overview

L3VPN Service in OpenDaylight provides a framework to create L3VPN based on BGP-MP. It also helps to create Network Virtualization for DC Cloud environment.

Modules & Interfaces

L3VPN service can be realized using the following modules -

VPN Service Modules

1. **VPN Manager** : Creates and manages VPNs and VPN Interfaces
2. **BGP Manager** : Configures BGP routing stack and provides interface to routing services
3. **FIB Manager** : Provides interface to FIB, creates and manages forwarding rules in Dataplane
4. **Nexthop Manager** : Creates and manages nexthop egress pointer, creates egress rules in Dataplane
5. **Interface Manager** : Creates and manages different type of network interfaces, e.g., VLAN, l3tunnel etc.,
6. **Id Manager** : Provides cluster-wide unique ID for a given key. Used by different modules to get unique IDs for different entities.
7. **MD-SAL Util** : Provides interface to MD-SAL. Used by service modules to access MD-SAL Datastore and services.

All the above modules can function independently and can be utilized by other services as well.

Configuration Interfaces

The following modules expose configuration interfaces through which user can configure L3VPN Service.

1. BGP Manager
2. VPN Manager
3. Interface Manager
4. FIB Manager

Configuration Interface Details

1. Data Node Path : */config/bgp:bgp-router/*
   (a) Fields :
      i. local-as-identifier
      ii. local-as-number
   (b) REST Methods : GET, PUT, DELETE, POST
2. Data Node Path : */config/bgp:bgp-neighbors/*
   (a) Fields :
      i. List of bgp-neighbor
   (b) REST Methods : GET, PUT, DELETE, POST
3. Data Node Path : */config/bgp:bgp-neighbors/bgp-neighbor/‘‘{as-number}‘‘/*
   (a) Fields :
      i. as-number
      ii. ip-address
   (b) REST Methods : GET, PUT, DELETE, POST

1. Data Node Path : */config/l3vpn:vpn-instances/*
   (a) Fields :
      i. List of vpn-instance
   (b) REST Methods : GET, PUT, DELETE, POST
2. Data Node Path : */config/l3vpn:vpn-interfaces/vpn-instance*
   (a) Fields :
      i. name
      ii. route-distinguisher
      iii. import-route-policy
      iv. export-route-policy
   (b) REST Methods : GET, PUT, DELETE, POST
3. Data Node Path : */config/l3vpn:vpn-interfaces/*
   (a) Fields :
      i. List of vpn-interface
   (b) REST Methods : GET, PUT, DELETE, POST
4. Data Node Path : */config/l3vpn:vpn-interfaces/vpn-interface*
   (a) Fields :
      i. name
      ii. vpn-instance-name
   (b) REST Methods : GET, PUT, DELETE, POST
5. Data Node Path : */config/l3vpn:vpn-interfaces/vpn-interface/‘‘{name}‘‘/adjacency*

(a) Fields :

i. ip-address

ii. mac-address

(b) REST Methods : GET, PUT, DELETE, POST

1. Data Node Path : */config/if:interfaces/interface*

(a) Fields :

i. name

ii. type

iii. enabled

iv. of-port-id

v. tenant-id

vi. base-interface

(b) type specific fields

i. when type = *l2vlan*

A. vlan-id

ii. when type = *stacked_vlan*

A. stacked-vlan-id

iii. when type = *l3tunnel*

A. tunnel-type

B. local-ip

C. remote-ip

D. gateway-ip

iv. when type = *mpls*

A. list labelStack

B. num-labels

(c) REST Methods : GET, PUT, DELETE, POST

1. Data Node Path : */config/odl-fib:fibEntries/vrfTables*

(a) Fields :

i. List of vrfTables

(b) REST Methods : GET, PUT, DELETE, POST

2. Data Node Path : */config/odl-fib:fibEntries/vrfTables/‘‘{routeDistinguisher}‘‘/*

(a) Fields :

i. route-distinguisher

ii. list vrfEntries

A. destPrefix

B. label

C. nexthopAddress

(b) REST Methods : GET, PUT, DELETE, POST

3. Data Node Path : */config/odl-fib:fibEntries/ipv4Table*

(a) Fields :

i. list ipv4Entry

A. destPrefix

B. nexthopAddress

(b) REST Methods : GET, PUT, DELETE, POST

Provisioning Sequence & Sample Configurations

Installation

1. Edit *etc/custom.properties* and set the following property: *vpnservice.bgpspeaker.host.name = <bgpserver-ip>* *<bgpserver-ip>* here refers to the IP address of the host where BGP is running.
2. Run ODL and install VPN Service *feature:install odl-vpnservice-core*

Use REST interface to configure L3VPN service

Pre-requisites:

1. BGP stack with VRF support needs to installed and configured
   (a) *Configure BGP as specified in Step 1 below.*
2. Create pairs of GRE/VxLAN Tunnels (using ovsdb/ovs-vsctl) between each switch and between each switch to the Gateway node
   (a) *Create *l3tunnel* interfaces corresponding to each tunnel in interfaces DS as specified in Step 2 below.**

Step 1 : Configure BGP

1. Configure BGP Router

REST API : *PUT /config/bgp:bgp-router/*

Sample JSON Data

```
{
    "bgp-router": {
        "local-as-identifier": "10.10.10.10",
        "local-as-number": 108
    }
}
```

2. Configure BGP Neighbors

REST API : *PUT /config/bgp:bgp-neighbors/*

Sample JSON Data

```
{
   "bgp-neighbor" : [
          {
              "as-number": 105,
              "ip-address": "169.144.42.168"
          }
     ]
 }
```

Step 2 : Create Tunnel Interfaces

Create l3tunnel interfaces corresponding to all GRE/VxLAN tunnels created with ovsdb (*refer Prerequisites*). Use following REST Interface -

REST API : *PUT /config/if:interfaces/if:interfacce*

Sample JSON Data

```
{
    "interface": [
        {
            "name" : "GRE_192.168.57.101_192.168.57.102",
            "type" : "odl-interface:l3tunnel",
            "odl-interface:tunnel-type": "odl-interface:tunnel-type-gre",
            "odl-interface:local-ip" : "192.168.57.101",
            "odl-interface:remote-ip" : "192.168.57.102",
            "odl-interface:portId" : "openflow:1:3",
            "enabled" : "true"
        }
    ]
}
```

Following is expected as a result of these configurations

1. Unique If-index is generated
2. *Interface-state* operational DS is updated
3. Corresponding Nexthop Group Entry is created

Step 3 : OS Create Neutron Ports and attach VMs

At this step user creates VMs.

Step 4 : Create VM Interfaces

Create l2vlan interfaces corresponding to VM created in step 3

REST API : *PUT /config/if:interfaces/if:interface*

Sample JSON Data

```
{
    "interface": [
        {
            "name" : "dpn1-dp1.2",
            "type" : "l2vlan",
            "odl-interface:of-port-id" : "openflow:1:2",
            "odl-interface:vlan-id" : "1",
            "enabled" : "true"
        }
    ]
}
```

Step 5: Create VPN Instance

REST API : *PUT /config/l3vpn:vpn-instances/l3vpn:vpn-instance/*

Sample JSON Data

```
{
  "vpn-instance": [
    {
        "description": "Test VPN Instance 1",
        "vpn-instance-name": "testVpn1",
        "ipv4-family": {
            "route-distinguisher": "4000:1",
            "export-route-policy": "4000:1,5000:1",
            "import-route-policy": "4000:1,5000:1",
        }
    }
  ]
}
```

Following is expected as a result of these configurations

1. VPN ID is allocated and updated in data-store
2. Corresponding VRF is created in BGP
3. If there are vpn-interface configurations for this VPN, corresponding action is taken as defined in step 5

Step 5 : Create VPN-Interface and Local Adjacency

this can be done in two steps as well

1. Create vpn-interface

REST API : *PUT /config/l3vpn:vpn-interfaces/l3vpn:vpn-interface/*

Sample JSON Data

```
{
  "vpn-interface": [
    {
      "vpn-instance-name": "testVpn1",
      "name": "dpn1-dp1.2",
    }
  ]
}
```

Note: name here is the name of VM interface created in step 3, 4

2. Add Adjacencies on vpn-interafce

REST API : *PUT /config/l3vpn:vpn-interfaces/l3vpn:vpn-interface/dpn1-dp1.3/adjacency*

Sample JSON Data

```
  {
     "adjacency" : [
            {
                "ip-address" : "169.144.42.168",
                "mac-address" : "11:22:33:44:55:66"
            }
       ]
   }

its a list, user can define more than one adjacency on a
vpn\_interface
```

Above steps can be carried out in a single step as following

```
{
    "vpn-interface": [
        {
            "vpn-instance-name": "testVpn1",
            "name": "dpn1-dp1.3",
            "odl-l3vpn:adjacency": [
                {
                    "odl-l3vpn:mac_address": "11:22:33:44:55:66",
                    "odl-l3vpn:ip_address": "11.11.11.2",
                }
            ]
        }
    ]
}
```

Following is expected as a result of these configurations

1. Prefix label is generated and stored in DS
2. Ingress table is programmed with flow corresponding to interface
3. Local Egress Group is created

4. Prefix is added to BGP for advertisement
5. BGP pushes route update to FIB YANG Interface
6. FIB Entry flow is added to FIB Table in OF pipeline

Neutron Service User Guide

Overview

This Karaf feature (`odl-neutron-service`) provides integration support for OpenStack Neutron via the OpenDaylight ML2 mechanism driver. The Neutron Service is only one of the components necessary for OpenStack integration. For those related components please refer to documentations of each component:

- https://wiki.openstack.org/wiki/Neutron
- https://launchpad.net/networking-odl
- http://git.openstack.org/cgit/openstack/networking-odl/
- https://wiki.opendaylight.org/view/NeutronNorthbound:Main

Use cases and who will use the feature

If you want OpenStack integration with OpenDaylight, you will need this feature with an OpenDaylight provider feature like ovsdb/netvirt, group based policy, VTN, and lisp mapper. For provider configuration, please refer to each individual provider's documentation. Since the Neutron service only provides the northbound API for the OpenStack Neutron ML2 mechanism driver. Without those provider features, the Neutron service itself isn't useful.

Neutron Service feature Architecture

The Neutron service provides northbound API for OpenStack Neutron via RESTCONF and also its dedicated REST API. It communicates through its YANG model with providers.

Configuring Neutron Service feature

As the Karaf feature includes everything necessary for communicating northbound, no special configuration is needed. Usually this feature is used with an OpenDaylight southbound plugin that implements actual network virtualization functionality and OpenStack Neutron. The user wants to setup those configurations. Refer to each related documentations for each configurations.

Administering or Managing `odl-neutron-service`

There is no specific configuration regarding to Neutron service itself. For related configuration, please refer to OpenStack Neutron configuration and OpenDaylight related services which are providers for OpenStack.

installing `odl-neutron-service` while the controller running

1. While OpenDaylight is running, in Karaf prompt, type: `feature:install odl-neutron-service`.
2. Wait a while until the initialization is done and the controller stabilizes.

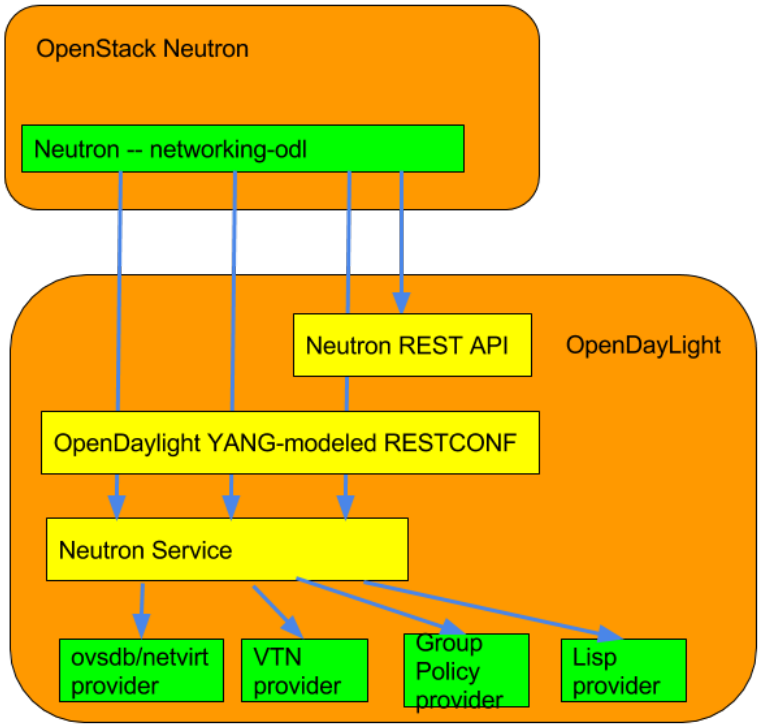

Fig. 69: Neutron Service Architecture

`odl-neutron-service` provides only a unified interface for OpenStack Neutron. It doesn't provide actual functionality for network virtualization. Refer to each OpenDaylight project documentation for actual configuration with OpenStack Neutron.

Neutron Logger

Another service, the Neutron Logger, is provided for debugging/logging purposes. It logs changes on Neutron YANG models.

```
feature:install odl-neutron-logger
```

Network Intent Composition (NIC) User Guide

Overview

Network Intent Composition (NIC) is an interface that allows clients to express a desired state in an implementation-neutral form that will be enforced via modification of available resources under the control of the OpenDaylight system.

This description is purposely abstract as an intent interface might encompass network services, virtual devices, storage, etc.

The intent interface is meant to be a controller-agnostic interface so that "intents" are portable across implementations, such as OpenDaylight and ONOS. Thus an intent specification should not contain implementation or technology specifics.

The intent specification will be implemented by decomposing the intent and augmenting it with implementation specifics that are driven by local implementation rules, policies, and/or settings.

Network Intent Composition (NIC) Architecture

The core of the NIC architecture is the intent model, which specifies the details of the desired state. It is the responsibility of the NIC implementation transforms this desired state to the resources under the control of OpenDaylight. The component that transforms the intent to the implementation is typically referred to as a renderer.

For the Boron release, multiple, simultaneous renderers will not be supported. Instead either the VTN or GBP renderer feature can be installed, but not both.

For the Boron release, the only actions supported are "ALLOW" and "BLOCK". The "ALLOW" action indicates that traffic can flow between the source and destination end points, while "BLOCK" prevents that flow; although it is possible that an given implementation may augment the available actions with additional actions.

Besides transforming a desired state to an actual state it is the responsibility of a renderer to update the operational state tree for the NIC data model in OpenDaylight to reflect the intent which the renderer implemented.

Configuring Network Intent Composition (NIC)

For the Boron release there is no default implementation of a renderer, thus without an additional module installed the NIC will not function.

Administering or Managing Network Intent Composition (NIC)

There is no additional administration of management capabilities related to the Network Intent Composition features.

Interactions

A user can interact with the Network Intent Composition (NIC) either through the RESTful interface using standard RESTCONF operations and syntax or via the Karaf console CLI.

REST

Configuration

The Network Intent Composition (NIC) feature supports the following REST operations against the configuration data store.

- POST - creates a new instance of an intent in the configuration store, which will trigger the realization of that intent. An ID *must* be specified as part of this request as an attribute of the intent.
- GET - fetches a list of all configured intents or a specific configured intent.
- DELETE - removes a configured intent from the configuration store, which triggers the removal of the intent from the network.

Operational

The Network Intent Composition (NIC) feature supports the following REST operations against the operational data store.

- GET - fetches a list of all operational intents or a specific operational intent.

Karaf Console CLI

This feature provides karaf console CLI command to manipulate the intent data model. The CLI essentailly invokes the equivalent data operations.

intent:add

Creates a new intent in the configuration data tree

```
DESCRIPTION
        intent:add

    Adds an intent to the controller.

Examples: --actions [ALLOW] --from <subject> --to <subject>
          --actions [BLOCK] --from <subject>

SYNTAX
        intent:add [options]
```

(continues on next page)

(continued from previous page)

```
OPTIONS
        -a, --actions
                Action to be performed.
                -a / --actions BLOCK/ALLOW
                (defaults to [BLOCK])
        --help
                Display this help message
        -t, --to
                Second Subject.
                -t / --to <subject>
                (defaults to any)
        -f, --from
                First subject.
                -f / --from <subject>
                (defaults to any)
```

intent:delete

Removes an existing intent from the system

```
DESCRIPTION
        intent:remove

    Removes an intent from the controller.

SYNTAX
        intent:remove id

ARGUMENTS
        id  Intent Id
```

intent:list

Lists all the intents in the system

```
DESCRIPTION
        intent:list

    Lists all intents in the controller.

SYNTAX
        intent:list [options]

OPTIONS
        -c, --config
                List Configuration Data (optional).
                -c / --config <ENTER>
        --help
                Display this help message
```

intent:show

Displayes the details of a single intent

```
DESCRIPTION
        intent:show

    Shows detailed information about an intent.

SYNTAX
        intent:show id

ARGUMENTS
        id  Intent Id
```

intent:map

List/Add/Delete current state from/to the mapping service.

```
DESCRIPTION
        intent:map

        List/Add/Delete current state from/to the mapping service.

SYNTAX
        intent:map [options]

         Examples: --list, -l [ENTER], to retrieve all keys.
                   --add-key <key> [ENTER], to add a new key with empty contents.
                   --del-key <key> [ENTER], to remove a key with it's values."
                   --add-key <key> --value [<value 1>, <value 2>, ...] [ENTER],
                     to add a new key with some values (json format).
OPTIONS
       --help
           Display this help message
       -l, --list
           List values associated with a particular key.
       -l / --filter <regular expression> [ENTER]
       --add-key
           Adds a new key to the mapping service.
       --add-key <key name> [ENTER]
       --value
           Specifies which value should be added/delete from the mapping service.
       --value "key=>value"... --value "key=>value" [ENTER]
           (defaults to [])
       --del-key
           Deletes a key from the mapping service.
       --del-key <key name> [ENTER]
```

NIC Usage Examples

Default Requirements

Start mininet, and create three switches (s1, s2, and s3) and four hosts (h1, h2, h3, and h4) in it.

Replace <Controller IP> based on your environment.

```
$  sudo mn --mac --topo single,2 --controller=remote,ip=<Controller IP>
```

```
mininet> net
h1 h1-eth0:s2-eth1
h2 h2-eth0:s2-eth2
h3 h3-eth0:s3-eth1
h4 h4-eth0:s3-eth2
s1 lo:  s1-eth1:s2-eth3 s1-eth2:s3-eth3
s2 lo:  s2-eth1:h1-eth0 s2-eth2:h2-eth0 s2-eth3:s1-eth1
s3 lo:  s3-eth1:h3-eth0 s3-eth2:h4-eth0 s3-eth3:s1-eth2
```

Downloading and deploy Karaf distribution

- Get the Boron distribution.
- Unzip the downloaded zip distribution.
- To run the Karaf.

```
./bin/karaf
```

- Once the console is up, type as below to install feature.

```
feature:install odl-nic-core-mdsal odl-nic-console odl-nic-listeners
```

Simple Mininet topology

```
!/usr/bin/python

from mininet.topo import Topo

class SimpleTopology( Topo ):
    "Simple topology example."

    def __init__( self ):
        "Create custom topo."


    Topo.__init__( self )


        Switch1 = self.addSwitch( 's1' )
        Switch2 = self.addSwitch( 's2' )
        Switch3 = self.addSwitch( 's3' )
        Switch4 = self.addSwitch( 's4' )
        Host11 = self.addHost( 'h1' )
        Host12 = self.addHost( 'h2' )
        Host21 = self.addHost( 'h3' )
        Host22 = self.addHost( 'h4' )
        Host23 = self.addHost( 'h5' )
        Service1 = self.addHost( 'srvc1' )
```

(continues on next page)

(continued from previous page)

```
        self.addLink( Host11, Switch1 )
        self.addLink( Host12, Switch1 )
        self.addLink( Host21, Switch2 )
        self.addLink( Host22, Switch2 )
        self.addLink( Host23, Switch2 )
        self.addLink( Switch1, Switch2 )
        self.addLink( Switch2, Switch4 )
        self.addLink( Switch4, Switch3 )
        self.addLink( Switch3, Switch1 )
        self.addLink( Switch3, Service1 )
        self.addLink( Switch4, Service1 )


topos = { 'simpletopology': ( lambda: SimpleTopology() ) }
```

- Initialize topology
- Add hosts and switches
- Host used to represent the service
- Add links

Source: https://gist.github.com/vinothgithub15/315d0a427d5afc39f2d7

How to configure VTN Renderer

The section demonstrates allow or block packets of the traffic within a VTN Renderer, according to the specified flow conditions.

The table below lists the actions to be applied when a packet matches the condition:

Action	Function
Allow	Permits the packet to be forwarded normally.
Block	Discards the packet preventing it from being forwarded.

Requirement

- Before execute the follow steps, please, use default requirements. See section *Default Requirements*.

Configuration

Please execute the following curl commands to test network intent using mininet:

Create Intent

To provision the network for the two hosts(h1 and h2) and demonstrates the action allow.

```
curl -v --user "admin":"admin" -H "Accept: application/json" -H "Content-type:␣
↪application/json" -X PUT http://localhost:8181/restconf/config/intent:intents/
↪intent/b9a13232-525e-4d8c-be21-cd65e3436034 -d '{ "intent:intent" : { "intent:id":
↪"b9a13232-525e-4d8c-be21-cd65e3436034", "intent:actions" : [ { "order" :␣
↪: {} } ], "intent:subjects" : [ { "order":1 , "end-point-group" : {"name":"10.0.0.1
↪"} }, { "order":2 , "end-point-group" : {"name":"10.0.0.2"}} ] } }'
```

(continues on next page)

(continued from previous page)

To provision the network for the two hosts(h2 and h3) and demonstrates the action allow.

```
curl -v --user "admin":"admin" -H "Accept: application/json" -H "Content-type:␣
↪application/json" -X PUT http://localhost:8181/restconf/config/intent:intents/
↪intent/b9a13232-525e-4d8c-be21-cd65e3436035 -d '{ "intent:intent" : { "intent:id":
↪"b9a13232-525e-4d8c-be21-cd65e3436035", "intent:actions" : [ { "order" : 2, "allow"␣
↪: {} } ], "intent:subjects" : [ { "order":1 , "end-point-group" : {"name":"10.0.0.2
↪"} }, { "order":2 , "end-point-group" : {"name":"10.0.0.3"}} ] } }'
```

Verification

As we have applied action type allow now ping should happen between hosts (h1 and h2) and (h2 and h3).

```
mininet> pingall
Ping: testing ping reachability
h1 -> h2 X X
h2 -> h1 h3 X
h3 -> X h2 X
h4 -> X X X
```

Update the intent

To provision block action that indicates traffic is not allowed between h1 and h2.

```
curl -v --user "admin":"admin" -H "Accept: application/json" -H "Content-type:␣
↪application/json" -X PUT http://localhost:8181/restconf/config/intent:intents/
↪intent/b9a13232-525e-4d8c-be21-cd65e3436034 -d '{ "intent:intent" : { "intent:id":
↪"b9a13232-525e-4d8c-be21-cd65e3436034", "intent:actions" : [ { "order" : 2, "block"␣
↪: {} } ], "intent:subjects" : [ { "order":1 , "end-point-group" : {"name":"10.0.0.1
↪"} }, { "order":2 , "end-point-group" : {"name":"10.0.0.2"}} ] } }'
```

Verification

As we have applied action type block now ping should not happen between hosts (h1 and h2).

```
mininet> pingall
Ping: testing ping reachability
h1 -> X X X
h2 -> X h3 X
h3 -> X h2 X
h4 -> X X X
```

Note: Old actions and hosts are replaced by the new action and hosts.

Delete the intent

Respective intent and the traffics will be deleted.

```
curl -v --user "admin":"admin" -H "Accept: application/json" -H     "Content-type:␣
↪application/json" -X DELETE http://localhost:8181/restconf/config/intent:intents/
↪intent/b9a13232-525e-4d8c-be21-cd65e3436035
```

Verification

Deletion of intent and flow.

```
mininet> pingall
Ping: testing ping reachability
h1 -> X X X
h2 -> X X X
h3 -> X X X
h4 -> X X X
```

Note: Ping between two hosts can also be done using MAC Address

To provision the network for the two hosts(h1 MAC address and h2 MAC address).

```
curl -v --user "admin":"admin" -H "Accept: application/json" -H "Content-type:␣
↪application/json" -X PUT http://localhost:8181/restconf/config/intent:intents/
↪intent/b9a13232-525e-4d8c-be21-cd65e3436035 -d '{ "intent:intent" : { "intent:id":
↪"b9a13232-525e-4d8c-be21-cd65e3436035", "intent:actions" : [ { "order" : 2, "allow"␣
↪: {} } ], "intent:subjects" : [ { "order":1 , "end-point-group" : {"name":
↪"6e:4f:f7:27:15:c9"} }, { "order":2 , "end-point-group" : {"name":"aa:7d:1f:4a:70:81
↪"}} ] } }'
```

How to configure Redirect Action

The section explains the redirect action supported in NIC. The redirect functionality supports forwarding (to redirect) the traffic to a service configured in SFC before forwarding it to the destination.

Following steps explain Redirect action function:

- Configure the service in SFC using the SFC APIs.
- Configure the intent with redirect action and the service information where the traffic needs to be redirected.
- The flows are computed as below
  1. First flow entry between the source host connected node and the ingress node of the configured service.
  2. Second flow entry between the egress Node id the configured service and the ID and destination host connected host.
  3. Third flow entry between the destination host node and the source host node.

Requirement

- Save the mininet *Simple Mininet topology* script as redirect_test.py
- Start mininet, and create switches in it.

Replace <Controller IP> based on your environment.

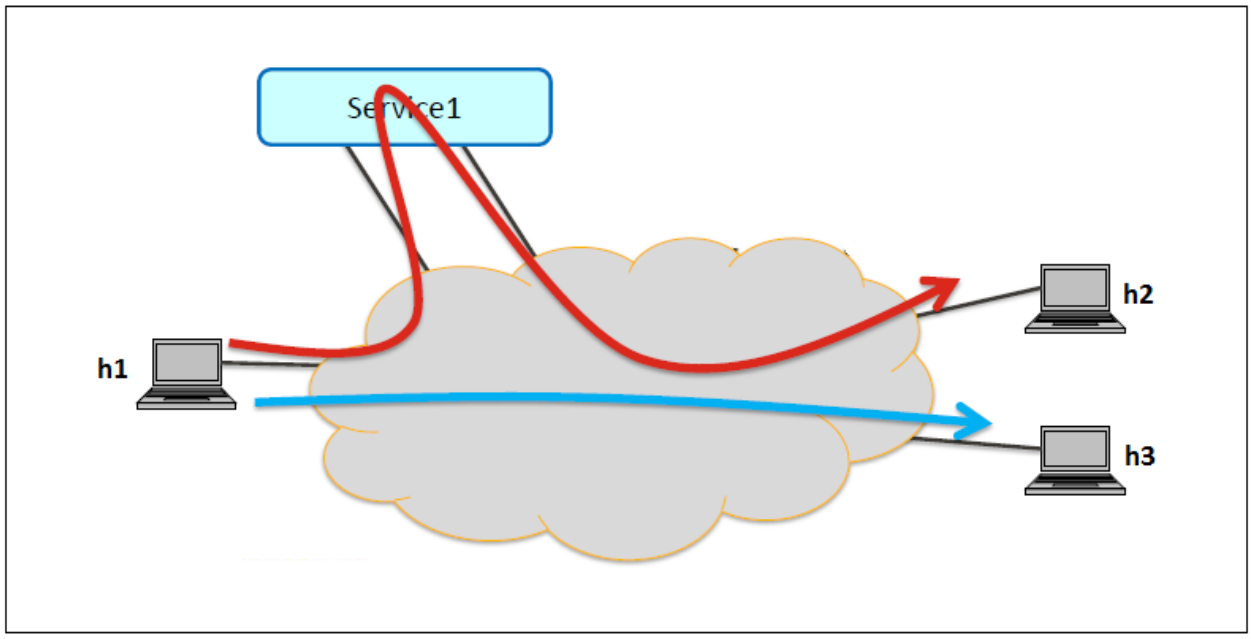

Fig. 70: REDIRECT SERVICE

```
sudo mn --controller=remote,ip=<Controller IP>--custom redirect_test.py --topo mytopo2
```

```
mininet> net
h1 h1-eth0:s1-eth1
h2 h2-eth0:s1-eth2
h3 h3-eth0:s2-eth1
h4 h4-eth0:s2-eth2
h5 h5-eth0:s2-eth3
srvc1 srvc1-eth0:s3-eth3 srvc1-eth1:s4-eth3
s1 lo:  s1-eth1:h1-eth0 s1-eth2:h2-eth0 s1-eth3:s2-eth4 s1-eth4:s3-eth2
s2 lo:  s2-eth1:h3-eth0 s2-eth2:h4-eth0 s2-eth3:h5-eth0 s2-eth4:s1-eth3 s2-eth5:s4-
↪eth1
s3 lo:  s3-eth1:s4-eth2 s3-eth2:s1-eth4 s3-eth3:srvc1-eth0
s4 lo:  s4-eth1:s2-eth5 s4-eth2:s3-eth1 s4-eth3:srvc1-eth1
c0
```

Starting the Karaf

- Before execute the following steps, please, use the default requirements. See section *Downloading and deploy Karaf distribution*.

Configuration

Mininet

- Configure srvc1 as service node in the mininet environment.

Please execute the following commands in the mininet console (where mininet script is executed).

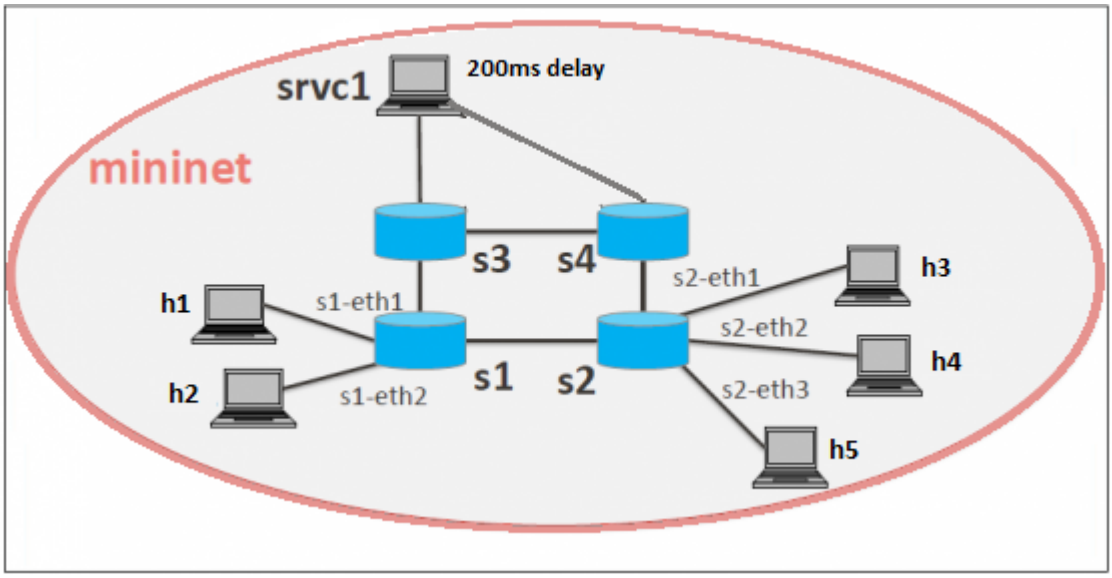

Fig. 71: CONFIGURATION THE NETWORK IN MININET

```
srvc1 ip addr del 10.0.0.6/8 dev srvc1-eth0
srvc1 brctl addbr br0
srvc1 brctl addif br0 srvc1-eth0
srvc1 brctl addif br0 srvc1-eth1
srvc1 ifconfig br0 up
srvc1 tc qdisc add dev srvc1-eth1 root netem delay 200ms
```

Configure service in SFC

The service (srvc1) is configured using SFC REST API. As part of the configuration the ingress and egress node connected the service is configured.

```
curl -i -H "Content-Type: application/json" -H "Cache-Control: no-cache" --data '{
  "service-functions": {
    "service-function": [
      {
        "name": "srvc1",
        "sf-data-plane-locator": [
          {
            "name": "Egress",
            "service-function-forwarder": "openflow:4"
          },
          {
            "name": "Ingress",
            "service-function-forwarder": "openflow:3"
          }
        ],
        "nsh-aware": false,
        "type": "delay"
      }
    ]
  }
```

(continues on next page)

(continued from previous page)

```
}' -X PUT --user admin:admin http://localhost:8181/restconf/config/service-
↪function:service-functions/
```

SFF RESTCONF Request

Configuring switch and port information for the service functions.

```
curl -i -H "Content-Type: application/json" -H "Cache-Control: no-cache" --data '{
  "service-function-forwarders": {
    "service-function-forwarder": [
      {
        "name": "openflow:3",
        "service-node": "OVSDB2",
        "sff-data-plane-locator": [
          {
            "name": "Ingress",
            "data-plane-locator":
            {
                "vlan-id": 100,
                "mac": "11:11:11:11:11:11",
                "transport": "service-locator:mac"
            },
            "service-function-forwarder-ofs:ofs-port":
            {
                "port-id" : "3"
            }
          }
        ],
        "service-function-dictionary": [
          {
            "name": "srvc1",
            "sff-sf-data-plane-locator":
            {
                "sf-dpl-name" : "openflow:3",
                "sff-dpl-name" : "Ingress"
            }
          }
        ]
      },
      {
        "name": "openflow:4",
        "service-node": "OVSDB3",
        "sff-data-plane-locator": [
          {
            "name": "Egress",
            "data-plane-locator":
            {
                "vlan-id": 200,
                "mac": "44:44:44:44:44:44",
                "transport": "service-locator:mac"
            },
            "service-function-forwarder-ofs:ofs-port":
            {
                "port-id" : "3"
            }
          }
        ],
```

(continues on next page)

(continued from previous page)

```
        "service-function-dictionary": [
          {
            "name": "srvc1",
            "sff-sf-data-plane-locator":
            {
                "sf-dpl-name" : "openflow:4",
                "sff-dpl-name" : "Egress"
            }
          }
        ]
      }
    ]
  }
}' -X PUT --user admin:admin http://localhost:8181/restconf/config/service-function-
↪forwarder:service-function-forwarders/
```

CLI Command

To provision the network for the two hosts (h1 and h5).

Demonstrates the redirect action with service name srvc1.

```
intent:add -f <SOURCE_MAC> -t <DESTINATION_MAC> -a REDIRECT -s <SERVICE_NAME>
```

Example:

```
intent:add -f 32:bc:ec:65:a7:d1 -t c2:80:1f:77:41:ed -a REDIRECT -s srvc1
```

Verification

- As we have applied action type redirect now ping should happen between hosts h1 and h5.

```
mininet> h1 ping h5
PING 10.0.0.5 (10.0.0.5) 56(84) bytes of data.
64 bytes from 10.0.0.5: icmp_seq=2 ttl=64 time=201 ms
64 bytes from 10.0.0.5: icmp_seq=3 ttl=64 time=200 ms
64 bytes from 10.0.0.5: icmp_seq=4 ttl=64 time=200 ms
```

The redirect functionality can be verified by the time taken by the ping operation (200ms). The service srvc1 configured using SFC introduces 200ms delay. As the traffic from h1 to h5 is redirected via the srvc1, the time taken by the traffic from h1 to h5 will take about 200ms.

- Flow entries added to nodes for the redirect action.

```
mininet> dpctl dump-flows
*** s1 ------------------------------------------------------------------------
NXST_FLOW reply (xid=0x4):
cookie=0x0, duration=9.406s, table=0, n_packets=6, n_bytes=588, idle_age=3,␣
↪priority=9000,in_port=1,dl_src=32:bc:ec:65:a7:d1, dl_dst=c2:80:1f:77:41:ed␣
↪actions=output:4
cookie=0x0, duration=9.475s, table=0, n_packets=6, n_bytes=588, idle_age=3,␣
↪priority=9000,in_port=3,dl_src=c2:80:1f:77:41:ed, dl_dst=32:bc:ec:65:a7:d1␣
↪actions=output:1
cookie=0x1, duration=362.315s, table=0, n_packets=144, n_bytes=12240, idle_age=4,␣
↪priority=9500,dl_type=0x88cc actions=CONTROLLER:65535
```

(continues on next page)

(continued from previous page)

```
cookie=0x1, duration=362.324s, table=0, n_packets=4, n_bytes=168, idle_age=3,␣
↪priority=10000,arp actions=CONTROLLER:65535,NORMAL
*** s2 ----------------------------------------------------------------------
NXST_FLOW reply (xid=0x4):
cookie=0x0, duration=9.503s, table=0, n_packets=6, n_bytes=588, idle_age=3,␣
↪priority=9000,in_port=3,dl_src=c2:80:1f:77:41:ed, dl_dst=32:bc:ec:65:a7:d1␣
↪actions=output:4
cookie=0x0, duration=9.437s, table=0, n_packets=6, n_bytes=588, idle_age=3,␣
↪priority=9000,in_port=5,dl_src=32:bc:ec:65:a7:d1, dl_dst=c2:80:1f:77:41:ed␣
↪actions=output:3
cookie=0x3, duration=362.317s, table=0, n_packets=144, n_bytes=12240, idle_age=4,␣
↪priority=9500,dl_type=0x88cc actions=CONTROLLER:65535
cookie=0x3, duration=362.32s, table=0, n_packets=4, n_bytes=168, idle_age=3,␣
↪priority=10000,arp actions=CONTROLLER:65535,NORMAL
*** s3 ----------------------------------------------------------------------
NXST_FLOW reply (xid=0x4):
cookie=0x0, duration=9.41s, table=0, n_packets=6, n_bytes=588, idle_age=3,␣
↪priority=9000,in_port=2,dl_src=32:bc:ec:65:a7:d1, dl_dst=c2:80:1f:77:41:ed␣
↪actions=output:3
*** s4 ----------------------------------------------------------------------
NXST_FLOW reply (xid=0x4):
cookie=0x0, duration=9.486s, table=0, n_packets=6, n_bytes=588, idle_age=3,␣
↪priority=9000,in_port=3,dl_src=32:bc:ec:65:a7:d1, dl_dst=c2:80:1f:77:41:ed␣
↪actions=output:1
```

How to configure QoS Attribute Mapping

This section explains how to provision QoS attribute mapping constraint using NIC OF-Renderer.

The QoS attribute mapping currently supports DiffServ. It uses a 6-bit differentiated services code point (DSCP) in the 8-bit differentiated services field (DS field) in the IP header.

Action	Function
Allow	Permits the packet to be forwarded normally, but allows for packet header fields, e.g., DSCP, to be modified.

The following steps explain QoS Attribute Mapping function:

- Initially configure the QoS profile which contains profile name and DSCP value.
- When a packet is transferred from a source to destination, the flow builder evaluates whether the transferred packet matches the condition such as action, endpoints in the flow.
- If the packet matches the endpoints, the flow builder applies the flow matching action and DSCP value.

Requirement

- Before execute the following steps, please, use the default requirements. See section *Default Requirements*.

Configuration

Please execute the following CLI commands to test network intent using mininet:

- To apply the QoS constraint, configure the QoS profile.

```
intent:qosConfig -p <qos_profile_name> -d <valid_dscp_value>
```

Example:

```
intent:qosConfig -p High_Quality -d 46
```

Note: Valid DSCP value ranges from 0-63.

- To provision the network for the two hosts (h1 and h3), add intents that allows traffic in both directions by execute the following CLI command.

Demonstrates the ALLOW action with constraint QoS and QoS profile name.

```
intent:add -a ALLOW -t <DESTINATION_MAC> -f <SOURCE_MAC> -q QOS -p <qos_profile_name>
```

Example:

```
intent:add -a ALLOW -t 00:00:00:00:00:03 -f 00:00:00:00:00:01 -q QOS -p High_Quality
intent:add -a ALLOW -t 00:00:00:00:00:01 -f 00:00:00:00:00:03 -q QOS -p High_Quality
```

Verification

- As we have applied action type ALLOW now ping should happen between hosts h1 and h3.

```
mininet> h1 ping h3
PING 10.0.0.3 (10.0.0.3) 56(84) bytes of data.
64 bytes from 10.0.0.3: icmp_req=1 ttl=64 time=0.984 ms
64 bytes from 10.0.0.3: icmp_req=2 ttl=64 time=0.110 ms
64 bytes from 10.0.0.3: icmp_req=3 ttl=64 time=0.098 ms
```

- Verification of the flow entry and ensuring the mod_nw_tos is part of actions.

```
mininet> dpctl dump-flows
*** s1 ------------------------------------------------------------------------
NXST_FLOW reply (xid=0x4):
cookie=0x0, duration=21.873s, table=0, n_packets=3, n_bytes=294, idle_age=21,␣
↪priority=9000,dl_src=00:00:00:00:00:03,dl_dst=00:00:00:00:00:01 actions=NORMAL,mod_
↪nw_tos:184
cookie=0x0, duration=41.252s, table=0, n_packets=3, n_bytes=294, idle_age=41,␣
↪priority=9000,dl_src=00:00:00:00:00:01,dl_dst=00:00:00:00:00:03 actions=NORMAL,mod_
↪nw_tos:184
```

Requirement

- Before execute the follow steps, please, use default requirements. See section *Default Requirements*.

How to configure Log Action

This section demonstrates log action in OF Renderer. This demonstration aims at enabling communication between two hosts and logging the flow statistics details of the particular traffic.

Configuration

Please execute the following CLI commands to test network intent using mininet:

- To provision the network for the two hosts (h1 and h3), add intents that allows traffic in both directions by execute the following CLI command.

```
intent:add -a ALLOW -t <DESTINATION_MAC> -f <SOURCE_MAC>
```

Example:

```
intent:add -a ALLOW -t 00:00:00:00:00:03 -f 00:00:00:00:00:01
intent:add -a ALLOW -t 00:00:00:00:00:01 -f 00:00:00:00:00:03
```

- To log the flow statistics details of the particular traffic.

```
intent:add -a LOG -t <DESTINATION_MAC> -f <SOURCE_MAC>
```

Example:

```
intent:add -a LOG -t 00:00:00:00:00:03 -f 00:00:00:00:00:01
```

Verification

- As we have applied action type ALLOW now ping should happen between hosts h1 and h3.

```
mininet> h1 ping h3
PING 10.0.0.3 (10.0.0.3) 56(84) bytes of data.
64 bytes from 10.0.0.3: icmp_req=1 ttl=64 time=0.984 ms
64 bytes from 10.0.0.3: icmp_req=2 ttl=64 time=0.110 ms
64 bytes from 10.0.0.3: icmp_req=3 ttl=64 time=0.098 ms
```

- To view the flow statistics log details such as, byte count, packet count and duration, check the karaf.log.

```
2015-12-15 22:56:20,256 | INFO | lt-dispatcher-23 | IntentFlowManager | 264 - org.
↪opendaylight.nic.of-renderer - 1.1.0.SNAPSHOT | Creating block intent for␣
↪endpoints: source00:00:00:00:00:01 destination 00:00:00:00:00:03
2015-12-15 22:56:20,252 | INFO | lt-dispatcher-29 | FlowStatisticsListener | 264 -␣
↪org.opendaylight.nic.of-renderer - 1.1.0.SNAPSHOT | Flow Statistics gathering for␣
↪Byte Count:Counter64 [_value=238]
2015-12-15 22:56:20,252 | INFO | lt-dispatcher-29 | FlowStatisticsListener | 264 -␣
↪org.opendaylight.nic.of-renderer - 1.1.0.SNAPSHOT | Flow Statistics gathering for␣
↪Packet Count:Counter64 [_value=3]
2015-12-15 22:56:20,252 | INFO | lt-dispatcher-29 | FlowStatisticsListener | 264 -␣
↪org.opendaylight.nic.of-renderer - 1.1.0.SNAPSHOT | Flow Statistics gathering for␣
↪Duration in Nano second:Counter32 [_value=678000000]
2015-12-15 22:56:20,252 | INFO | lt-dispatcher-29 | FlowStatisticsListener | 264 -␣
↪org.opendaylight.nic.of-renderer - 1.1.0.SNAPSHOT | Flow Statistics gathering for␣
↪Duration in Second:Counter32 [_value=49]
```

OCP Plugin User Guide

This document describes how to use the ORI Control & Management Protocol (OCP) feature in OpenDaylight. This document contains overview, scope, architecture and design, installation, configuration and tutorial sections for the feature.

Overview

OCP is an ETSI standard protocol for control and management of Remote Radio Head (RRH) equipment. The OCP Project addresses the need for a southbound plugin that allows applications and controller services to interact with RRHs using OCP. The OCP southbound plugin will allow applications acting as a Radio Equipment Control (REC) to interact with RRHs that support an OCP agent.

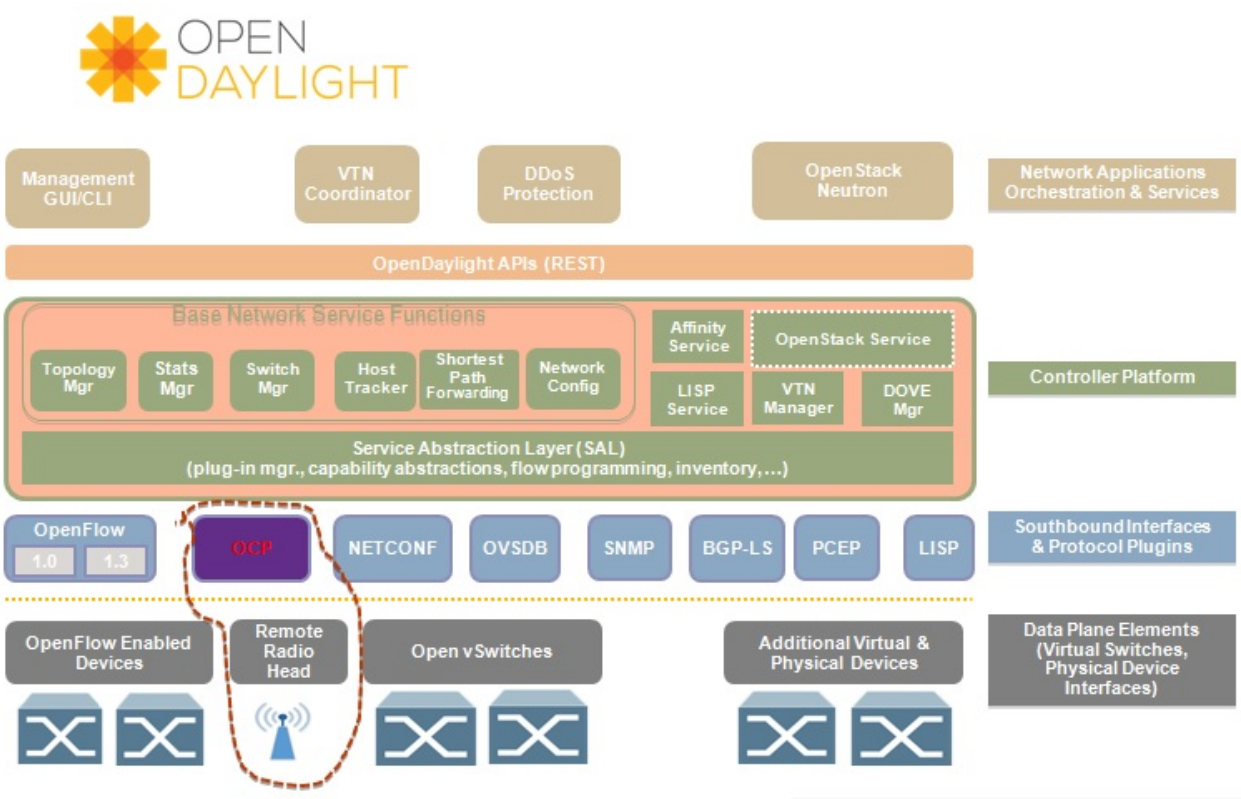

Fig. 72: OCP southbound plugin

It is foreseen that, in 5G, C-RAN will use the packet-based Transport-SDN (T-SDN) as the fronthaul network to transport both control plane and user plane data between RRHs and BBUs. As a result, the addition of the OCP plugin to OpenDaylight will make it possible to build an RRH controller on top of OpenDaylight to centrally manage deployed RRHs, as well as integrating the RRH controller with T-SDN on one single platform, achieving the joint RRH and fronthaul network provisioning in C-RAN.

Scope

The OCP Plugin project includes:

- OCP v4.1.1 support
- Integration of OCP protocol library
- Simple API invoked as a RPC
- Simple API that allows applications to perform elementary functions of the following categories:
  - Device management
  - Config management

  - Object lifecycle
  - Object state management
  - Fault management
  - Software management (not implemented as of Boron)
- Indication processing
- Logging (not implemented as of Boron)
- AISG/Iuant interface message tunnelling (not implemented as of Boron)
- ALD connection management (not implemented as of Boron)

Architecture and Design

OCP is a vendor-neutral standard communications interface defined to enable control and management between RE and REC of an ORI architecture. The OCP Plugin supports the implementation of the OCP specification; it is based on the Model Driven Service Abstraction Layer (MD-SAL) architecture.

OCP Plugin will support the following functionality:

- Connection handling
- Session management
- State management
- Error handling
- Connection establishment will be handled by OCP library using opensource netty.io library
- Message handling
- Event/indication handling and propagation to upper layers

Activities in OCP plugin module

- Integration with OCP protocol library
- Integration with corresponding MD-SAL infrastructure

OCP protocol library is a component in OpenDaylight that mediates communication between OpenDaylight controller and RRHs supporting OCP protocol. Its primary goal is to provide the OCP Plugin with communication channel that can be used for managing RRHs.

Key objectives:

- Immutable transfer objects generation (transformation of OCP protocol library's POJO objects into OpenDaylight DTO objects)
- Scalable non-blocking implementation
- Pipeline processing
- Scatter buffer
- TLS support

OCP Service addresses the need for a northbound interface that allows applications and other controller services to interact with RRHs using OCP, by providing API for abstracting OCP operations.

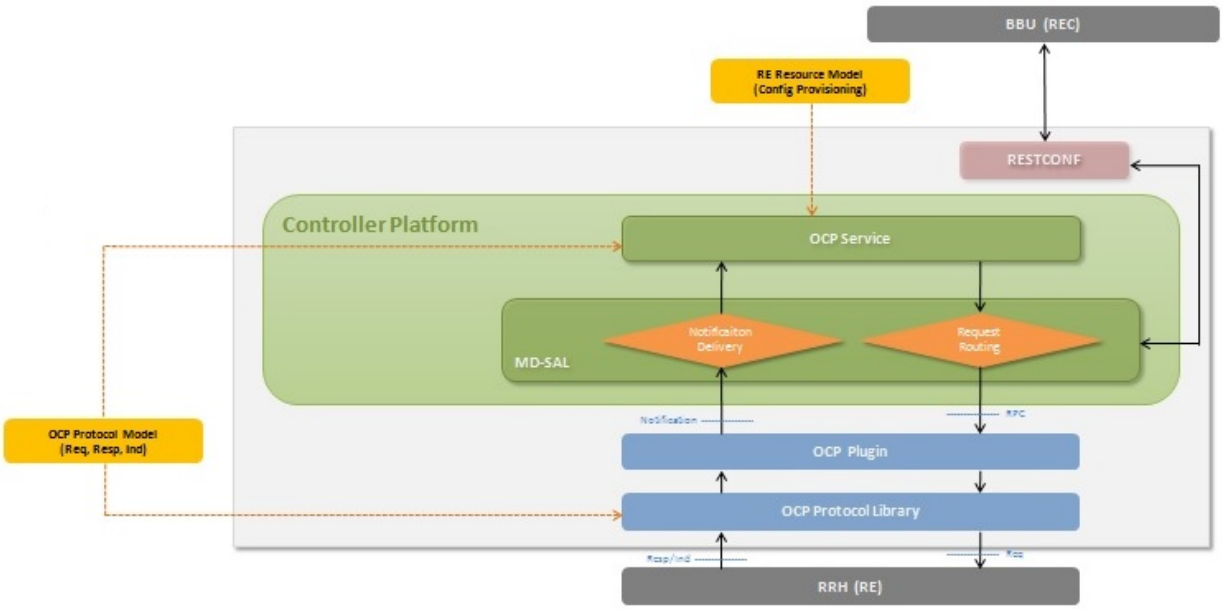

Fig. 73: Overall architecture

Message Flow

Installation

The OCP Plugin project has two top level Karaf features, odl-ocpplugin-all and odl-ocpjava-all, which contain the following sub-features:

- odl-ocpplugin-southbound
- odl-ocpplugin-app-ocp-service
- odl-ocpjava-protocol

The OCP service (odl-ocpplugin-app-ocp-service), together with the OCP southbound (odl-ocpplugin-southbound) and OCP protocol library (odl-ocpjava-protocol), provides OpenDaylight with basic OCP v4.1.1 functionality.

There are two ways to interact with OCP service: one is via RESTCONF (programmatic) and the other is using DLUX web interface (manual), so you have to install the following features to enable RESTCONF and DLUX.

```
karaf#>feature:install odl-restconf odl-l2switch-switch odl-mdsal-apidocs odl-dlux-
↪core odl-dlux-all
```

Then install the odl-ocpplugin-all feature which includes the odl-ocpplugin-southbound and odl-ocpplugin-app-ocp-service features. Note that the odl-ocpjava-all feature will be installed automatically as the odl-ocpplugin-southbound feature is dependent on the odl-ocpjava-protocol feature.

```
karaf#>feature:install odl-ocpplugin-all
```

After all required features are installed, use following command from karaf console to check and make sure features are correctly installed and initialized.

```
karaf#>feature:list | grep ocp
```

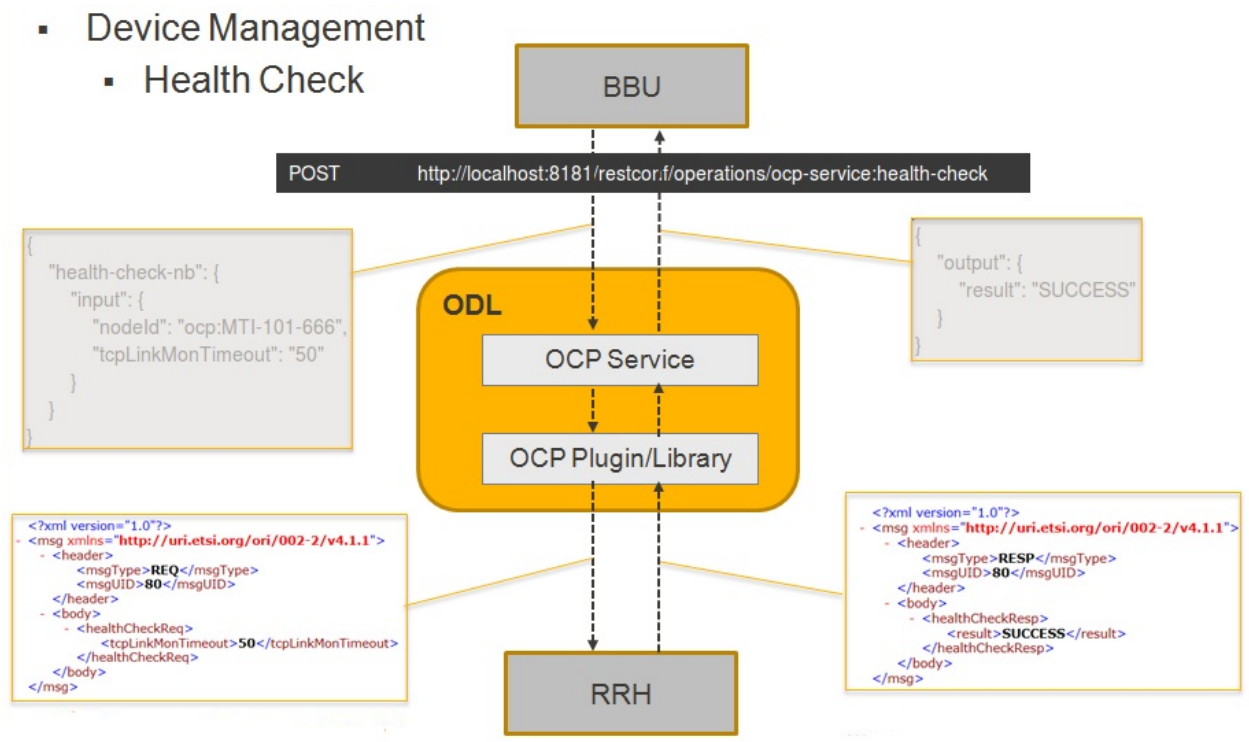

Fig. 74: Message flow example

Configuration

Configuring the OCP plugin can be done via its configuration file, 62-ocpplugin.xml, which can be found in the <odl-install-dir>/etc/opendaylight/karaf/ directory.

As of Boron, there are the following settings that are configurable:

1. **port** specifies the port number on which the OCP plugin listens for connection requests
2. **radioHead-idle-timeout** determines the time duration (unit: milliseconds) for which a radio head has been idle before the idle event is triggered to perform health check
3. **ocp-version** specifies the OCP protocol version supported by the OCP plugin
4. **rpc-requests-quota** sets the maximum number of concurrent rpc requests allowed
5. **global-notification-quota** sets the maximum number of concurrent notifications allowed

Test Environment

The OCP Plugin project contains a simple OCP agent for testing purposes; the agent has been designed specifically to act as a fake radio head device, giving you an idea of what it would look like during the OCP handshake taking place between the OCP agent and OpenDaylight (OCP plugin).

To run the simple OCP agent, you have to first download its JAR file from OpenDaylight Nexus Repository.

```
wget https://nexus.opendaylight.org/content/repositories/opendaylight.release/org/
↪opendaylight/ocpplugin/simple-agent/0.1.0-Boron/simple-agent-0.1.0-Boron.jar
```

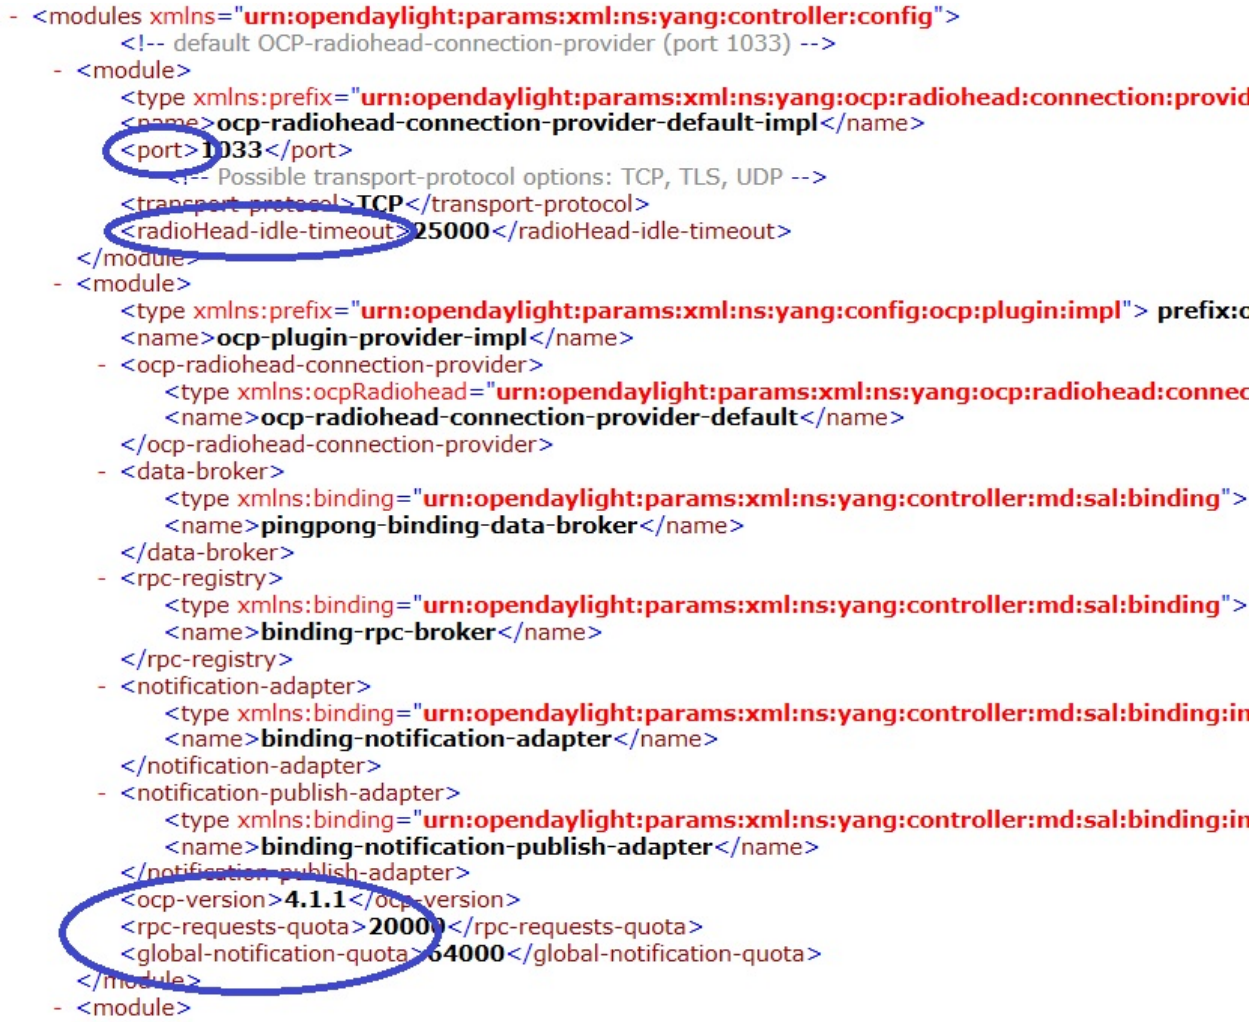

Fig. 75: OCP plugin configuration

Then run the agent with no arguments (assuming you already have JDK 1.8 or above installed) and it should display the usage that lists the expected arguments.

```
java -classpath simple-agent-0.1.0-Boron.jar org.opendaylight.ocpplugin.OcpAgent

Usage: java org.opendaylight.ocpplugin.OcpAgent <controller's ip address> <port␣
↪number> <vendor id> <serial number>
```

Here is an example:

```
java -classpath simple-agent-0.1.0-Boron.jar org.opendaylight.ocpplugin.OcpAgent 127.
↪0.0.1 1033 XYZ 123
```

Web / Graphical Interface

Once you enable the DLUX feature, you can access the Controller GUI using following URL.

```
http://<controller-ip>:8080/index.html
```

Expand Nodes. You should see all the radio head devices that are connected to the controller running at <controller-ip>.

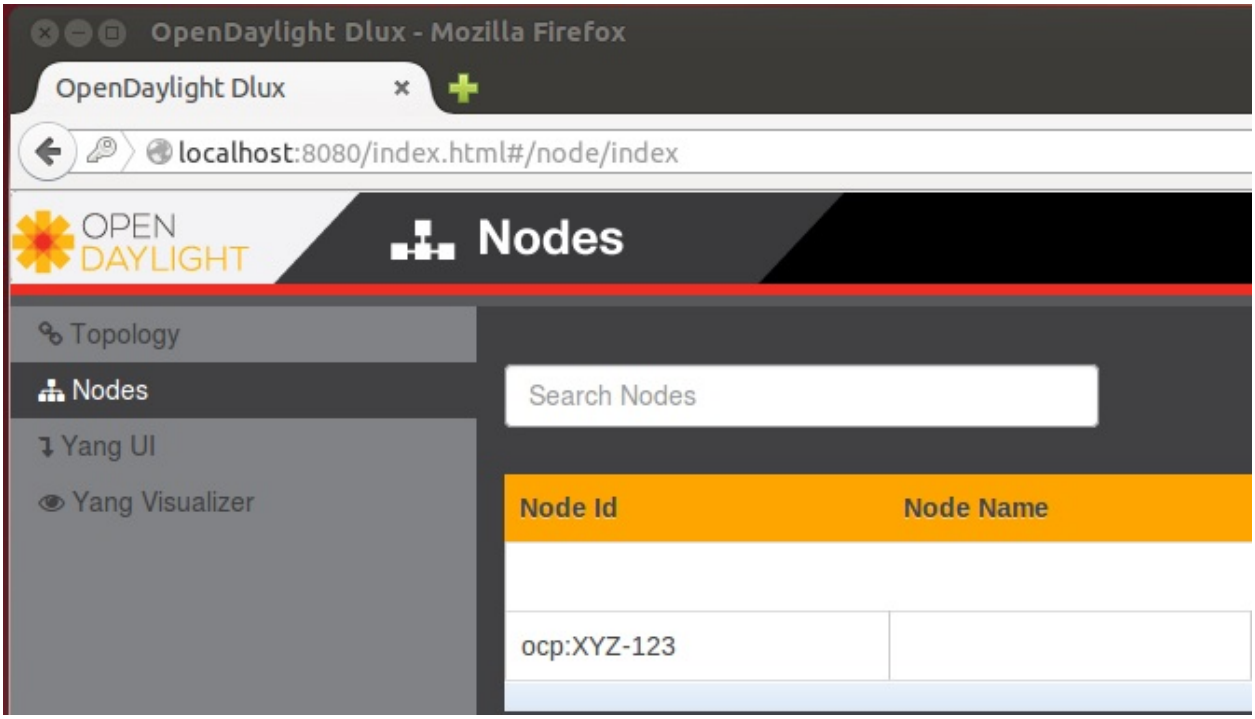

Fig. 76: DLUX Nodes

And expand Yang UI if you want to browse the various northbound APIs exposed by the OCP service.

For information on how to use these northbound APIs, please refer to the OCP Plugin Developer Guide.

Fig. 77: DLUX Yang UI

Programmatic Interface

The OCP Plugin project has implemented a complete set of the C&M operations (elementary functions) defined in the OCP specification, in the form of both northbound and southbound APIs, including:

- health-check
- set-time
- re-reset
- get-param
- modify-param
- create-obj
- delete-obj
- get-state
- modify-state
- get-fault

The API is documented in the OCP Plugin Developer Guide under the section Southbound API and Northbound API, respectively.

ODL-SDNi User Guide

Introduction

This user guide will help to setup the ODL-SDNi application.

Components

SDNiAggregator, SDNi API, SDNiWrapper, and SDNiUI are the four components in ODL-SDNi App:

- SDNiAggregator: Connects with switch, topology, hosttracker managers of controller to get the topology and other related data.
- SDNi REST API: It is a part of controller northbound, which gives the required information by quering SDNi-Aggregator through RESTCONF.
- SDNiWrapper: This component uses the SDNi REST API and gathers the information required to be shared among controllers.
- SDNiUI:This component displays all the SDN controllers which are connected to each other.

Troubleshooting

To work with multiple controllers, change some of the configuration in config.ini file. For example change the listening port of one controller to 6653 and other controller to 6663 in /root/controller/opendaylight/distribution/opendaylight/target/distribution.opendaylight-osgipackage/opendaylight/configuration/config.ini (i.e., of.listenPort=6653).

OpenFlow related system parameters.

TCP port on which the controller is listening (default 6633) of.listenPort=6653

OF-CONFIG User Guide

Overview

OF-CONFIG defines an OpenFlow switch as an abstraction called an OpenFlow Logical Switch. The OF-CONFIG protocol enables configuration of essential artifacts of an OpenFlow Logical Switch so that an OpenFlow controller can communicate and control the OpenFlow Logical switch via the OpenFlow protocol. OF-CONFIG introduces an operating context for one or more OpenFlow data paths called an OpenFlow Capable Switch for one or more switches. An OpenFlow Capable Switch is intended to be equivalent to an actual physical or virtual network element (e.g. an Ethernet switch) which is hosting one or more OpenFlow data paths by partitioning a set of OpenFlow related resources such as ports and queues among the hosted OpenFlow data paths. The OF-CONFIG protocol enables dynamic association of the OpenFlow related resources of an OpenFlow Capable Switch with specific OpenFlow Logical Switches which are being hosted on the OpenFlow Capable Switch. OF-CONFIG does not specify or report how the partitioning of resources on an OpenFlow Capable Switch is achieved. OF-CONFIG assumes that resources such as ports and queues are partitioned amongst multiple OpenFlow Logical Switches such that each OpenFlow Logical Switch can assume full control over the resources that is assigned to it.

How to start

- start OF-CONFIG feature as below:

```
feature:install odl-of-config-all
```

Configuration on the OVS supporting OF-CONFIG

Note: OVS is not supported by OF-CONFIG temporarily because the OpenDaylight version of OF-CONFIG is 1.2 while the OVS version of OF-CONFIG is not standard.

The introduction of configuring the OVS can be referred to:

https://github.com/openvswitch/of-config.

Connection Establishment between the Capable/Logical Switch and OF-CONFIG

The OF-CONFIG protocol is based on NETCONF. So the switches supporting OF-CONFIG can also access OpenDaylight using the functions provided by NETCONF. This is the preparation step before connecting to OF-CONFIG. How to access the switch to OpenDaylight using the NETCONF can be referred to the *NETCONF Southbound User Guide* or NETCONF Southbound examples on the wiki.

Now the switches supporting OF-CONFIG and they have connected to the controller using NETCONF as described in preparation phase. OF-CONFIG can check whether the switch can support OF-CONFIG by reading the capability list in NETCONF.

The OF-CONFIG will get the information of the capable switch and logical switch via the NETCONF connection, and creates separate topologies for the capable and logical switches in the OpenDaylight Topology module.

The Connection between the capable/logical switches and OF-CONFIG is finished.

Configuration On Capable Switch

Here is an example showing how to make the configuration to modify-controller-connection on the capable switch using OF-CONFIG. Other configurations can follow the same way of the example.

- Example: modify-controller-connection

Note: this configuration can execute via the NETCONF, which can be referred to the *NETCONF Southbound User Guide* or NETCONF Southbound examples on the wiki.

OpenFlow Plugin Project User Guide

Overview and Architecture

Overview and Architecture

Overview

OpenFlow is a vendor-neutral standard communications interface defined to enable interaction between the control and forwarding layers of an SDN architecture. The OpenFlow plugin project intends to develop a plugin to support implementations of the OpenFlow specification as it develops and evolves. Specifically the project has developed a plugin aiming to support OpenFlow 1.0 and 1.3.x. It can be extended to add support for subsequent OpenFlow specifications. The plugin is based on the Model Driven Service Abstraction Layer (MD-SAL) architecture (https://wiki.opendaylight.org/view/OpenDaylight_Controller:MD-SAL). This new OpenFlow 1.0/1.3 MD-SAL based plugin is distinct from the old OpenFlow 1.0 plugin which was based on the API driven SAL (AD-SAL) architecture.

Scope

- Southbound plugin and integration of OpenFlow 1.0/1.3.x library project
- Ongoing support and integration of the OpenFlow specification
- The plugin should be implemented in an easily extensible manner
- Protocol verification activities will be performed on supported OpenFlow specifications

Architecture and Design

Functionality

OpenFlow 1.3 Plugin will support the following functionality

- Connection Handling
- Session Management
- State Management.
- Error Handling.
- Mapping function(Infrastructure to OF structures).
- Connection establishment will be handled by OpenFlow library using opensource netty.io library.

- Message handling(Ex: Packet in).
- Event handling and propagation to upper layers.
- Plugin will support both MD-SAL and Hard SAL.
- Will be backward compatible with OF 1.0.

Activities in OF plugin module

- New OF plugin bundle will support both OF 1.0 and OF 1.3.
- Integration with OpenFlow library.
- Integration with corresponding MD-SAL infrastructure.
- Hard SAL will be supported as adapter on top of MD-SAL plugin.
- OF 1.3 and OF 1.0 plugin will be integrated as single bundle.

Design

Overall Architecture

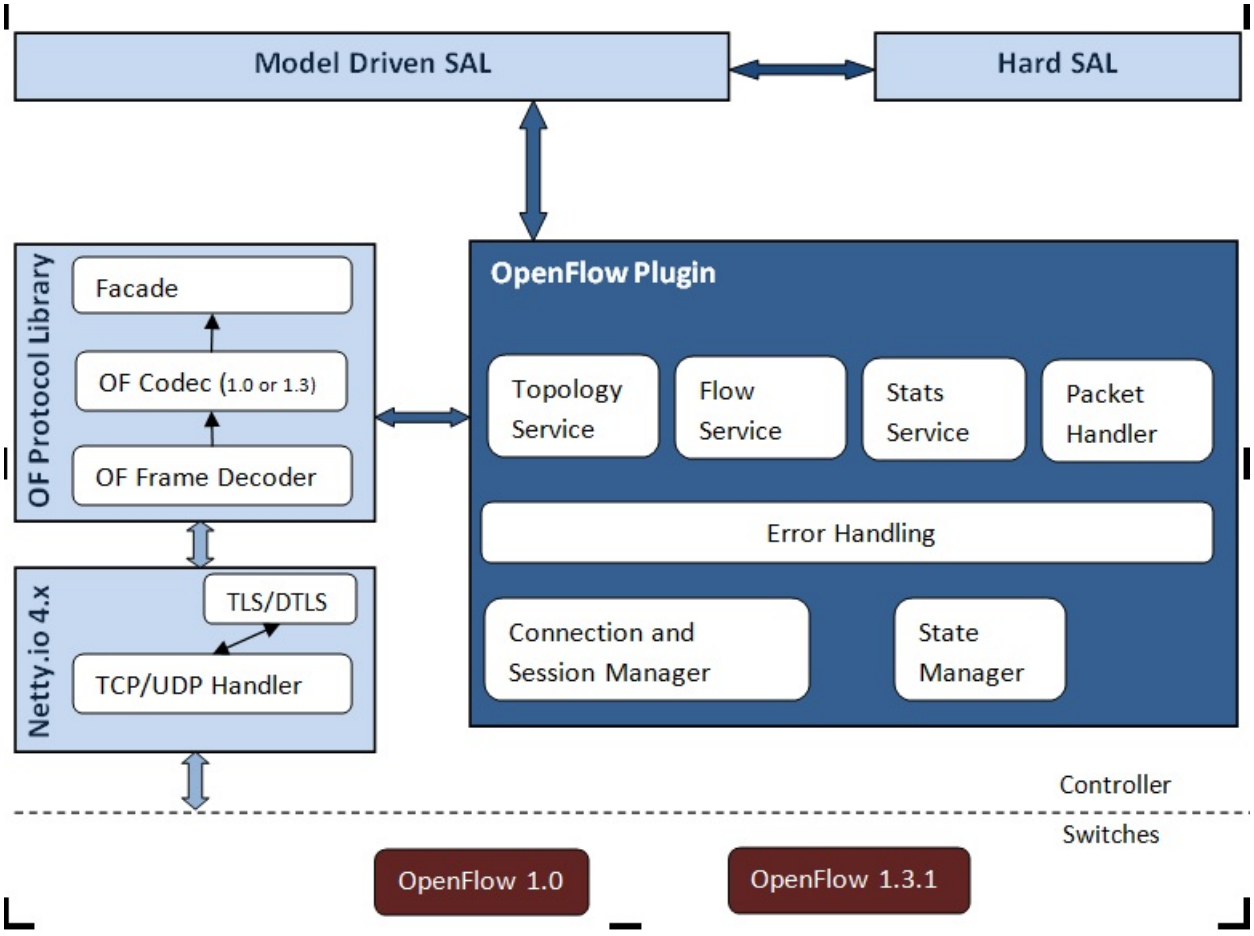

Fig. 78: overal architecture

Coverage

Intro

This page is to catalog the things that have been tested and confirmed to work:

Coverage

Coverage has been moved to a GoogleDoc Spreadsheet

OF 1.3 Considerations

The baseline model is a OF 1.3 model, and the coverage tables primarily deal with OF 1.3. However for OF 1.0, we have a column to indicate either N/A if it doesn't apply, or whether its been confirmed working.

OF 1.0 Considerations

OF 1.0 is being considered as a switch with: * 1 Table * 0 Groups * 0 Meters * 1 Instruction (Apply Actions) * and a limited vocabulary of matches and actions.

Tutorial / How-To

Running the controller with the new OpenFlow Plugin

How to start

There are all helium features (from features-openflowplugin) duplicated into features-openflowplugin-li. The duplicates got suffix *-li* and provide Lithium codebase functionality.

These are most used:

- odl-openflowplugin-app-lldp-speaker-li
- odl-openflowplugin-flow-services-rest-li
- odl-openflowplugin-drop-test-li

In case topology is required then the first one should be installed.

```
feature:install odl-openflowplugin-app-lldp-speaker-li
```

The Li-southbound currently provides:

- flow management
- group management
- meter management
- statistics polling

What to log

In order to see really low level messages enter these in karaf console:

```
log:set TRACE org.opendaylight.openflowplugin.openflow.md.core
log:set TRACE org.opendaylight.openflowplugin.impl
```

How enable topology

In order for topology to work (fill dataStore/operational with links) there must be LLDP responses delivered back to controller. This requires table-miss-entries. Table-miss-entry is a flow in table.id=0 with low priority, empty match and one output action = send to controller. Having this flow installed on every node will enable for gathering and exporting links between nodes into dataStore/operational. This is done if you use for example l2 switch application.

```
<?xml version="1.0" encoding="UTF-8" standalone="no"?>
<flow xmlns="urn:opendaylight:flow:inventory">
   <barrier>false</barrier>
   <cookie>54</cookie>
   <flags>SEND_FLOW_REM</flags>
   <flow-name>FooXf54</flow-name>
   <hard-timeout>0</hard-timeout>
   <id>4242</id>
   <idle-timeout>0</idle-timeout>
   <installHw>false</installHw>
   <instructions>
       <instruction>
           <apply-actions>
               <action>
                   <output-action>
                       <max-length>65535</max-length>
                       <output-node-connector>CONTROLLER</output-node-connector>
                   </output-action>
                   <order>0</order>
               </action>
           </apply-actions>
           <order>0</order>
       </instruction>
   </instructions>
   <match/>
   <priority>0</priority>
   <strict>false</strict>
   <table_id>0</table_id>
</flow>
```

Enable RESTCONF and Controller GUI

If you want to use RESTCONF with openflowplugin project, you have to install *odl-restconf* feature to enable that. To install *odl-restconf* feature run the following command

```
karaf#>feature:install odl-restconf
```

If you want to access the Controller GUI, you have to install *odl-dlux-core* feature to enable that. Run following command to install it

```
karaf#>feature:install odl-dlux-core
```

Once you enable the feature, access the Controller GUI using following URL

```
http://<controller-ip>:8181/dlux/index.html
```

OpenFlow 1.3 Enabled Software Switches / Environment

Getting Mininet with OF 1.3

Download Mininet VM Upgraded to OF 1.3 (or the newer mininet-2.1.0 with OVS-2.0 that works with VMware Player. For using this on VirtualBox, import this to VMware Player and then export the .vmdk) or you could build one yourself Openflow Protocol Library:OpenVirtualSwitch[Instructions for setting up Mininet with OF 1.3].

Installing under VirtualBox

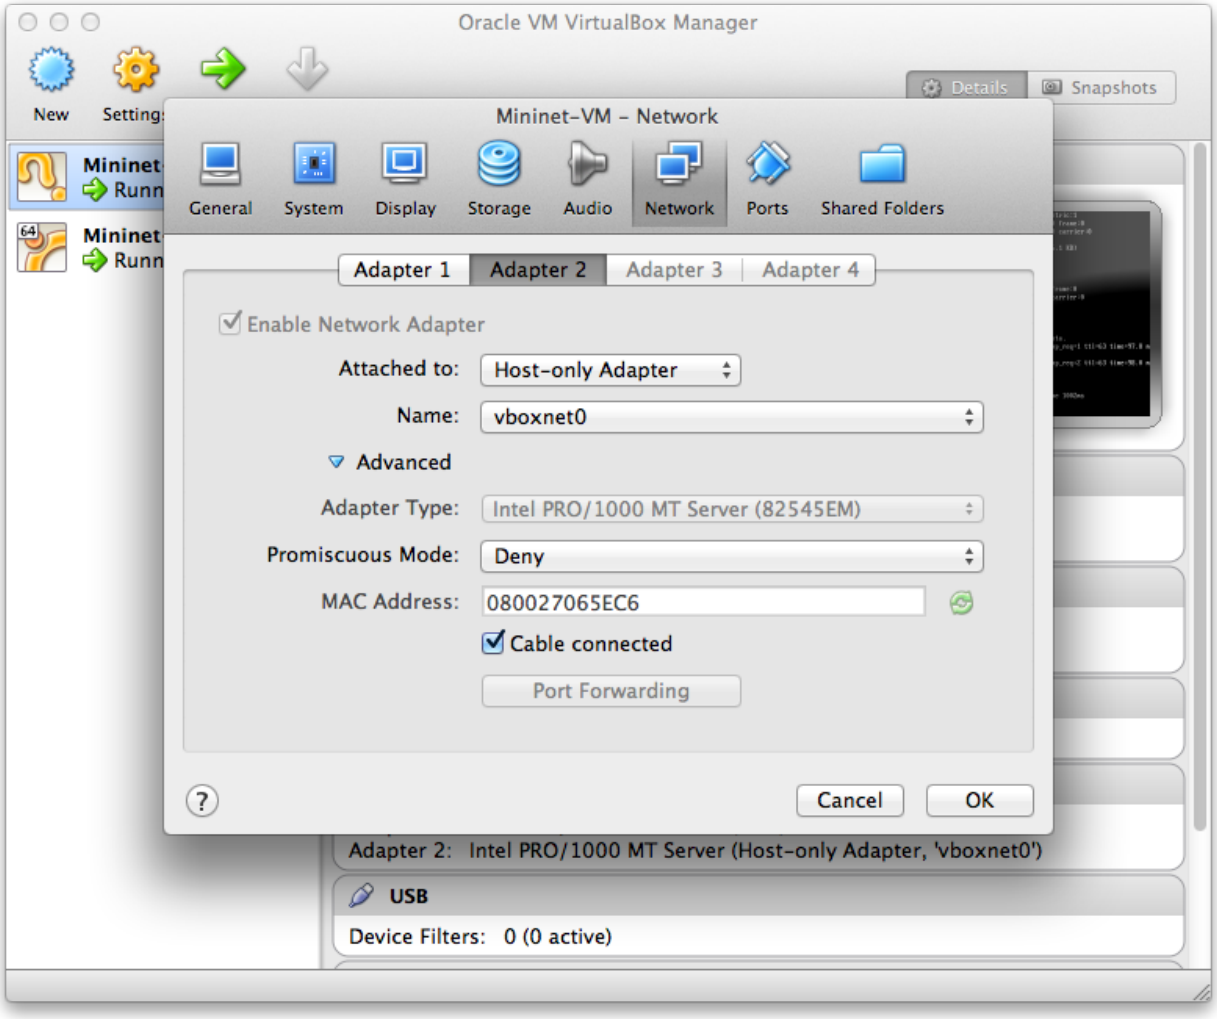

Fig. 79: configuring a host-only adapter

For whatever reason, at least on the Mac, NATed interfaces in VirtualBox don't actually seem to allow for connections from the host to the VM. Instead, you need to configure a host-only network and set it up. Do this by:

- Go to the VM's settings in VirtualBox then to network and add a second adapter attached to "Host-only Adapter" (see the screenshot to the right)
- Edit the /etc/network/interfaces file to configure the adapter properly by adding these two lines

```
auto eth1
iface eth1 inet dhcp
```

- Reboot the VM

At this point you should have two interfaces one which gives you NATed access to the internet and another that gives you access between your mac and the VMs. At least for me, the NATed interface gets a 10.0.2.x address and the the host-only interface gets a 192.168.56.x address.

Your simplest choice: Use Vagrant

Download Virtual Box and install it Download Vagrant and install it

```
cd openflowplugin/vagrant/mininet-2.1.0-of-1.3/
vagrant up
vagrant ssh
```

This will leave you sshed into a fully provisioned Ubuntu Trusty box with mininet-2.1.0 and OVS 2.0 patches to work with OF 1.3.

Setup CPqD Openflow 1.3 Soft Switch

Latest version of Openvswitch (v2.0.0) doesn't support all the openflow 1.3 features, e.g group multipart statistics request. Alternate options is CPqD Openflow 1.3 soft switch, It supports most of the openflow 1.3 features.

- You can setup the switch as per the instructions given on the following URL

https://github.com/CPqD/ofsoftswitch13

- Fire following command to start the switch

Start the datapath :

```
$ sudo udatapath/ofdatapath --datapath-id=<dpid> --interfaces=<if-list> ptcp:<port>
 e.g $ sudo udatapath/ofdatapath --datapath-id=000000000001 --interfaces=ethX␣
↪ptcp:6680
```

ethX should not be associated with ip address and ipv6 should be disabled on it. If you are installing the switch on your local machine, you can use following command (for Ubuntu) to create virtual interface.

```
ip link add link ethX address 00:19:d1:29:d2:58 macvlan0 type macvlan
```

ethX - Any existing interface.

Or if you are using mininet VM for installing this switch, you can simply add one more adaptor to your VM.

Start Openflow protocol agent:

```
$secchan/ofprotocol tcp:<switch-host>:<switch-port> tcp:<ctrl-host>:<ctrl-port>
 e.g $secchan/ofprotocol tcp:127.0.0.1:6680 tcp:127.0.0.1:6653
```

Commands to add entries to various tables of the switch

- Add meter

```
$utilities/dpctl tcp:<switch-host>:<switch-port> meter-mod cmd=add,meter=1␣
↪drop:rate=50
```

- Add Groups

```
$utilities/dpctl tcp:127.0.0.1:6680 group-mod cmd=add,type=all,group=1
```

```
$utilities/dpctl tcp:127.0.0.1:6680 group-mod cmd=add,type=sel,group=2 weight=10␣
↪output:1
```

- Create queue

```
$utilities/dpctl tcp:<ip>:<switch port> queue-mod <port-number> <queue-number>
↪<minimum-bandwidth>
  e.g - $utilities/dpctl tcp:127.0.0.1:6680 queue-mod 1 1 23
```

"dpctl" –help is not very intuitive, so please keep adding any new command you figured out while your experiment with the switch.

Using the built-in Wireshark

Mininet comes with pre-installed Wireshark, but for some reason it does not include the Openflow protocol dissector. You may want to get and install it in the */.wireshark/plugins/* directory.

First login to your mininet VM

```
ssh mininet@<your mininet vm ip> -X
```

The -X option in ssh will enable x-session over ssh so that the wireshark window can be shown on your host machine's display. when prompted, enter the password (mininet).

From the mininet vm shell, set the wireshark capture privileges (http://wiki.wireshark.org/CaptureSetup/CapturePrivileges):

```
sudo chgrp mininet /usr/bin/dumpcap
sudo chmod 754 /usr/bin/dumpcap
sudo setcap 'CAP_NET_RAW+eip CAP_NET_ADMIN+eip' /usr/bin/dumpcap
```

Finally, start wireshark:

```
wireshark
```

The wireshark window should show up.

To see only Openflow packets, you may want to apply the following filter in the Filter window:

```
tcp.port == 6633 and tcp.flags.push == 1
```

Start the capture on *any* port.

Running Mininet with OF 1.3

From within the Mininet VM, run:

```
sudo mn --topo single,3  --controller 'remote,ip=<your controller ip>,port=6653' --
↪switch ovsk,protocols=OpenFlow13
```

End to End Inventory

Introduction

The purpose of this page is to walk you through how to see the Inventory Manager working end to end with the openflowplugin using OpenFlow 1.3.

Basically, you will learn how to:

1. Run the Base/Virtualization/Service provider Edition with the new openflowplugin: OpenDaylight_OpenFlow_Plugin::Running_controller_with_the_new_OF_plugin[Running the controller with the new OpenFlow Plugin]
2. Start mininet to use OF 1.3: OpenDaylight_OpenFlow_Plugin::Test_Environment[OpenFlow 1.3 Enabled Software Switches / Environment]
3. Use RESTCONF to see the nodes appear in inventory.

Restconf for Inventory

The REST url for listing all the nodes is:

```
http://localhost:8181/restconf/operational/opendaylight-inventory:nodes/
```

You will need to set the Accept header:

```
Accept: application/xml
```

You will also need to use HTTP Basic Auth with username: admin password: admin.

Alternately, if you have a node's id you can address it as

```
http://localhost:8181/restconf/operational/opendaylight-inventory:nodes/node/<id>
```

for example

```
http://localhost:8181/restconf/operational/opendaylight-inventory:nodes/node/
↪openflow:1
```

How to hit RestConf with Postman

Install Postman for Chrome

In the chrome browser bar enter

```
chrome://apps/
```

And click on Postman.

Enter the URL. Click on the Headers button on the far right. Enter the Accept: header. Click on the Basic Auth Tab at the top and setup the username and password. Send.

Known Bug

If you have not had any switches come up, and though no children for http://localhost:8080/restconf/datastore/opendaylight-inventory:nodes/ and exception will be thrown. I'm pretty sure I know how to fix this bug, just need to get to it :)

End to End Flows

Instructions

Learn End to End for Inventory

See *End to End Inventory*

Check inventory

- Run mininet with support for OF 1.3 as described in *End to End Inventory*
- Make sure you see the openflow:1 node come up as described in *End to End Inventory*

Flow Strategy

Current way to flush a flow to switch looks like this:

1. Create MD-SAL modeled flow and commit it into dataStore using two phase commit MD-SAL FAQ
2. FRM gets notified and invokes corresponding rpc (addFlow) on particular service provider (if suitable provider for given node registered)
3. The provider (plugin in this case) transforms MD-SAL modeled flow into OF-API modeled flow
4. OF-API modeled flow is then flushed into OFLibrary
5. OFLibrary encodes flow into particular version of wire protocol and sends it to particular switch
6. Check on mininet side if flow is set

Push your flow

- With PostMan:
  - Set headers:
    - Content-Type: application/xml
    - Accept: application/xml
    - Authentication
  - Use URL: "http://<controller IP>:8181/restconf/config/opendaylight-inventory:nodes/node/openflow:1/table/0/flow/1"
  - PUT
  - Use Body:

```
<?xml version="1.0" encoding="UTF-8" standalone="no"?>
<flow xmlns="urn:opendaylight:flow:inventory">
    <priority>2</priority>
    <flow-name>Foo</flow-name>
    <match>
        <ethernet-match>
            <ethernet-type>
                <type>2048</type>
            </ethernet-type>
        </ethernet-match>
        <ipv4-destination>10.0.10.2/24</ipv4-destination>
    </match>
    <id>1</id>
    <table_id>0</table_id>
    <instructions>
        <instruction>
            <order>0</order>
            <apply-actions>
                <action>
                   <order>0</order>
                   <dec-nw-ttl/>
                </action>
            </apply-actions>
        </instruction>
    </instructions>
</flow>
```

***Note**: If you want to try a different flow id or a different table, make sure the URL and the body stay in sync. For example, if you wanted to try: table 2 flow 20 you'd change the URL to:

"http://<controller IP>:8181/restconf/config/opendaylight-inventory:nodes/node/openflow:1/table/2/flow/20"

but you would also need to update the 20 and 2 in the body of the XML.

Other caveat, we have a known bug with updates, so please only write to a given flow id and table id on a given node once at this time until we resolve it. Or you can use the DELETE method with the same URL in PostMan to delete the flow information on switch and controller cache.

Check for your flow on the switch

- See your flow on your mininet:

```
mininet@mininet-vm:~$ sudo ovs-ofctl -O OpenFlow13 dump-flows s1
OFPST_FLOW reply (OF1.3) (xid=0x2):
cookie=0x0, duration=7.325s, table=0, n_packets=0, n_bytes=0, idle_timeout=300, hard_
↪timeout=600, send_flow_rem priority=2,ip,nw_dst=10.0.10.0/24 actions=dec_ttl
```

If you want to see the above information from the mininet prompt - use "sh" instead of "sudo" i.e. use "sh ovs-ofctl -O OpenFlow13 dump-flows s1".

Check for your flow in the controller config via RESTCONF

- See your configured flow in POSTMAN with
  - URL http://<controller IP>:8181/restconf/operational/opendaylight-inventory:nodes/node/openflow:1/table/0/
  - GET

- You no longer need to set Accept header

Return Response:


```
{
  "flow-node-inventory:table": [
    {
      "flow-node-inventory:id": 0,
      "flow-node-inventory:flow": [
        {
          "flow-node-inventory:priority": 1,
          "flow-node-inventory:id": "10b1a23c-5299-4f7b-83d6-563bab472754",
          "flow-node-inventory:table_id": 0,
          "flow-node-inventory:hard-timeout": 0,
          "flow-node-inventory:idle-timeout": 0,
          "flow-node-inventory:instructions": {
            "flow-node-inventory:instruction": [
              {
                "flow-node-inventory:apply-actions": {
                  "flow-node-inventory:action": [
                    {
                      "flow-node-inventory:output-action": {
                        "flow-node-inventory:output-node-connector": "openflow:1:1"
                      },
                      "flow-node-inventory:order": 0
                    }
                  ]
                },
                "flow-node-inventory:order": 0
              }
            ]
          },
          "flow-node-inventory:match": {
            "flow-node-inventory:ethernet-match": {
              "flow-node-inventory:ethernet-type": {
                "flow-node-inventory:type": 2048
              }
            },
            "flow-node-inventory:ipv4-destination": "10.0.0.2"
          },
          "flow-node-inventory:cookie": 0
        },
        {
          "flow-node-inventory:priority": 1,
          "flow-node-inventory:id": "020bf359-1299-4da6-b4f7-368bd83b5841",
          "flow-node-inventory:table_id": 0,
          "flow-node-inventory:hard-timeout": 0,
          "flow-node-inventory:idle-timeout": 0,
          "flow-node-inventory:instructions": {
            "flow-node-inventory:instruction": [
              {
                "flow-node-inventory:apply-actions": {
                  "flow-node-inventory:action": [
                    {
                      "flow-node-inventory:output-action": {
                        "flow-node-inventory:output-node-connector": "openflow:1:1"
                      },
                      "flow-node-inventory:order": 0
```


(continues on next page)

(continued from previous page)


```
                    }
                  ]
                },
                "flow-node-inventory:order": 0
              }
            ]
          },
          "flow-node-inventory:match": {
            "flow-node-inventory:ethernet-match": {
              "flow-node-inventory:ethernet-type": {
                "flow-node-inventory:type": 2048
              }
            },
            "flow-node-inventory:ipv4-destination": "10.0.0.1"
          },
          "flow-node-inventory:cookie": 0
        },
        {
          "flow-node-inventory:priority": 1,
          "flow-node-inventory:id": "42172bfc-9142-4a92-9e90-ee62529b1e85",
          "flow-node-inventory:table_id": 0,
          "flow-node-inventory:hard-timeout": 0,
          "flow-node-inventory:idle-timeout": 0,
          "flow-node-inventory:instructions": {
            "flow-node-inventory:instruction": [
              {
                "flow-node-inventory:apply-actions": {
                  "flow-node-inventory:action": [
                    {
                      "flow-node-inventory:output-action": {
                        "flow-node-inventory:output-node-connector": "openflow:1:1"
                      },
                      "flow-node-inventory:order": 0
                    }
                  ]
                },
                "flow-node-inventory:order": 0
              }
            ]
          },
          "flow-node-inventory:match": {
            "flow-node-inventory:ethernet-match": {
              "flow-node-inventory:ethernet-type": {
                "flow-node-inventory:type": 2048
              }
            },
            "flow-node-inventory:ipv4-destination": "10.0.0.3"
          },
          "flow-node-inventory:cookie": 0
        },
        {
          "flow-node-inventory:priority": 1,
          "flow-node-inventory:id": "99bf566e-89f3-4c6f-ae9e-e26012ceb1e4",
          "flow-node-inventory:table_id": 0,
          "flow-node-inventory:hard-timeout": 0,
          "flow-node-inventory:idle-timeout": 0,
          "flow-node-inventory:instructions": {
```


(continues on next page)

(continued from previous page)


```
            "flow-node-inventory:instruction": [
              {
                "flow-node-inventory:apply-actions": {
                  "flow-node-inventory:action": [
                    {
                      "flow-node-inventory:output-action": {
                        "flow-node-inventory:output-node-connector": "openflow:1:1"
                      },
                      "flow-node-inventory:order": 0
                    }
                  ]
                },
                "flow-node-inventory:order": 0
              }
            ]
          },
          "flow-node-inventory:match": {
            "flow-node-inventory:ethernet-match": {
              "flow-node-inventory:ethernet-type": {
                "flow-node-inventory:type": 2048
              }
            },
            "flow-node-inventory:ipv4-destination": "10.0.0.4"
          },
          "flow-node-inventory:cookie": 0
        },
        {
          "flow-node-inventory:priority": 1,
          "flow-node-inventory:id": "019dcc2e-5b4f-44f0-90cc-de490294b862",
          "flow-node-inventory:table_id": 0,
          "flow-node-inventory:hard-timeout": 0,
          "flow-node-inventory:idle-timeout": 0,
          "flow-node-inventory:instructions": {
            "flow-node-inventory:instruction": [
              {
                "flow-node-inventory:apply-actions": {
                  "flow-node-inventory:action": [
                    {
                      "flow-node-inventory:output-action": {
                        "flow-node-inventory:output-node-connector": "openflow:1:2"
                      },
                      "flow-node-inventory:order": 0
                    }
                  ]
                },
                "flow-node-inventory:order": 0
              }
            ]
          },
          "flow-node-inventory:match": {
            "flow-node-inventory:ethernet-match": {
              "flow-node-inventory:ethernet-type": {
                "flow-node-inventory:type": 2048
              }
            },
            "flow-node-inventory:ipv4-destination": "10.0.0.5"
          },
```


(continues on next page)

(continued from previous page)


```
          "flow-node-inventory:cookie": 0
        },
        {
          "flow-node-inventory:priority": 1,
          "flow-node-inventory:id": "968cf81e-3f16-42f1-8b16-d01ff719c63c",
          "flow-node-inventory:table_id": 0,
          "flow-node-inventory:hard-timeout": 0,
          "flow-node-inventory:idle-timeout": 0,
          "flow-node-inventory:instructions": {
            "flow-node-inventory:instruction": [
              {
                "flow-node-inventory:apply-actions": {
                  "flow-node-inventory:action": [
                    {
                      "flow-node-inventory:output-action": {
                        "flow-node-inventory:output-node-connector": "openflow:1:2"
                      },
                      "flow-node-inventory:order": 0
                    }
                  ]
                },
                "flow-node-inventory:order": 0
              }
            ]
          },
          "flow-node-inventory:match": {
            "flow-node-inventory:ethernet-match": {
              "flow-node-inventory:ethernet-type": {
                "flow-node-inventory:type": 2048
              }
            },
            "flow-node-inventory:ipv4-destination": "10.0.0.8"
          },
          "flow-node-inventory:cookie": 0
        },
        {
          "flow-node-inventory:priority": 1,
          "flow-node-inventory:id": "1c14ea3c-9dcc-4434-b566-7e99033ea252",
          "flow-node-inventory:table_id": 0,
          "flow-node-inventory:hard-timeout": 0,
          "flow-node-inventory:idle-timeout": 0,
          "flow-node-inventory:instructions": {
            "flow-node-inventory:instruction": [
              {
                "flow-node-inventory:apply-actions": {
                  "flow-node-inventory:action": [
                    {
                      "flow-node-inventory:output-action": {
                        "flow-node-inventory:output-node-connector": "openflow:1:2"
                      },
                      "flow-node-inventory:order": 0
                    }
                  ]
                },
                "flow-node-inventory:order": 0
              }
            ]
```


(continues on next page)

(continued from previous page)


```
          },
          "flow-node-inventory:match": {
            "flow-node-inventory:ethernet-match": {
              "flow-node-inventory:ethernet-type": {
                "flow-node-inventory:type": 2048
              }
            },
            "flow-node-inventory:ipv4-destination": "10.0.0.6"
          },
          "flow-node-inventory:cookie": 0
        },
        {
          "flow-node-inventory:priority": 1,
          "flow-node-inventory:id": "ed9deeb2-be8f-4b84-bcd8-9d12049383d6",
          "flow-node-inventory:table_id": 0,
          "flow-node-inventory:hard-timeout": 0,
          "flow-node-inventory:idle-timeout": 0,
          "flow-node-inventory:instructions": {
            "flow-node-inventory:instruction": [
              {
                "flow-node-inventory:apply-actions": {
                  "flow-node-inventory:action": [
                    {
                      "flow-node-inventory:output-action": {
                        "flow-node-inventory:output-node-connector": "openflow:1:2"
                      },
                      "flow-node-inventory:order": 0
                    }
                  ]
                },
                "flow-node-inventory:order": 0
              }
            ]
          },
          "flow-node-inventory:match": {
            "flow-node-inventory:ethernet-match": {
              "flow-node-inventory:ethernet-type": {
                "flow-node-inventory:type": 2048
              }
            },
            "flow-node-inventory:ipv4-destination": "10.0.0.7"
          },
          "flow-node-inventory:cookie": 0
        }
      ]
    }
  ]
}
```


Look for your flow stats in the controller operational data via

RESTCONF

- See your operational flow stats in POSTMAN with
  - URL "http://<controller IP>:8181/restconf/operational/opendaylight-inventory:nodes/node/openflow:1/table/0/"

- GET

Return Response:


```
{
  "flow-node-inventory:table": [
    {
      "flow-node-inventory:id": 0,
      "flow-node-inventory:flow": [
        {
          "flow-node-inventory:id": "10b1a23c-5299-4f7b-83d6-563bab472754",
          "opendaylight-flow-statistics:flow-statistics": {
            "opendaylight-flow-statistics:cookie": 0,
            "opendaylight-flow-statistics:duration": {
              "opendaylight-flow-statistics:nanosecond": 886000000,
              "opendaylight-flow-statistics:second": 2707
            },
            "opendaylight-flow-statistics:hard-timeout": 0,
            "opendaylight-flow-statistics:byte-count": 784,
            "opendaylight-flow-statistics:match": {
              "opendaylight-flow-statistics:ethernet-match": {
                "opendaylight-flow-statistics:ethernet-type": {
                  "opendaylight-flow-statistics:type": 2048
                }
              },
              "opendaylight-flow-statistics:ipv4-destination": "10.0.0.2/32"
            },
            "opendaylight-flow-statistics:priority": 1,
            "opendaylight-flow-statistics:packet-count": 8,
            "opendaylight-flow-statistics:table_id": 0,
            "opendaylight-flow-statistics:idle-timeout": 0,
            "opendaylight-flow-statistics:instructions": {
              "opendaylight-flow-statistics:instruction": [
                {
                  "opendaylight-flow-statistics:order": 0,
                  "opendaylight-flow-statistics:apply-actions": {
                    "opendaylight-flow-statistics:action": [
                      {
                        "opendaylight-flow-statistics:order": 0,
                        "opendaylight-flow-statistics:output-action": {
                          "opendaylight-flow-statistics:output-node-connector": "1",
                          "opendaylight-flow-statistics:max-length": 0
                        }
                      }
                    ]
                  }
                }
              ]
            }
          }
        },
        {
          "flow-node-inventory:id": "020bf359-1299-4da6-b4f7-368bd83b5841",
          "opendaylight-flow-statistics:flow-statistics": {
            "opendaylight-flow-statistics:cookie": 0,
            "opendaylight-flow-statistics:duration": {
              "opendaylight-flow-statistics:nanosecond": 826000000,
              "opendaylight-flow-statistics:second": 2711
```


(continues on next page)

(continued from previous page)


```
            },
            "opendaylight-flow-statistics:hard-timeout": 0,
            "opendaylight-flow-statistics:byte-count": 1568,
            "opendaylight-flow-statistics:match": {
              "opendaylight-flow-statistics:ethernet-match": {
                "opendaylight-flow-statistics:ethernet-type": {
                  "opendaylight-flow-statistics:type": 2048
                }
              },
              "opendaylight-flow-statistics:ipv4-destination": "10.0.0.1/32"
            },
            "opendaylight-flow-statistics:priority": 1,
            "opendaylight-flow-statistics:packet-count": 16,
            "opendaylight-flow-statistics:table_id": 0,
            "opendaylight-flow-statistics:idle-timeout": 0,
            "opendaylight-flow-statistics:instructions": {
              "opendaylight-flow-statistics:instruction": [
                {
                  "opendaylight-flow-statistics:order": 0,
                  "opendaylight-flow-statistics:apply-actions": {
                    "opendaylight-flow-statistics:action": [
                      {
                        "opendaylight-flow-statistics:order": 0,
                        "opendaylight-flow-statistics:output-action": {
                          "opendaylight-flow-statistics:output-node-connector": "1",
                          "opendaylight-flow-statistics:max-length": 0
                        }
                      }
                    ]
                  }
                }
              ]
            }
          }
        },
        {
          "flow-node-inventory:id": "42172bfc-9142-4a92-9e90-ee62529b1e85",
          "opendaylight-flow-statistics:flow-statistics": {
            "opendaylight-flow-statistics:cookie": 0,
            "opendaylight-flow-statistics:duration": {
              "opendaylight-flow-statistics:nanosecond": 548000000,
              "opendaylight-flow-statistics:second": 2708
            },
            "opendaylight-flow-statistics:hard-timeout": 0,
            "opendaylight-flow-statistics:byte-count": 784,
            "opendaylight-flow-statistics:match": {
              "opendaylight-flow-statistics:ethernet-match": {
                "opendaylight-flow-statistics:ethernet-type": {
                  "opendaylight-flow-statistics:type": 2048
                }
              },
              "opendaylight-flow-statistics:ipv4-destination": "10.0.0.3/32"
            },
            "opendaylight-flow-statistics:priority": 1,
            "opendaylight-flow-statistics:packet-count": 8,
            "opendaylight-flow-statistics:table_id": 0,
            "opendaylight-flow-statistics:idle-timeout": 0,
```


(continues on next page)

(continued from previous page)


```
            "opendaylight-flow-statistics:instructions": {
              "opendaylight-flow-statistics:instruction": [
                {
                  "opendaylight-flow-statistics:order": 0,
                  "opendaylight-flow-statistics:apply-actions": {
                    "opendaylight-flow-statistics:action": [
                      {
                        "opendaylight-flow-statistics:order": 0,
                        "opendaylight-flow-statistics:output-action": {
                          "opendaylight-flow-statistics:output-node-connector": "1",
                          "opendaylight-flow-statistics:max-length": 0
                        }
                      }
                    ]
                  }
                }
              ]
            }
          }
        },
        {
          "flow-node-inventory:id": "99bf566e-89f3-4c6f-ae9e-e26012ceb1e4",
          "opendaylight-flow-statistics:flow-statistics": {
            "opendaylight-flow-statistics:cookie": 0,
            "opendaylight-flow-statistics:duration": {
              "opendaylight-flow-statistics:nanosecond": 296000000,
              "opendaylight-flow-statistics:second": 2710
            },
            "opendaylight-flow-statistics:hard-timeout": 0,
            "opendaylight-flow-statistics:byte-count": 1274,
            "opendaylight-flow-statistics:match": {
              "opendaylight-flow-statistics:ethernet-match": {
                "opendaylight-flow-statistics:ethernet-type": {
                  "opendaylight-flow-statistics:type": 2048
                }
              },
              "opendaylight-flow-statistics:ipv4-destination": "10.0.0.4/32"
            },
            "opendaylight-flow-statistics:priority": 1,
            "opendaylight-flow-statistics:packet-count": 13,
            "opendaylight-flow-statistics:table_id": 0,
            "opendaylight-flow-statistics:idle-timeout": 0,
            "opendaylight-flow-statistics:instructions": {
              "opendaylight-flow-statistics:instruction": [
                {
                  "opendaylight-flow-statistics:order": 0,
                  "opendaylight-flow-statistics:apply-actions": {
                    "opendaylight-flow-statistics:action": [
                      {
                        "opendaylight-flow-statistics:order": 0,
                        "opendaylight-flow-statistics:output-action": {
                          "opendaylight-flow-statistics:output-node-connector": "1",
                          "opendaylight-flow-statistics:max-length": 0
                        }
                      }
                    ]
                  }
```


(continues on next page)

(continued from previous page)


```
            }
          ]
        }
      }
    },
    {
      "flow-node-inventory:id": "019dcc2e-5b4f-44f0-90cc-de490294b862",
      "opendaylight-flow-statistics:flow-statistics": {
        "opendaylight-flow-statistics:cookie": 0,
        "opendaylight-flow-statistics:duration": {
          "opendaylight-flow-statistics:nanosecond": 392000000,
          "opendaylight-flow-statistics:second": 2711
        },
        "opendaylight-flow-statistics:hard-timeout": 0,
        "opendaylight-flow-statistics:byte-count": 1470,
        "opendaylight-flow-statistics:match": {
          "opendaylight-flow-statistics:ethernet-match": {
            "opendaylight-flow-statistics:ethernet-type": {
              "opendaylight-flow-statistics:type": 2048
            }
          },
          "opendaylight-flow-statistics:ipv4-destination": "10.0.0.5/32"
        },
        "opendaylight-flow-statistics:priority": 1,
        "opendaylight-flow-statistics:packet-count": 15,
        "opendaylight-flow-statistics:table_id": 0,
        "opendaylight-flow-statistics:idle-timeout": 0,
        "opendaylight-flow-statistics:instructions": {
          "opendaylight-flow-statistics:instruction": [
            {
              "opendaylight-flow-statistics:order": 0,
              "opendaylight-flow-statistics:apply-actions": {
                "opendaylight-flow-statistics:action": [
                  {
                    "opendaylight-flow-statistics:order": 0,
                    "opendaylight-flow-statistics:output-action": {
                      "opendaylight-flow-statistics:output-node-connector": "2",
                      "opendaylight-flow-statistics:max-length": 0
                    }
                  }
                ]
              }
            }
          ]
        }
      }
    },
    {
      "flow-node-inventory:id": "968cf81e-3f16-42f1-8b16-d01ff719c63c",
      "opendaylight-flow-statistics:flow-statistics": {
        "opendaylight-flow-statistics:cookie": 0,
        "opendaylight-flow-statistics:duration": {
          "opendaylight-flow-statistics:nanosecond": 344000000,
          "opendaylight-flow-statistics:second": 2707
        },
        "opendaylight-flow-statistics:hard-timeout": 0,
        "opendaylight-flow-statistics:byte-count": 784,
```


(continues on next page)

(continued from previous page)


```
            "opendaylight-flow-statistics:match": {
              "opendaylight-flow-statistics:ethernet-match": {
                "opendaylight-flow-statistics:ethernet-type": {
                  "opendaylight-flow-statistics:type": 2048
                }
              },
              "opendaylight-flow-statistics:ipv4-destination": "10.0.0.8/32"
            },
            "opendaylight-flow-statistics:priority": 1,
            "opendaylight-flow-statistics:packet-count": 8,
            "opendaylight-flow-statistics:table_id": 0,
            "opendaylight-flow-statistics:idle-timeout": 0,
            "opendaylight-flow-statistics:instructions": {
              "opendaylight-flow-statistics:instruction": [
                {
                  "opendaylight-flow-statistics:order": 0,
                  "opendaylight-flow-statistics:apply-actions": {
                    "opendaylight-flow-statistics:action": [
                      {
                        "opendaylight-flow-statistics:order": 0,
                        "opendaylight-flow-statistics:output-action": {
                          "opendaylight-flow-statistics:output-node-connector": "2",
                          "opendaylight-flow-statistics:max-length": 0
                        }
                      }
                    ]
                  }
                }
              ]
            }
          }
        },
        {
          "flow-node-inventory:id": "ed9deeb2-be8f-4b84-bcd8-9d12049383d6",
          "opendaylight-flow-statistics:flow-statistics": {
            "opendaylight-flow-statistics:cookie": 0,
            "opendaylight-flow-statistics:duration": {
              "opendaylight-flow-statistics:nanosecond": 577000000,
              "opendaylight-flow-statistics:second": 2706
            },
            "opendaylight-flow-statistics:hard-timeout": 0,
            "opendaylight-flow-statistics:byte-count": 784,
            "opendaylight-flow-statistics:match": {
              "opendaylight-flow-statistics:ethernet-match": {
                "opendaylight-flow-statistics:ethernet-type": {
                  "opendaylight-flow-statistics:type": 2048
                }
              },
              "opendaylight-flow-statistics:ipv4-destination": "10.0.0.7/32"
            },
            "opendaylight-flow-statistics:priority": 1,
            "opendaylight-flow-statistics:packet-count": 8,
            "opendaylight-flow-statistics:table_id": 0,
            "opendaylight-flow-statistics:idle-timeout": 0,
            "opendaylight-flow-statistics:instructions": {
              "opendaylight-flow-statistics:instruction": [
                {
```


(continues on next page)

(continued from previous page)


```
                "opendaylight-flow-statistics:order": 0,
                "opendaylight-flow-statistics:apply-actions": {
                  "opendaylight-flow-statistics:action": [
                    {
                      "opendaylight-flow-statistics:order": 0,
                      "opendaylight-flow-statistics:output-action": {
                        "opendaylight-flow-statistics:output-node-connector": "2",
                        "opendaylight-flow-statistics:max-length": 0
                      }
                    }
                  ]
                }
              }
            ]
          }
        }
      },
      {
        "flow-node-inventory:id": "1c14ea3c-9dcc-4434-b566-7e99033ea252",
        "opendaylight-flow-statistics:flow-statistics": {
          "opendaylight-flow-statistics:cookie": 0,
          "opendaylight-flow-statistics:duration": {
            "opendaylight-flow-statistics:nanosecond": 659000000,
            "opendaylight-flow-statistics:second": 2705
          },
          "opendaylight-flow-statistics:hard-timeout": 0,
          "opendaylight-flow-statistics:byte-count": 784,
          "opendaylight-flow-statistics:match": {
            "opendaylight-flow-statistics:ethernet-match": {
              "opendaylight-flow-statistics:ethernet-type": {
                "opendaylight-flow-statistics:type": 2048
              }
            },
            "opendaylight-flow-statistics:ipv4-destination": "10.0.0.6/32"
          },
          "opendaylight-flow-statistics:priority": 1,
          "opendaylight-flow-statistics:packet-count": 8,
          "opendaylight-flow-statistics:table_id": 0,
          "opendaylight-flow-statistics:idle-timeout": 0,
          "opendaylight-flow-statistics:instructions": {
            "opendaylight-flow-statistics:instruction": [
              {
                "opendaylight-flow-statistics:order": 0,
                "opendaylight-flow-statistics:apply-actions": {
                  "opendaylight-flow-statistics:action": [
                    {
                      "opendaylight-flow-statistics:order": 0,
                      "opendaylight-flow-statistics:output-action": {
                        "opendaylight-flow-statistics:output-node-connector": "2",
                        "opendaylight-flow-statistics:max-length": 0
                      }
                    }
                  ]
                }
              }
            ]
          }
```


(continues on next page)

(continued from previous page)

```
            }
          }
        ],
        "opendaylight-flow-table-statistics:flow-table-statistics": {
          "opendaylight-flow-table-statistics:active-flows": 8,
          "opendaylight-flow-table-statistics:packets-matched": 97683,
          "opendaylight-flow-table-statistics:packets-looked-up": 101772
        }
      }
    ]
}
```

Discovering and testing new Flow Types

Currently, the openflowplugin has a test-provider that allows you to push various flows through the system from the OSGI command line. Once those flows have been pushed through, you can see them as examples and then use them to see in the config what a particular flow example looks like.

Using addMDFlow

From the

```
cd openflowplugin/distribution/base/target/distributions-openflowplugin-base-0.0.1-
↪SNAPSHOT-osgipackage/opendaylight
./run.sh
```

Point your mininet at the controller as described above.

once you can see your node (probably openflow:1 if you've been following along) in the inventory, at the OSGI command line try running:

```
addMDFlow openflow:1 f#
```

Where # is a number between 1 and 80. This will create one of 80 possible flows. You can go confirm they were created on the switch.

Once you've done that, use

- GET
- Accept: application/xml
- URL: "http://192.168.195.157:8181/restconf/config/opendaylight-inventory:nodes/node/openflow:1/table/2/"

To see a full listing of the flows in table 2 (where they will be put). If you want to see a particular flow, look at

- URL: "http://192.168.195.157:8181/restconf/config/opendaylight-inventory:nodes/node/openflow:1/table/2/flow/#"

Where # is 123 + the f# you used. So for example, for f22, your url would be

- URL: "http://192.168.195.157:8181/restconf/config/opendaylight-inventory:nodes/node/openflow:1/table/2/flow/145"

Note: You may have to trim out some of the sections like that contain bitfields and binary types that are not correctly modeled.

Note: Before attempting to PUT a flow you have created via addMDFlow, please change its URL and body to, for example, use table 1 instead of table 2 or another Flow Id, so you don't collide.

Note: There are several test command providers and the one handling flows is **OpenflowpluginTestCommandProvider**. Methods, which can be use as **commands in OSGI-console** have prefix _.

Example Flows

Examples for XML for various flow matches, instructions & actions can be found in following section *here*.

End to End Topology

Introduction

The purpose of this page is to walk you through how to see the Topology Manager working end to end with the openflowplugin using OpenFlow 1.3.

Basically, you will learn how to:

1. Run the Base/Virtualization/Service provider Edition with the new openflowplugin: *Running the controller with the new OpenFlow Plugin*
2. Start mininet to use OF 1.3: OpenFlow 1.3 Enabled Software Switches / Environment
3. Use RESTCONF to see the topology information.

Restconf for Topology

The REST url for listing all the nodes is:

```
http://localhost:8080/restconf/operational/network-topology:network-topology/
```

You will need to set the Accept header:

```
Accept: application/xml
```

You will also need to use HTTP Basic Auth with username: admin password: admin.

Alternately, if you have a node's id you can address it as

```
http://localhost:8080/restconf/operational/network-topology:network-topology/topology/
↪<id>
```

for example

```
http://localhost:8080/restconf/operational/network-topology:network-topology/topology/
↪flow:1/
```

How to hit RestConf with Postman

Install postman for Chrome

In the chrome browser bar enter

```
chrome://apps/
```

And click on Postman.

Enter the URL. Click on the Headers button on the far right. Enter the Accept: header. Click on the Basic Auth Tab at the top and setup the username and password. Send.

End to End Groups

NOTE

Groups are NOT SUPPORTED in current (2.0.0) version of openvswitch. See

- http://openvswitch.org/releases/NEWS-2.0.0
- http://comments.gmane.org/gmane.linux.network.openvswitch.general/3251

For testing group feature please use for example CPQD virtual switch in the *End to End Inventory* section.

Instructions

Learn End to End for Inventory

End to End Inventory

Check inventory

Run CPqD with support for OF 1.3 as described in *End to End Inventory*

Make sure you see the openflow:1 node come up as described in *End to End Inventory*

Group Strategy

Current way to flush a group to switch looks like this:

1. create MD-SAL modeled group and commit it into dataStore using two phase commit
2. FRM gets notified and invokes corresponding rpc (addGroup) on particular service provider (if suitable provider for given node registered)
3. the provider (plugin in this case) transforms MD-SAL modeled group into OF-API modeled group
4. OF-API modeled group is then flushed into OFLibrary
5. OFLibrary encodes group into particular version of wire protocol and sends it to particular switch
6. check on CPqD if group is installed

Push your Group

- With PostMan:
  - Set

        * Content-Type: application/xml
        * Accept: application/xml
    - Use URL: "http://<ip-address>:8080/restconf/config/opendaylight-inventory:nodes/node/openflow:1/group/1"
    - PUT
    - Use Body:

```
<?xml version="1.0" encoding="UTF-8" standalone="no"?>
<group xmlns="urn:opendaylight:flow:inventory">
    <group-type>group-all</group-type>
    <buckets>
        <bucket>
            <action>
                <pop-vlan-action/>
                <order>0</order>
            </action>
            <bucket-id>12</bucket-id>
            <watch_group>14</watch_group>
            <watch_port>1234</watch_port>
        </bucket>
        <bucket>
            <action>
                <set-field>
                    <ipv4-source>100.1.1.1</ipv4-source>
                </set-field>
                <order>0</order>
            </action>
            <action>
                <set-field>
                    <ipv4-destination>200.71.9.5210</ipv4-destination>
                </set-field>
                <order>1</order>
            </action>
            <bucket-id>13</bucket-id>
            <watch_group>14</watch_group>
            <watch_port>1234</watch_port>
        </bucket>
    </buckets>
    <barrier>false</barrier>
    <group-name>Foo</group-name>
    <group-id>1</group-id>
</group>
```

Note: If you want to try a different group id, make sure the URL and the body stay in sync. For example, if you wanted to try: group-id 20 you'd change the URL to "http://<ip-address>:8080/restconf/config/opendaylight-inventory:nodes/node/openflow:1/group/20" but you would also need to update the <group-id>20</group-id> in the body to match.

Note: <ip-address> :Provide the IP Address of the machine on which the controller is running.

Check for your group on the switch

- See your group on your cpqd switch:

```
COMMAND: sudo dpctl tcp:127.0.0.1:6000 stats-group

SENDING:
stat_req{type="grp", flags="0x0", group="all"}


RECEIVED:
stat_repl{type="grp", flags="0x0", stats=[
{group="1", ref_cnt="0", pkt_cnt="0", byte_cnt="0", cntrs=[{pkt_cnt="0", byte_cnt="0"}
↪, {pkt_cnt="0", byte_cnt="0"}]}]}
```

Check for your group in the controller config via RESTCONF

- See your configured group in POSTMAN with
  - URL http://<ip-address>:8080/restconf/config/opendaylight-inventory:nodes/node/openflow:1/group/1
  - GET
  - You should no longer need to set Accept
  - Note: <ip-address> :Provide the IP Address of the machine on which the controller is running.

Look for your group stats in the controller operational data via RESTCONF

- See your operational group stats in POSTMAN with
  - URL http://<ip-address>:8080/restconf/operational/opendaylight-inventory:nodes/node/openflow:1/group/1
  - GET
  - Note: <ip-address> :Provide the IP Address of the machine on which the controller is running.

Discovering and testing Group Types

Currently, the openflowplugin has a test-provider that allows you to push various groups through the system from the OSGI command line. Once those groups have been pushed through, you can see them as examples and then use them to see in the config what a particular group example looks like.

Using addGroup

From the

```
cd openflowplugin/distribution/base/target/distributions-openflowplugin-base-0.0.1-
↪SNAPSHOT-osgipackage/opendaylight
./run.sh
```

Point your CPqD at the controller as described above.

once you can see your node (probably openflow:1 if you've been following along) in the inventory, at the OSGI command line try running:

```
addGroup openflow:1
```

This will install a group in the switch. You can check whether the group is installed or not.

Once you've done that, use

- GET
- Accept: application/xml
- URL: "http://<ip-address>:8080/restconf/config/opendaylight-inventory:nodes/node/openflow:1/group/1"
  - Note: <ip-address> :Provide the IP Address of the machine on which the controller is running.

Note: Before attempting to PUT a group you have created via addGroup, please change its URL and body to, for example, use group 1 instead of group 2 or another Group Id, so that they don't collide.

Note: There are several test command providers and the one handling groups is OpenflowpluginGroupTestCommandProvider. Methods, which can be use as commands in OSGI-console have prefix _.

Example Group

Examples for XML for various Group Types can be found in the test-scripts bundle of the plugin code with names g1.xml, g2.xml and g3.xml.

End to End Meters

Instructions

Learn End to End for Inventory

- *End to End Inventory*

Check inventory

- Run mininet with support for OF 1.3 as described in *End to End Inventory*
- Make sure you see the openflow:1 node come up as described in *End to End Inventory*

Meter Strategy

Current way to flush a meter to switch looks like this:

1. create MD-SAL modeled flow and commit it into dataStore using two phase commit
2. FRM gets notified and invokes corresponding rpc (addMeter) on particular service provider (if suitable provider for given node registered)
3. the provider (plugin in this case) transforms MD-SAL modeled meter into OF-API modeled meter
4. OF-API modeled meter is then flushed into OFLibrary

5. OFLibrary encodes meter into particular version of wire protocol and sends it to particular switch
6. check on mininet side if meter is installed

Push your Meter

- Using PostMan:
  - Set Request Headers
    - Content-Type: application/xml
    - Accept: application/xml
  - Use URL: "http://:8080/restconf/config/opendaylight-inventory:nodes/node/openflow:1/meter/1"
  - Method:PUT
  - Request Body:

```
<?xml version="1.0" encoding="UTF-8" standalone="no"?>
<meter xmlns="urn:opendaylight:flow:inventory">
    <container-name>abcd</container-name>
    <flags>meter-burst</flags>
    <meter-band-headers>
        <meter-band-header>
            <band-burst-size>444</band-burst-size>
            <band-id>0</band-id>
            <band-rate>234</band-rate>
            <dscp-remark-burst-size>5</dscp-remark-burst-size>
            <dscp-remark-rate>12</dscp-remark-rate>
            <prec_level>1</prec_level>
            <meter-band-types>
                <flags>ofpmbt-dscp-remark</flags>
            </meter-band-types>
        </meter-band-header>
    </meter-band-headers>
    <meter-id>1</meter-id>
    <meter-name>Foo</meter-name>
</meter>
```

Note: If you want to try a different meter id, make sure the URL and the body stay in sync. For example, if you wanted to try: meter-id 20 you'd change the URL to "http://:8080/restconf/config/opendaylight-inventory:nodes/node/openflow:1/meter/20" but you would also need to update the 20 in the body to match.

Note: :Provide the IP Address of the machine on which the controller is running.

Check for your meter on the switch

- See your meter on your CPqD switch:

```
COMMAND: $ sudo dpctl tcp:127.0.0.1:6000 meter-config
```

(continues on next page)

(continued from previous page)

```
SENDING:
stat_req{type="mconf", flags="0x0"{meter_id= ffffffff"}


RECEIVED:
stat_repl{type="mconf", flags="0x0", stats=[{meter= c"", flags="4", bands=[{type =␣
↪dscp_remark, rate="12", burst_size="5", prec_level="1"}]}]}
```

Check for your meter in the controller config via RESTCONF

- See your configured flow in POSTMAN with
  - URL "http://:8080/restconf/config/opendaylight-inventory:nodes/node/openflow:1/meter/1"
  - Method: GET
  - You should no longer need to set Request Headers for Accept
  - Note: :Provide the IP Address of the machine on which the controller is running.

Look for your meter stats in the controller operational data via RESTCONF

- See your operational meter stats in POSTMAN with
  - URL "http://:8080/restconfig/operational/opendaylight-inventory:nodes/node/openflow:1/meter/1"
  - Method: GET
  - Note: :Provide the IP Address of the machine on which the controller is running.

Discovering and testing Meter Types

Currently, the openflowplugin has a test-provider that allows you to push various meters through the system from the OSGI command line. Once those meters have been pushed through, you can see them as examples and then use them to see in the config what a particular meter example looks like.

Using addMeter

From the

```
cd openflowplugin/distribution/base/target/distributions-openflowplugin-base-0.0.1-
↪SNAPSHOT-osgipackage/opendaylight
./run.sh
```

Point your CPqD at the controller as described above.

Once you can see your CPqD connected to the controller, at the OSGI command line try running:

```
addMeter openflow:1
```

Once you've done that, use

- GET
- Accept: application/xml

- URL: "http://:8080/restconf/config/opendaylight-inventory:nodes/node/openflow:1/meter/12"
  - Note: :Provide the IP Address of the machine on which the controller is running.

Note: Before attempting to PUT a meter you have created via addMeter, please change its URL and body to, for example, use meter 1 instead of meter 2 or another Meter Id, so you don't collide.

Note: There are several test command providers and the one handling Meter is **OpenflowpluginMeterTestCommandProvider**. Methods, which can be used as **commands in OSGI-console** have prefix _. Examples: addMeter, modifyMeter and removeMeter.

Example Meter

Examples for XML for various Meter Types can be found in the test-scripts bundle of the plugin code with names m1.xml, m2.xml and m3.xml.

Statistics

Overview

This page contains high level detail about the statistics collection mechanism in new OpenFlow plugin.

Statistics collection in new OpenFlow plugin

New OpenFlow plugin collects following statistics from OpenFlow enabled node(switch):

1. Individual Flow Statistics
2. Aggregate Flow Statistics
3. Flow Table Statistics
4. Port Statistics
5. Group Description
6. Group Statistics
7. Meter Configuration
8. Meter Statistics
9. Queue Statistics
10. Node Description
11. Flow Table Features
12. Port Description
13. Group Features
14. Meter Features

At a high level statistics collection mechanism is divided into following three parts

1. Statistics related YANG models, service APIs and notification interfaces are defined in the MD-SAL.
2. Service APIs (RPCs) defined in yang models are implemented by OpenFlow plugin. Notification interfaces are wired up by OpenFlow plugin to MD-SAL.
3. Statistics Manager Module: This module use service APIs implemented by OpenFlow plugin to send statistics requests to all the connected OpenFlow enabled nodes. Module also implements notification interfaces to receive statistics response from nodes. Once it receives statistics response, it augment all the statistics data to the relevant element of the node (like node-connector, flow, table,group, meter) and store it in MD-SAL operational data store.

Details of statistics collection

- Current implementation collects above mentioned statistics (except 10-14) at a periodic interval of 15 seconds.
- Statistics mentioned in 10 to 14 are only fetched when any node connects to the controller because these statistics are just static details about the respective elements.
- Whenever any new element is added to node (like flow, group, meter, queue) it sends statistics request immediately to fetch the latest statistics and store it in the operational data store.
- Whenever any element is deleted from the node, it immediately remove the relevant statistics from operational data store.
- Statistics data are augmented to their respective element stored in the configuration data store. E.g Controller installed flows are stored in configuration data store. Whenever Statistics Manager receive statistics data related to these flow, it search the corresponding flow in the configuration data store and augment statistics in the corresponding location in operational data store. Similar approach is used for other elements of the node.
- Statistics Manager stores flow statistics as an unaccounted flow statistics in operational data store if there is no corresponding flow exist in configuration data store. ID format of unaccounted flow statistics is as follows - [#UF$TABLE**Unaccounted-flow-count - e.g #UF$TABLE*2*1].
- All the unaccounted flows will be cleaned up periodically after every two cycle of flow statistics collection, given that there is no update for these flows in the last two cycles.
- Statistics Manager only entertains statistics response for the request sent by itself. User can write its own statistics collector using the statistics service APIs and notification defined in yang models, it won't effect the functioning of Statistics Manager.
- OpenFlow 1.0 don't have concept of Meter and Group, so Statistics Manager don't send any group & meter related statistics request to OpenFlow 1.0 enabled switch.

RESTCONF Uris to access statistics of various node elements

- Aggregate Flow Statistics & Flow Table Statistics

```
GET  http://<controller-ip>:8080/restconf/operational/opendaylight-inventory:nodes/
↪node/{node-id}/table/{table-id}
```

- Individual Flow Statistics from specific table

```
GET  http://<controller-ip>:8080/restconf/operational/opendaylight-inventory:nodes/
↪node/{node-id}/table/{table-id}/flow/{flow-id}
```

- Group Features & Meter Features Statistics

```
GET  http://<controller-ip>:8080/restconf/operational/opendaylight-inventory:nodes/
↪node/{node-id}
```

- Group Description & Group Statistics

```
GET  http://<controller-ip>:8080/restconf/operational/opendaylight-inventory:nodes/
↪node/{node-id}/group/{group-id}
```

- Meter Configuration & Meter Statistics

```
GET  http://<controller-ip>:8080/restconf/operational/opendaylight-inventory:nodes/
↪node/{node-id}/meter/{meter-id}
```

- Node Connector Statistics

```
GET  http://<controller-ip>:8080/restconf/operational/opendaylight-inventory:nodes/
↪node/{node-id}/node-connector/{node-connector-id}
```

- Queue Statistics

```
GET  http://<controller-ip>:8080/restconf/operational/opendaylight-inventory:nodes/
↪node/{node-id}/node-connector/{node-connector-id}/queue/{queue-id}
```

Bugs

For more details and queuries, please send mail to openflowplugin-dev@lists.opendaylight.org or avishnoi@in.ibm.com If you want to report any bug in statistics collection, please use bugzilla.

Web / Graphical Interface

In the Hydrogen & Helium release, the current Web UI does not support the new OpenFlow 1.3 constructs such as groups, meters, new fields in the flows, multiple flow tables, etc.

Command Line Interface

The following is not exactly CLI - just a set of test commands which can be executed on the OSGI console testing various features in OpenFlow 1.3 spec.

- *OSGI Console Test Provider Commands: Flows*
- *OSGI Console Test Provider Commands: Groups*
- *OSGI Console Test Provider Commands: Meters*
- *OSGI Console Test Provider Commands: Topology Events*

Flows : Test Provider

Currently, the openflowplugin has a test-provider that allows you to push various flows through the system from the OSGI command line. Once those flows have been pushed through, you can see them as examples and then use them to see in the config what a particular flow example looks like.

AddFlow : addMDFlow

Run the controller by executing:

```
cd openflowplugin/distribution/base/target/distributions-openflowplugin-base-0.0.1-
↪SNAPSHOT-osgipackage/opendaylight
./run.sh
```

Point your mininet to the controller by giving the parameters –controller=remote,ip=.

Once you see your node (probably openflow:1 if you've been following along) in the inventory, at the OSGI command line try running:

```
addMDFlow openflow:1 f#
```

Where # is a number between 1 and 80 and openflow:1 is the of the switch. This will create one of 80 possible flows. You can confirm that they were created on the switch.

RemoveFlow : removeMDFlow

Similar to addMDFlow, from the controller OSGi prompt, while your switch is connected to the controller, try running:

```
removeMDFlow openflow:1 f#
```

where # is a number between 1 and 80 and openflow:1 is the of the switch. The flow to be deleted should have same flowid and Nodeid as used for flow add.

ModifyFlow : modifyMDFlow

Similar to addMDFlow, from the controller OSGi prompt, while your switch is connected to the controller, try running:

```
modifyMDFlow openflow:1 f#
```

where # is a number between 1 and 80 and openflow:1 is the of the switch. The flow to be deleted should have same flowid and Nodeid as used for flow add.

Group : Test Provider

Currently, the openflowplugin has a test-provider that allows you to push various flows through the system from the OSGI command line. Once those flows have been pushed through, you can see them as examples and then use them to see in the config what a particular flow example looks like.

AddGroup : addGroup

Run the controller by executing:

```
cd openflowplugin/distribution/base/target/distributions-openflowplugin-base-0.0.1-
↪SNAPSHOT-osgipackage/opendaylight
./run.sh
```

Point your mininet to the controller by giving the parameters –controller=remote,ip=.

Once you see your node (probably openflow:1 if you've been following along) in the inventory, at the OSGI command line try running:

```
addGroup openflow:1 a# g#
```

Where # is a number between 1 and 4 for grouptype(g#) and 1 and 28 for actiontype(a#). You can confirm that they were created on the switch.

RemoveGroup : removeGroup

Run the controller by executing:

```
cd openflowplugin/distribution/base/target/distributions-openflowplugin-base-0.0.1-
↪SNAPSHOT-osgipackage/opendaylight
./run.sh
```

Point your mininet at the controller as described above.

Once you see your node (probably openflow:1 if you've been following along) in the inventory, at the OSGI command line try running:

```
removeGroup openflow:1 a# g#
```

Where # is a number between 1 and 4 for grouptype(g#) and 1 and 28 for actiontype(a#). GroupId should be same as that used for adding the flow. You can confirm that it was removed from the switch.

ModifyGroup : modifyGroup

Run the controller by executing:

```
cd openflowplugin/distribution/base/target/distributions-openflowplugin-base-0.0.1-
↪SNAPSHOT-osgipackage/opendaylight
./run.sh
```

Point your mininet at the controller as described above.

Once you see your node (probably openflow:1 if you've been following along) in the inventory, at the OSGI command line try running:

```
modifyGroup openflow:1 a# g#
```

Where # is a number between 1 and 4 for grouptype(g#) and 1 and 28 for actiontype(a#). GroupId should be same as that used for adding the flow. You can confirm that it was modified on the switch.

Meters : Test Provider

Currently, the openflowplugin has a test-provider that allows you to push various flows through the system from the OSGI command line. Once those flows have been pushed through, you can see them as examples and then use them to see in the config what a particular flow example looks like.

AddMeter : addMeter

Run the controller by executing:

```
cd openflowplugin/distribution/base/target/distributions-openflowplugin-base-0.0.1-
↪SNAPSHOT-osgipackage/opendaylight
./run.sh
```

Point your mininet to the controller by giving the parameters –controller=remote,ip=.

Once you see your node (probably openflow:1 if you've been following along) in the inventory, at the OSGI command line try running:

```
addMeter openflow:1
```

You can now confirm that meter has been created on the switch.

RemoveMeter : removeMeter

Run the controller by executing:

```
cd openflowplugin/distribution/base/target/distributions-openflowplugin-base-0.0.1-
↪SNAPSHOT-osgipackage/opendaylight
./run.sh
```

Point your mininet to the controller by giving the parameters –controller=remote,ip=.

Once you see your node (probably openflow:1 if you've been following along) in the inventory, at the OSGI command line try running:

```
removeMeter openflow:1
```

The CLI takes care of using the same meterId and nodeId as used for meter add. You can confirm that it was removed from the switch.

ModifyMeter : modifyMeter

Run the controller by executing:

```
cd openflowplugin/distribution/base/target/distributions-openflowplugin-base-0.0.1-
↪SNAPSHOT-osgipackage/opendaylight
./run.sh
```

Point your mininet to the controller by giving the parameters –controller=remote,ip=.

Once you see your node (probably openflow:1 if you've been following along) in the inventory, at the OSGI command line try running:

```
modifyMeter openflow:1
```

The CLI takes care of using the same meterId and nodeId as used for meter add. You can confirm that it was modified on the switch.

Topology : Notification

Currently, the openflowplugin has a test-provider that allows you to get notifications for the topology related events like Link-Discovered , Link-Removed events.

Link Discovered Event : Testing

Run the controller by executing:

```
cd openflowplugin/distribution/base/target/distributions-openflowplugin-base-0.0.1-
↪SNAPSHOT-osgipackage/opendaylight
./run.sh
```

Point your mininet to the controller by giving the parameters –controller=remote,ip=. Once the controller is connected to the switch, Link-Discovered event can be tested by initially configuring the specific flows on the switch. For Link Discovered event either table-miss flow or LLDP ether-type flow can be configured.

Configuring Table-Miss flow using OpenflowpluginTestCommandProvider

```
addMDFlow Openflow:1 fTM
```

as per this OpenDaylight_OpenFlow_Plugin:Test_Provider#Flows_:_Test_Provider[link]. *fTM* is the table-miss scenario here.

Once the table-miss flow is configured through above command, we can see the Link-Discovered event in the debug logs on the controller console.

Configuring LLDP ether-type flow using OpenflowpluginTestCommandProvider

```
addMDFlow Openflow:1 0(table-id) f81
```

You can confirm that they were created on the switch.

Once the LLDP ether-type flow is configured through above command, we can see the Link-Discovered event in the debug logs on the controller console.

Link Removed Event : Testing

Having configured either table-miss or lldp ether-type flow on switch, once the switch is disconnected we see the Link-Removed event

Programmatic Interface

The API is documented in the model documentation under the section OpenFlow Services at:

- Models Documentation (OpenFlow Services Section)

Example flows

Overview

The flow examples on this page are tested to work with OVS.

Use, for example, POSTMAN with the following parameters:

```
PUT http://<ctrl-addr>:8080/restconf/config/opendaylight-inventory:nodes/node/<Node-
↪id>/table/<Table-#>/flow/<Flow-#>

- Accept: application/xml
- Content-Type: application/xml
```

For example:

```
PUT http://localhost:8080/restconf/config/opendaylight-inventory:nodes/node/
↪openflow:1/table/2/flow/127
```

Make sure that the Table-# and Flow-# in the URL and in the XML match.

The format of the flow-programming XML is determined by by the grouping *flow* in the opendaylight-flow-types yang model: MISSING LINK.

Match Examples

The format of the XML that describes OpenFlow matches is determined by the opendaylight-match-types yang model: .

IPv4 Dest Address

- Flow=124, Table=2, Priority=2, Instructions=\{Apply_Actions={dec_nw_ttl}}, match=\{ipv4_destination_address=10.0.1.1/24}
- Note that ethernet-type MUST be 2048 (0x800)

```
<?xml version="1.0" encoding="UTF-8" standalone="no"?>
<flow xmlns="urn:opendaylight:flow:inventory">
    <strict>false</strict>
    <instructions>
        <instruction>
            <order>0</order>
            <apply-actions>
                <action>
                    <order>0</order>
                    <dec-nw-ttl/>
                </action>
            </apply-actions>
        </instruction>
    </instructions>
    <table_id>2</table_id>
    <id>124</id>
    <cookie_mask>255</cookie_mask>
    <installHw>false</installHw>
    <match>
        <ethernet-match>
            <ethernet-type>
                <type>2048</type>
            </ethernet-type>
        </ethernet-match>
        <ipv4-destination>10.0.1.1/24</ipv4-destination>
    </match>
    <hard-timeout>12</hard-timeout>
```

(continues on next page)

(continued from previous page)

```
    <cookie>1</cookie>
    <idle-timeout>34</idle-timeout>
    <flow-name>FooXf1</flow-name>
    <priority>2</priority>
    <barrier>false</barrier>
</flow>
```

Ethernet Src Address

- Flow=126, Table=2, Priority=2, Instructions=\{Apply_Actions={drop}}, match=\{ethernet-source=00:00:00:00:00:01}

```
<?xml version="1.0" encoding="UTF-8" standalone="no"?>
<flow xmlns="urn:opendaylight:flow:inventory">
    <strict>false</strict>
    <instructions>
        <instruction>
            <order>0</order>
            <apply-actions>
                <action>
                    <order>0</order>
                    <drop-action/>
                </action>
            </apply-actions>
        </instruction>
    </instructions>
    <table_id>2</table_id>
    <id>126</id>
    <cookie_mask>255</cookie_mask>
    <installHw>false</installHw>
    <match>
        <ethernet-match>
            <ethernet-source>
                <address>00:00:00:00:00:01</address>
            </ethernet-source>
        </ethernet-match>
    </match>
    <hard-timeout>12</hard-timeout>
    <cookie>3</cookie>
    <idle-timeout>34</idle-timeout>
    <flow-name>FooXf3</flow-name>
    <priority>2</priority>
    <barrier>false</barrier>
</flow>
```

Ethernet Src & Dest Addresses, Ethernet Type

- Flow=127, Table=2, Priority=2, Instructions=\{Apply_Actions={drop}}, match=\{ethernet-source=00:00:00:00:23:ae, ethernet-destination=ff:ff:ff:ff:ff:ff, ethernet-type=45}

```
<?xml version="1.0" encoding="UTF-8" standalone="no"?>
<flow xmlns="urn:opendaylight:flow:inventory">
    <strict>false</strict>
```

(continues on next page)

(continued from previous page)

```
    <instructions>
        <instruction>
            <order>0</order>
            <apply-actions>
                <action>
                    <order>0</order>
                    <dec-mpls-ttl/>
                </action>
            </apply-actions>
        </instruction>
    </instructions>
    <table_id>2</table_id>
    <id>127</id>
    <cookie_mask>255</cookie_mask>
    <installHw>false</installHw>
    <match>
        <ethernet-match>
            <ethernet-type>
                <type>45</type>
            </ethernet-type>
            <ethernet-destination>
                <address>ff:ff:ff:ff:ff:ff</address>
            </ethernet-destination>
            <ethernet-source>
                <address>00:00:00:00:23:ae</address>
            </ethernet-source>
        </ethernet-match>
    </match>
    <hard-timeout>12</hard-timeout>
    <cookie>4</cookie>
    <idle-timeout>34</idle-timeout>
    <flow-name>FooXf4</flow-name>
    <priority>2</priority>
    <barrier>false</barrier>
</flow>
```

Ethernet Src & Dest Addresses, IPv4 Src & Dest Addresses, Input Port

- Note that ethernet-type MUST be 34887 (0x8847)

```
<?xml version="1.0" encoding="UTF-8" standalone="no"?>
<flow xmlns="urn:opendaylight:flow:inventory">
    <strict>false</strict>
    <instructions>
        <instruction>
            <order>0</order>
            <apply-actions>
                <action>
                    <order>0</order>
                    <dec-mpls-ttl/>
                </action>
            </apply-actions>
        </instruction>
    </instructions>
    <table_id>2</table_id>
```

(continues on next page)

(continued from previous page)

```
    <id>128</id>
    <cookie_mask>255</cookie_mask>
    <match>
        <ethernet-match>
            <ethernet-type>
                <type>34887</type>
            </ethernet-type>
            <ethernet-destination>
                <address>ff:ff:ff:ff:ff:ff</address>
            </ethernet-destination>
            <ethernet-source>
                <address>00:00:00:00:23:ae</address>
            </ethernet-source>
        </ethernet-match>
        <ipv4-source>10.1.2.3/24</ipv4-source>
        <ipv4-destination>20.4.5.6/16</ipv4-destination>
        <in-port>0</in-port>
    </match>
    <hard-timeout>12</hard-timeout>
    <cookie>5</cookie>
    <idle-timeout>34</idle-timeout>
    <flow-name>FooXf5</flow-name>
    <priority>2</priority>
    <barrier>false</barrier>
</flow>
```

Ethernet Src & Dest Addresses, IPv4 Src & Dest Addresses, IP

Protocol #, IP DSCP, IP ECN, Input Port

- Note that ethernet-type MUST be 2048 (0x800)

```
<?xml version="1.0" encoding="UTF-8" standalone="no"?>
<flow xmlns="urn:opendaylight:flow:inventory">
    <strict>false</strict>
    <instructions>
        <instruction>
            <order>0</order>
            <apply-actions>
                <action>
                    <order>0</order>
                    <dec-nw-ttl/>
                </action>
            </apply-actions>
        </instruction>
    </instructions>
    <table_id>2</table_id>
    <id>130</id>
    <cookie_mask>255</cookie_mask>
    <match>
        <ethernet-match>
            <ethernet-type>
                <type>2048</type>
            </ethernet-type>
            <ethernet-destination>
```

(continues on next page)

(continued from previous page)

```
                <address>ff:ff:ff:ff:ff:aa</address>
            </ethernet-destination>
            <ethernet-source>
                <address>00:00:00:11:23:ae</address>
            </ethernet-source>
        </ethernet-match>
        <ipv4-source>10.1.2.3/24</ipv4-source>
        <ipv4-destination>20.4.5.6/16</ipv4-destination>
        <ip-match>
            <ip-protocol>56</ip-protocol>
            <ip-dscp>15</ip-dscp>
            <ip-ecn>1</ip-ecn>
        </ip-match>
        <in-port>0</in-port>
    </match>
    <hard-timeout>12000</hard-timeout>
    <cookie>7</cookie>
    <idle-timeout>12000</idle-timeout>
    <flow-name>FooXf7</flow-name>
    <priority>2</priority>
    <barrier>false</barrier>
</flow>
```

Ethernet Src & Dest Addresses, IPv4 Src & Dest Addresses, TCP Src &

Dest Ports, IP DSCP, IP ECN, Input Port

- Note that ethernet-type MUST be 2048 (0x800)
- Note that IP Protocol Type MUST be 6

```
<?xml version="1.0" encoding="UTF-8" standalone="no"?>
<flow xmlns="urn:opendaylight:flow:inventory">
    <strict>false</strict>
    <instructions>
        <instruction>
            <order>0</order>
            <apply-actions>
                <action>
                    <order>0</order>
                    <dec-nw-ttl/>
                </action>
            </apply-actions>
        </instruction>
    </instructions>
    <table_id>2</table_id>
    <id>131</id>
    <cookie_mask>255</cookie_mask>
    <match>
        <ethernet-match>
            <ethernet-type>
                <type>2048</type>
            </ethernet-type>
            <ethernet-destination>
                <address>ff:ff:29:01:19:61</address>
```

(continues on next page)

(continued from previous page)

```
            </ethernet-destination>
            <ethernet-source>
                <address>00:00:00:11:23:ae</address>
            </ethernet-source>
        </ethernet-match>
        <ipv4-source>17.1.2.3/8</ipv4-source>
        <ipv4-destination>172.168.5.6/16</ipv4-destination>
        <ip-match>
            <ip-protocol>6</ip-protocol>
            <ip-dscp>2</ip-dscp>
            <ip-ecn>2</ip-ecn>
        </ip-match>
        <tcp-source-port>25364</tcp-source-port>
        <tcp-destination-port>8080</tcp-destination-port>
        <in-port>0</in-port>
    </match>
    <hard-timeout>1200</hard-timeout>
    <cookie>8</cookie>
    <idle-timeout>3400</idle-timeout>
    <flow-name>FooXf8</flow-name>
    <priority>2</priority>
    <barrier>false</barrier>
</flow>
```

Ethernet Src & Dest Addresses, IPv4 Src & Dest Addresses, UDP Src &

Dest Ports, IP DSCP, IP ECN, Input Port

- Note that ethernet-type MUST be 2048 (0x800)
- Note that IP Protocol Type MUST be 17

```
<?xml version="1.0" encoding="UTF-8" standalone="no"?>
<flow xmlns="urn:opendaylight:flow:inventory">
    <strict>false</strict>
    <instructions>
        <instruction>
            <order>0</order>
            <apply-actions>
                <action>
                    <order>0</order>
                    <dec-nw-ttl/>
                </action>
            </apply-actions>
        </instruction>
    </instructions>
    <table_id>2</table_id>
    <id>132</id>
    <cookie_mask>255</cookie_mask>
    <match>
        <ethernet-match>
            <ethernet-type>
                <type>2048</type>
            </ethernet-type>
            <ethernet-destination>
```

(continues on next page)

(continued from previous page)

```
                <address>20:14:29:01:19:61</address>
            </ethernet-destination>
            <ethernet-source>
                <address>00:00:00:11:23:ae</address>
            </ethernet-source>
        </ethernet-match>
        <ipv4-source>19.1.2.3/10</ipv4-source>
        <ipv4-destination>172.168.5.6/18</ipv4-destination>
        <ip-match>
            <ip-protocol>17</ip-protocol>
            <ip-dscp>8</ip-dscp>
            <ip-ecn>3</ip-ecn>
        </ip-match>
        <udp-source-port>25364</udp-source-port>
        <udp-destination-port>8080</udp-destination-port>
        <in-port>0</in-port>
    </match>
    <hard-timeout>1200</hard-timeout>
    <cookie>9</cookie>
    <idle-timeout>3400</idle-timeout>
    <flow-name>FooXf9</flow-name>
    <priority>2</priority>
    <barrier>false</barrier>
```

Ethernet Src & Dest Addresses, IPv4 Src & Dest Addresses, ICMPv4

Type & Code, IP DSCP, IP ECN, Input Port

- Note that ethernet-type MUST be 2048 (0x800)
- Note that IP Protocol Type MUST be 1

```
<?xml version="1.0" encoding="UTF-8" standalone="no"?>
<flow xmlns="urn:opendaylight:flow:inventory">
    <strict>false</strict>
    <instructions>
        <instruction>
            <order>0</order>
            <apply-actions>
                <action>
                    <order>0</order>
                    <dec-nw-ttl/>
                </action>
            </apply-actions>
        </instruction>
    </instructions>
    <table_id>2</table_id>
    <id>134</id>
    <cookie_mask>255</cookie_mask>
    <match>
        <ethernet-match>
            <ethernet-type>
                <type>2048</type>
            </ethernet-type>
            <ethernet-destination>
```

(continues on next page)

(continued from previous page)

```
                <address>ff:ff:29:01:19:61</address>
            </ethernet-destination>
            <ethernet-source>
                <address>00:00:00:11:23:ae</address>
            </ethernet-source>
        </ethernet-match>
        <ipv4-source>17.1.2.3/8</ipv4-source>
        <ipv4-destination>172.168.5.6/16</ipv4-destination>
        <ip-match>
            <ip-protocol>1</ip-protocol>
            <ip-dscp>27</ip-dscp>
            <ip-ecn>3</ip-ecn>
        </ip-match>
        <icmpv4-match>
            <icmpv4-type>6</icmpv4-type>
            <icmpv4-code>3</icmpv4-code>
        </icmpv4-match>
        <in-port>0</in-port>
    </match>
    <hard-timeout>1200</hard-timeout>
    <cookie>11</cookie>
    <idle-timeout>3400</idle-timeout>
    <flow-name>FooXf11</flow-name>
    <priority>2</priority>
</flow>
```

Ethernet Src & Dest Addresses, ARP Operation, ARP Src & Target

Transport Addresses, ARP Src & Target Hw Addresses

- Note that ethernet-type MUST be 2054 (0x806)

```
<?xml version="1.0" encoding="UTF-8" standalone="no"?>
<flow xmlns="urn:opendaylight:flow:inventory">
    <strict>false</strict>
    <instructions>
        <instruction>
            <order>0</order>
            <apply-actions>
                <action>
                    <order>0</order>
                    <dec-nw-ttl/>
                </action>
                <action>
                    <order>1</order>
                    <dec-mpls-ttl/>
                </action>
            </apply-actions>
        </instruction>
    </instructions>
    <table_id>2</table_id>
    <id>137</id>
    <cookie_mask>255</cookie_mask>
    <match>
        <ethernet-match>
```

(continues on next page)

(continued from previous page)

```
            <ethernet-type>
                <type>2054</type>
            </ethernet-type>
            <ethernet-destination>
                <address>ff:ff:ff:ff:FF:ff</address>
            </ethernet-destination>
            <ethernet-source>
                <address>00:00:FC:01:23:ae</address>
            </ethernet-source>
        </ethernet-match>
        <arp-op>1</arp-op>
        <arp-source-transport-address>192.168.4.1</arp-source-transport-address>
        <arp-target-transport-address>10.21.22.23</arp-target-transport-address>
        <arp-source-hardware-address>
            <address>12:34:56:78:98:AB</address>
        </arp-source-hardware-address>
        <arp-target-hardware-address>
            <address>FE:DC:BA:98:76:54</address>
        </arp-target-hardware-address>
    </match>
    <hard-timeout>12</hard-timeout>
    <cookie>14</cookie>
    <idle-timeout>34</idle-timeout>
    <flow-name>FooXf14</flow-name>
    <priority>2</priority>
    <barrier>false</barrier>
```

Ethernet Src & Dest Addresses, Ethernet Type, VLAN ID, VLAN PCP

```
<?xml version="1.0" encoding="UTF-8" standalone="no"?>
<flow xmlns="urn:opendaylight:flow:inventory">
    <strict>false</strict>
    <instructions>
        <instruction>
            <order>0</order>
            <apply-actions>
                <action>
                    <order>0</order>
                    <dec-nw-ttl/>
                </action>
            </apply-actions>
        </instruction>
    </instructions>
    <table_id>2</table_id>
    <id>138</id>
    <cookie_mask>255</cookie_mask>
    <match>
        <ethernet-match>
            <ethernet-type>
                <type>2048</type>
            </ethernet-type>
            <ethernet-destination>
                <address>ff:ff:29:01:19:61</address>
            </ethernet-destination>
```

(continues on next page)

(continued from previous page)

```
            <ethernet-source>
                <address>00:00:00:11:23:ae</address>
            </ethernet-source>
        </ethernet-match>
        <vlan-match>
            <vlan-id>
                <vlan-id>78</vlan-id>
                <vlan-id-present>true</vlan-id-present>
            </vlan-id>
            <vlan-pcp>3</vlan-pcp>
      </vlan-match>
    </match>
    <hard-timeout>1200</hard-timeout>
    <cookie>15</cookie>
    <idle-timeout>3400</idle-timeout>
    <flow-name>FooXf15</flow-name>
    <priority>2</priority>
    <barrier>false</barrier>
</flow>
```

Ethernet Src & Dest Addresses, MPLS Label, MPLS TC, MPLS BoS

```
<?xml version="1.0" encoding="UTF-8" standalone="no"?>
<flow xmlns="urn:opendaylight:flow:inventory">
    <flow-name>FooXf17</flow-name>
    <id>140</id>
    <cookie_mask>255</cookie_mask>
    <cookie>17</cookie>
    <hard-timeout>1200</hard-timeout>
    <idle-timeout>3400</idle-timeout>
    <priority>2</priority>
    <table_id>2</table_id>
    <strict>false</strict>
    <instructions>
        <instruction>
            <order>0</order>
            <apply-actions>
                <action>
                    <order>0</order>
                    <dec-nw-ttl/>
                </action>
            </apply-actions>
        </instruction>
    </instructions>
    <match>
        <ethernet-match>
            <ethernet-type>
                <type>34887</type>
            </ethernet-type>
            <ethernet-destination>
                <address>ff:ff:29:01:19:61</address>
            </ethernet-destination>
            <ethernet-source>
                <address>00:00:00:11:23:ae</address>
```

(continues on next page)

(continued from previous page)

```
            </ethernet-source>
        </ethernet-match>
        <protocol-match-fields>
            <mpls-label>567</mpls-label>
            <mpls-tc>3</mpls-tc>
            <mpls-bos>1</mpls-bos>
        </protocol-match-fields>
    </match>
</flow>
```

IPv6 Src & Dest Addresses

- Note that ethernet-type MUST be 34525

```
<?xml version="1.0" encoding="UTF-8" standalone="no"?>
<flow xmlns="urn:opendaylight:flow:inventory">
    <strict>false</strict>
    <flow-name>FooXf18</flow-name>
    <id>141</id>
    <cookie_mask>255</cookie_mask>
    <cookie>18</cookie>
    <table_id>2</table_id>
    <priority>2</priority>
    <hard-timeout>1200</hard-timeout>
    <idle-timeout>3400</idle-timeout>
    <installHw>false</installHw>
    <instructions>
        <instruction>
            <order>0</order>
            <apply-actions>
                <action>
                    <order>0</order>
                    <dec-nw-ttl/>
                </action>
            </apply-actions>
        </instruction>
    </instructions>
    <match>
        <ethernet-match>
            <ethernet-type>
                <type>34525</type>
            </ethernet-type>
        </ethernet-match>
        <ipv6-source>fe80::2acf:e9ff:fe21:6431/128</ipv6-source>
        <ipv6-destination>aabb:1234:2acf:e9ff::fe21:6431/64</ipv6-destination>
    </match>
</flow>
```

Metadata

```
<?xml version="1.0" encoding="UTF-8" standalone="no"?>
<flow xmlns="urn:opendaylight:flow:inventory">
    <strict>false</strict>
```

(continues on next page)

(continued from previous page)

```
    <flow-name>FooXf19</flow-name>
    <id>142</id>
    <cookie_mask>255</cookie_mask>
    <cookie>19</cookie>
    <table_id>2</table_id>
    <priority>1</priority>
    <hard-timeout>1200</hard-timeout>
    <idle-timeout>3400</idle-timeout>
    <installHw>false</installHw>
    <instructions>
        <instruction>
            <order>0</order>
            <apply-actions>
                <action>
                    <order>0</order>
                    <dec-nw-ttl/>
                </action>
            </apply-actions>
        </instruction>
    </instructions>
    <match>
        <metadata>
            <metadata>12345</metadata>
        </metadata>
    </match>
</flow>
```

Metadata, Metadata Mask

```
<?xml version="1.0" encoding="UTF-8" standalone="no"?>
<flow xmlns="urn:opendaylight:flow:inventory">
    <strict>false</strict>
    <flow-name>FooXf20</flow-name>
    <id>143</id>
    <cookie_mask>255</cookie_mask>
    <cookie>20</cookie>
    <table_id>2</table_id>
    <priority>2</priority>
    <hard-timeout>1200</hard-timeout>
    <idle-timeout>3400</idle-timeout>
    <installHw>false</installHw>
    <instructions>
        <instruction>
            <order>0</order>
            <apply-actions>
                <action>
                    <order>0</order>
                    <dec-nw-ttl/>
                </action>
            </apply-actions>
        </instruction>
    </instructions>
    <match>
        <metadata>
```

(continues on next page)

(continued from previous page)

```
            <metadata>12345</metadata>
            <metadata-mask>//FF</metadata-mask>
        </metadata>
    </match>
</flow>
```

IPv6 Src & Dest Addresses, Metadata, IP DSCP, IP ECN, UDP Src & Dest Ports

- Note that ethernet-type MUST be 34525

```
<?xml version="1.0" encoding="UTF-8" standalone="no"?>
<flow xmlns="urn:opendaylight:flow:inventory">
    <strict>false</strict>
    <flow-name>FooXf21</flow-name>
    <id>144</id>
    <cookie_mask>255</cookie_mask>
    <cookie>21</cookie>
    <table_id>2</table_id>
    <priority>2</priority>
    <hard-timeout>1200</hard-timeout>
    <idle-timeout>3400</idle-timeout>
    <installHw>false</installHw>
    <instructions>
        <instruction>
            <order>0</order>
            <apply-actions>
                <action>
                    <order>0</order>
                    <dec-nw-ttl/>
                </action>
            </apply-actions>
        </instruction>
    </instructions>
    <match>
        <ethernet-match>
            <ethernet-type>
                <type>34525</type>
            </ethernet-type>
        </ethernet-match>
        <ipv6-source>1234:5678:9ABC:DEF0:FDCD:A987:6543:210F/76</ipv6-source>
        <ipv6-destination>fe80::2acf:e9ff:fe21:6431/128</ipv6-destination>
        <metadata>
            <metadata>12345</metadata>
        </metadata>
        <ip-match>
            <ip-protocol>17</ip-protocol>
            <ip-dscp>8</ip-dscp>
            <ip-ecn>3</ip-ecn>
        </ip-match>
        <udp-source-port>25364</udp-source-port>
        <udp-destination-port>8080</udp-destination-port>
    </match>
</flow>
```

IPv6 Src & Dest Addresses, Metadata, IP DSCP, IP ECN, TCP Src & Dest Ports

- Note that ethernet-type MUST be 34525
- Note that IP Protocol MUST be 6

```
<?xml version="1.0" encoding="UTF-8" standalone="no"?>
<flow xmlns="urn:opendaylight:flow:inventory">
    <strict>false</strict>
    <flow-name>FooXf22</flow-name>
    <id>145</id>
    <cookie_mask>255</cookie_mask>
    <cookie>22</cookie>
    <table_id>2</table_id>
    <priority>2</priority>
    <hard-timeout>1200</hard-timeout>
    <idle-timeout>3400</idle-timeout>
    <installHw>false</installHw>
    <instructions>
        <instruction>
            <order>0</order>
            <apply-actions>
                <action>
                    <order>0</order>
                    <dec-nw-ttl/>
                </action>
            </apply-actions>
        </instruction>
    </instructions>
    <match>
        <ethernet-match>
            <ethernet-type>
                <type>34525</type>
            </ethernet-type>
        </ethernet-match>
        <ipv6-source>1234:5678:9ABC:DEF0:FDCD:A987:6543:210F/76</ipv6-source>
        <ipv6-destination>fe80:2acf:e9ff:fe21::6431/94</ipv6-destination>
        <metadata>
            <metadata>12345</metadata>
        </metadata>
        <ip-match>
            <ip-protocol>6</ip-protocol>
            <ip-dscp>60</ip-dscp>
            <ip-ecn>3</ip-ecn>
        </ip-match>
        <tcp-source-port>183</tcp-source-port>
        <tcp-destination-port>8080</tcp-destination-port>
    </match>
</flow>
```

IPv6 Src & Dest Addresses, Metadata, IP DSCP, IP ECN, TCP Src & Dest Ports, IPv6 Label

- Note that ethernet-type MUST be 34525
- Note that IP Protocol MUST be 6

```
<?xml version="1.0" encoding="UTF-8" standalone="no"?>
<flow xmlns="urn:opendaylight:flow:inventory">
    <strict>false</strict>
    <flow-name>FooXf23</flow-name>
    <id>146</id>
    <cookie_mask>255</cookie_mask>
    <cookie>23</cookie>
    <table_id>2</table_id>
    <priority>2</priority>
    <hard-timeout>1200</hard-timeout>
    <idle-timeout>3400</idle-timeout>
    <installHw>false</installHw>
    <instructions>
        <instruction>
            <order>0</order>
            <apply-actions>
                <action>
                    <order>0</order>
                    <dec-nw-ttl/>
                </action>
            </apply-actions>
        </instruction>
    </instructions>
    <match>
        <ethernet-match>
            <ethernet-type>
                <type>34525</type>
            </ethernet-type>
        </ethernet-match>
        <ipv6-source>1234:5678:9ABC:DEF0:FDCD:A987:6543:210F/76</ipv6-source>
        <ipv6-destination>fe80:2acf:e9ff:fe21::6431/94</ipv6-destination>
        <metadata>
            <metadata>12345</metadata>
        </metadata>
        <ipv6-label>
            <ipv6-flabel>33</ipv6-flabel>
        </ipv6-label>
        <ip-match>
            <ip-protocol>6</ip-protocol>
            <ip-dscp>60</ip-dscp>
            <ip-ecn>3</ip-ecn>
        </ip-match>
        <tcp-source-port>183</tcp-source-port>
        <tcp-destination-port>8080</tcp-destination-port>
    </match>
</flow>
```

Tunnel ID

```
<?xml version="1.0" encoding="UTF-8" standalone="no"?>
<flow xmlns="urn:opendaylight:flow:inventory">
    <strict>false</strict>
    <flow-name>FooXf24</flow-name>
    <id>147</id>
    <cookie_mask>255</cookie_mask>
```

(continues on next page)

(continued from previous page)

```
    <cookie>24</cookie>
    <table_id>2</table_id>
    <priority>2</priority>
    <hard-timeout>1200</hard-timeout>
    <idle-timeout>3400</idle-timeout>
    <installHw>false</installHw>
    <instructions>
        <instruction>
            <order>0</order>
            <apply-actions>
                <action>
                    <order>0</order>
                    <dec-nw-ttl/>
                </action>
            </apply-actions>
        </instruction>
    </instructions>
    <match>
        <tunnel>
            <tunnel-id>2591</tunnel-id>
        </tunnel>
    </match>
</flow>
```

IPv6 Src & Dest Addresses, Metadata, IP DSCP, IP ECN, ICMPv6 Type & Code, IPv6 Label

- Note that ethernet-type MUST be 34525
- Note that IP Protocol MUST be 58

```
<?xml version="1.0" encoding="UTF-8" standalone="no"?>
<flow xmlns="urn:opendaylight:flow:inventory">
    <strict>false</strict>
    <flow-name>FooXf25</flow-name>
    <id>148</id>
    <cookie_mask>255</cookie_mask>
    <cookie>25</cookie>
    <table_id>2</table_id>
    <priority>2</priority>
    <hard-timeout>1200</hard-timeout>
    <idle-timeout>3400</idle-timeout>
    <installHw>false</installHw>
    <instructions>
        <instruction>
            <order>0</order>
            <apply-actions>
                <action>
                    <order>0</order>
                    <dec-nw-ttl/>
                </action>
            </apply-actions>
        </instruction>
    </instructions>
    <match>
        <ethernet-match>
```

(continues on next page)

(continued from previous page)

```
            <ethernet-type>
                <type>34525</type>
            </ethernet-type>
        </ethernet-match>
        <ipv6-source>1234:5678:9ABC:DEF0:FDCD:A987:6543:210F/76</ipv6-source>
        <ipv6-destination>fe80:2acf:e9ff:fe21::6431/94</ipv6-destination>
        <metadata>
            <metadata>12345</metadata>
        </metadata>
        <ipv6-label>
            <ipv6-flabel>33</ipv6-flabel>
        </ipv6-label>
        <ip-match>
            <ip-protocol>58</ip-protocol>
            <ip-dscp>60</ip-dscp>
            <ip-ecn>3</ip-ecn>
        </ip-match>
        <icmpv6-match>
            <icmpv6-type>6</icmpv6-type>
            <icmpv6-code>3</icmpv6-code>
        </icmpv6-match>
    </match>
</flow>
```

IPv6 Src & Dest Addresses, Metadata, IP DSCP, IP ECN, TCP Src & Dst Ports, IPv6 Label, IPv6 Ext Header

- Note that ethernet-type MUST be 34525
- Note that IP Protocol MUST be 58

```
<?xml version="1.0" encoding="UTF-8" standalone="no"?>
<flow xmlns="urn:opendaylight:flow:inventory">
    <strict>false</strict>
    <flow-name>FooXf27</flow-name>
    <id>150</id>
    <cookie_mask>255</cookie_mask>
    <cookie>27</cookie>
    <table_id>2</table_id>
    <priority>2</priority>
    <hard-timeout>1200</hard-timeout>
    <idle-timeout>3400</idle-timeout>
    <installHw>false</installHw>
    <instructions>
        <instruction>
            <order>0</order>
            <apply-actions>
                <action>
                    <order>0</order>
                    <dec-nw-ttl/>
                </action>
            </apply-actions>
        </instruction>
    </instructions>
    <match>
```

(continues on next page)

(continued from previous page)

```
        <ethernet-match>
            <ethernet-type>
                <type>34525</type>
            </ethernet-type>
        </ethernet-match>
        <ipv6-source>1234:5678:9ABC:DEF0:FDCD:A987:6543:210F/76</ipv6-source>
        <ipv6-destination>fe80:2acf:e9ff:fe21::6431/94</ipv6-destination>
        <metadata>
            <metadata>12345</metadata>
        </metadata>
        <ipv6-label>
            <ipv6-flabel>33</ipv6-flabel>
        </ipv6-label>
        <ipv6-ext-header>
            <ipv6-exthdr>0</ipv6-exthdr>
        </ipv6-ext-header>
        <ip-match>
            <ip-protocol>6</ip-protocol>
            <ip-dscp>60</ip-dscp>
            <ip-ecn>3</ip-ecn>
        </ip-match>
        <tcp-source-port>183</tcp-source-port>
        <tcp-destination-port>8080</tcp-destination-port>
    </match>
</flow>
```

Actions

The format of the XML that describes OpenFlow actions is determined by the opendaylight-action-types yang model: .

Apply Actions

Output to TABLE

```
<?xml version="1.0" encoding="UTF-8" standalone="no"?>
<flow xmlns="urn:opendaylight:flow:inventory">
    <strict>false</strict>
    <flow-name>FooXf101</flow-name>
    <id>256</id>
    <cookie_mask>255</cookie_mask>
    <cookie>101</cookie>
    <table_id>2</table_id>
    <priority>2</priority>
    <hard-timeout>1200</hard-timeout>
    <idle-timeout>3400</idle-timeout>
    <installHw>false</installHw>
    <instructions>
        <instruction>
            <order>0</order>
            <apply-actions>
                <action>
```

(continues on next page)

(continued from previous page)

```
                    <order>0</order>
                    <output-action>
                        <output-node-connector>TABLE</output-node-connector>
                        <max-length>60</max-length>
                    </output-action>
                </action>
            </apply-actions>
        </instruction>
    </instructions>
    <match>
        <ethernet-match>
            <ethernet-type>
                <type>34525</type>
            </ethernet-type>
        </ethernet-match>
        <ipv6-source>1234:5678:9ABC:DEF0:FDCD:A987:6543:210F/76</ipv6-source>
        <ipv6-destination>fe80:2acf:e9ff:fe21::6431/94</ipv6-destination>
        <metadata>
            <metadata>12345</metadata>
        </metadata>
        <ip-match>
            <ip-protocol>6</ip-protocol>
            <ip-dscp>60</ip-dscp>
            <ip-ecn>3</ip-ecn>
        </ip-match>
        <tcp-source-port>183</tcp-source-port>
        <tcp-destination-port>8080</tcp-destination-port>
    </match>
</flow>
```

Output to INPORT

```
<?xml version="1.0" encoding="UTF-8" standalone="no"?>
<flow xmlns="urn:opendaylight:flow:inventory">
    <strict>false</strict>
    <flow-name>FooXf102</flow-name>
    <id>257</id>
    <cookie_mask>255</cookie_mask>
    <cookie>102</cookie>
    <table_id>2</table_id>
    <priority>2</priority>
    <hard-timeout>1200</hard-timeout>
    <idle-timeout>3400</idle-timeout>
    <installHw>false</installHw>
    <instructions>
        <instruction>
            <order>0</order>
            <apply-actions>
                <action>
                    <order>0</order>
                    <output-action>
                        <output-node-connector>INPORT</output-node-connector>
                        <max-length>60</max-length>
                    </output-action>
```

(continues on next page)

(continued from previous page)

```
                </action>
7            </apply-actions>
        </instruction>
    </instructions>
    <match>
        <ethernet-match>
            <ethernet-type>
                <type>2048</type>
            </ethernet-type>
            <ethernet-destination>
                <address>ff:ff:29:01:19:61</address>
            </ethernet-destination>
            <ethernet-source>
                <address>00:00:00:11:23:ae</address>
            </ethernet-source>
        </ethernet-match>
        <ipv4-source>17.1.2.3/8</ipv4-source>
        <ipv4-destination>172.168.5.6/16</ipv4-destination>
        <ip-match>
            <ip-protocol>6</ip-protocol>
            <ip-dscp>2</ip-dscp>
            <ip-ecn>2</ip-ecn>
        </ip-match>
        <tcp-source-port>25364</tcp-source-port>
        <tcp-destination-port>8080</tcp-destination-port>
    </match>
</flow>
```

Output to Physical Port

```
<?xml version="1.0" encoding="UTF-8" standalone="no"?>
<flow xmlns="urn:opendaylight:flow:inventory">
    <strict>false</strict>
    <flow-name>FooXf103</flow-name>
    <id>258</id>
    <cookie_mask>255</cookie_mask>
    <cookie>103</cookie>
    <table_id>2</table_id>
    <priority>2</priority>
    <hard-timeout>1200</hard-timeout>
    <idle-timeout>3400</idle-timeout>
    <installHw>false</installHw>
    <instructions>
        <instruction>
            <order>0</order>
            <apply-actions>
                <action>
                    <order>0</order>
                    <output-action>
                        <output-node-connector>1</output-node-connector>
                        <max-length>60</max-length>
                    </output-action>
                </action>
            </apply-actions>
```

(continues on next page)

(continued from previous page)

```
        </instruction>
    </instructions>
    <match>
        <ethernet-match>
            <ethernet-type>
                <type>2048</type>
            </ethernet-type>
            <ethernet-destination>
                <address>ff:ff:29:01:19:61</address>
            </ethernet-destination>
            <ethernet-source>
                <address>00:00:00:11:23:ae</address>
            </ethernet-source>
        </ethernet-match>
        <ipv4-source>17.1.2.3/8</ipv4-source>
        <ipv4-destination>172.168.5.6/16</ipv4-destination>
        <ip-match>
            <ip-protocol>6</ip-protocol>
            <ip-dscp>2</ip-dscp>
            <ip-ecn>2</ip-ecn>
        </ip-match>
        <tcp-source-port>25364</tcp-source-port>
        <tcp-destination-port>8080</tcp-destination-port>
    </match>
</flow>
```

Output to LOCAL

```
<?xml version="1.0" encoding="UTF-8" standalone="no"?>
<flow xmlns="urn:opendaylight:flow:inventory">
    <strict>false</strict>
    <flow-name>FooXf104</flow-name>
    <id>259</id>
    <cookie_mask>255</cookie_mask>
    <cookie>104</cookie>
    <table_id>2</table_id>
    <priority>2</priority>
    <hard-timeout>1200</hard-timeout>
    <idle-timeout>3400</idle-timeout>
    <installHw>false</installHw>
    <instructions>
        <instruction>
            <order>0</order>
            <apply-actions>
                <action>
                    <order>0</order>
                    <output-action>
                        <output-node-connector>LOCAL</output-node-connector>
                        <max-length>60</max-length>
                    </output-action>
                </action>
            </apply-actions>
        </instruction>
    </instructions>
```

(continues on next page)

(continued from previous page)

```
    <match>
        <ethernet-match>
            <ethernet-type>
                <type>34525</type>
            </ethernet-type>
        </ethernet-match>
        <ipv6-source>1234:5678:9ABC:DEF0:FDCD:A987:6543:210F/76</ipv6-source>
        <ipv6-destination>fe80:2acf:e9ff:fe21::6431/94</ipv6-destination>
        <metadata>
            <metadata>12345</metadata>
        </metadata>
        <ip-match>
            <ip-protocol>6</ip-protocol>
            <ip-dscp>60</ip-dscp>
            <ip-ecn>3</ip-ecn>
        </ip-match>
        <tcp-source-port>183</tcp-source-port>
        <tcp-destination-port>8080</tcp-destination-port>
    </match>
</flow>
```

Output to NORMAL

```
<?xml version="1.0" encoding="UTF-8" standalone="no"?>
<flow xmlns="urn:opendaylight:flow:inventory">
    <strict>false</strict>
    <flow-name>FooXf105</flow-name>
    <id>260</id>
    <cookie_mask>255</cookie_mask>
    <cookie>105</cookie>
    <table_id>2</table_id>
    <priority>2</priority>
    <hard-timeout>1200</hard-timeout>
    <idle-timeout>3400</idle-timeout>
    <installHw>false</installHw>
    <instructions>
        <instruction>
            <order>0</order>
            <apply-actions>
                <action>
                    <order>0</order>
                    <output-action>
                        <output-node-connector>NORMAL</output-node-connector>
                        <max-length>60</max-length>
                    </output-action>
                </action>
            </apply-actions>
        </instruction>
    </instructions>
    <match>
        <ethernet-match>
            <ethernet-type>
                <type>34525</type>
            </ethernet-type>
```

(continues on next page)

(continued from previous page)

```
        </ethernet-match>
        <ipv6-source>1234:5678:9ABC:DEF0:FDCD:A987:6543:210F/84</ipv6-source>
        <ipv6-destination>fe80:2acf:e9ff:fe21::6431/90</ipv6-destination>
        <metadata>
            <metadata>12345</metadata>
        </metadata>
        <ip-match>
            <ip-protocol>6</ip-protocol>
            <ip-dscp>45</ip-dscp>
            <ip-ecn>2</ip-ecn>
        </ip-match>
        <tcp-source-port>20345</tcp-source-port>
        <tcp-destination-port>80</tcp-destination-port>
    </match>
</flow>
```

Output to FLOOD

```
<?xml version="1.0" encoding="UTF-8" standalone="no"?>
<flow xmlns="urn:opendaylight:flow:inventory">
    <strict>false</strict>
    <flow-name>FooXf106</flow-name>
    <id>261</id>
    <cookie_mask>255</cookie_mask>
    <cookie>106</cookie>
    <table_id>2</table_id>
    <priority>2</priority>
    <hard-timeout>1200</hard-timeout>
    <idle-timeout>3400</idle-timeout>
    <installHw>false</installHw>
    <instructions>
        <instruction>
            <order>0</order>
            <apply-actions>
                <action>
                    <order>0</order>
                    <output-action>
                        <output-node-connector>FLOOD</output-node-connector>
                        <max-length>60</max-length>
                    </output-action>
                </action>
            </apply-actions>
        </instruction>
    </instructions>
    <match>
        <ethernet-match>
            <ethernet-type>
                <type>34525</type>
            </ethernet-type>
        </ethernet-match>
        <ipv6-source>1234:5678:9ABC:DEF0:FDCD:A987:6543:210F/100</ipv6-source>
        <ipv6-destination>fe80:2acf:e9ff:fe21::6431/67</ipv6-destination>
        <metadata>
            <metadata>12345</metadata>
```

(continues on next page)

(continued from previous page)

```
        </metadata>
        <ip-match>
            <ip-protocol>6</ip-protocol>
            <ip-dscp>45</ip-dscp>
            <ip-ecn>2</ip-ecn>
        </ip-match>
        <tcp-source-port>20345</tcp-source-port>
        <tcp-destination-port>80</tcp-destination-port>
    </match>
</flow>
```

Output to ALL

```
<?xml version="1.0" encoding="UTF-8" standalone="no"?>
<flow xmlns="urn:opendaylight:flow:inventory">
    <strict>false</strict>
    <flow-name>FooXf107</flow-name>
    <id>262</id>
    <cookie_mask>255</cookie_mask>
    <cookie>107</cookie>
    <table_id>2</table_id>
    <priority>2</priority>
    <hard-timeout>1200</hard-timeout>
    <idle-timeout>3400</idle-timeout>
    <installHw>false</installHw>
    <instructions>
        <instruction>
            <order>0</order>
            <apply-actions>
                <action>
                    <order>0</order>
                    <output-action>
                        <output-node-connector>ALL</output-node-connector>
                        <max-length>60</max-length>
                    </output-action>
                </action>
            </apply-actions>
        </instruction>
    </instructions>
    <match>
        <ethernet-match>
            <ethernet-type>
                <type>2048</type>
            </ethernet-type>
            <ethernet-destination>
                <address>20:14:29:01:19:61</address>
            </ethernet-destination>
            <ethernet-source>
                <address>00:00:00:11:23:ae</address>
            </ethernet-source>
        </ethernet-match>
        <ipv4-source>19.1.2.3/10</ipv4-source>
        <ipv4-destination>172.168.5.6/18</ipv4-destination>
        <ip-match>
```

(continues on next page)

(continued from previous page)

```
            <ip-protocol>17</ip-protocol>
            <ip-dscp>8</ip-dscp>
            <ip-ecn>3</ip-ecn>
        </ip-match>
        <udp-source-port>25364</udp-source-port>
        <udp-destination-port>8080</udp-destination-port>
        <in-port>0</in-port>
    </match>
</flow>
```

Output to CONTROLLER

```
<?xml version="1.0" encoding="UTF-8" standalone="no"?>
<flow xmlns="urn:opendaylight:flow:inventory">
    <strict>false</strict>
    <flow-name>FooXf108</flow-name>
    <id>263</id>
    <cookie_mask>255</cookie_mask>
    <cookie>108</cookie>
    <table_id>2</table_id>
    <priority>2</priority>
    <hard-timeout>1200</hard-timeout>
    <idle-timeout>3400</idle-timeout>
    <installHw>false</installHw>
    <instructions>
        <instruction>
            <order>0</order>
            <apply-actions>
                <action>
                    <order>0</order>
                    <output-action>
                        <output-node-connector>CONTROLLER</output-node-connector>
                        <max-length>60</max-length>
                    </output-action>
                </action>
            </apply-actions>
        </instruction>
    </instructions>
    <match>
        <ethernet-match>
            <ethernet-type>
                <type>2048</type>
            </ethernet-type>
            <ethernet-destination>
                <address>20:14:29:01:19:61</address>
            </ethernet-destination>
            <ethernet-source>
                <address>00:00:00:11:23:ae</address>
            </ethernet-source>
        </ethernet-match>
        <ipv4-source>19.1.2.3/10</ipv4-source>
        <ipv4-destination>172.168.5.6/18</ipv4-destination>
        <ip-match>
            <ip-protocol>17</ip-protocol>
```

(continues on next page)

(continued from previous page)

```
            <ip-dscp>8</ip-dscp>
            <ip-ecn>3</ip-ecn>
        </ip-match>
        <udp-source-port>25364</udp-source-port>
        <udp-destination-port>8080</udp-destination-port>
        <in-port>0</in-port>
    </match>
</flow>
```

Output to ANY

```
<?xml version="1.0" encoding="UTF-8" standalone="no"?>
<flow xmlns="urn:opendaylight:flow:inventory">
    <strict>false</strict>
    <flow-name>FooXf109</flow-name>
    <id>264</id>
    <cookie_mask>255</cookie_mask>
    <cookie>109</cookie>
    <table_id>2</table_id>
    <priority>2</priority>
    <hard-timeout>1200</hard-timeout>
    <idle-timeout>3400</idle-timeout>
    <installHw>false</installHw>
    <instructions>
        <instruction>
            <order>0</order>
            <apply-actions>
                <action>
                    <order>0</order>
                    <output-action>
                        <output-node-connector>ANY</output-node-connector>
                        <max-length>60</max-length>
                    </output-action>
                </action>
            </apply-actions>
        </instruction>
    </instructions>
    <match>
        <ethernet-match>
            <ethernet-type>
                <type>2048</type>
            </ethernet-type>
            <ethernet-destination>
                <address>20:14:29:01:19:61</address>
            </ethernet-destination>
            <ethernet-source>
                <address>00:00:00:11:23:ae</address>
            </ethernet-source>
        </ethernet-match>
        <ipv4-source>19.1.2.3/10</ipv4-source>
        <ipv4-destination>172.168.5.6/18</ipv4-destination>
        <ip-match>
            <ip-protocol>17</ip-protocol>
            <ip-dscp>8</ip-dscp>
```

(continues on next page)

(continued from previous page)

```
            <ip-ecn>3</ip-ecn>
        </ip-match>
        <udp-source-port>25364</udp-source-port>
        <udp-destination-port>8080</udp-destination-port>
        <in-port>0</in-port>
    </match>
</flow>
```

Push VLAN

```
<?xml version="1.0" encoding="UTF-8" standalone="no"?>
<flow xmlns="urn:opendaylight:flow:inventory">
   <strict>false</strict>
   <instructions>
       <instruction>
           <order>0</order>
           <apply-actions>
              <action>
                 <push-vlan-action>
                     <ethernet-type>33024</ethernet-type>
                 </push-vlan-action>
                 <order>0</order>
              </action>
               <action>
                   <set-field>
                       <vlan-match>
                            <vlan-id>
                                <vlan-id>79</vlan-id>
                                <vlan-id-present>true</vlan-id-present>
                            </vlan-id>
                       </vlan-match>
                   </set-field>
                   <order>1</order>
               </action>
               <action>
                   <output-action>
                       <output-node-connector>5</output-node-connector>
                   </output-action>
                   <order>2</order>
               </action>
           </apply-actions>
       </instruction>
   </instructions>
   <table_id>0</table_id>
   <id>31</id>
   <match>
       <ethernet-match>
           <ethernet-type>
               <type>2048</type>
           </ethernet-type>
           <ethernet-destination>
               <address>FF:FF:29:01:19:61</address>
           </ethernet-destination>
           <ethernet-source>
```

(continues on next page)

(continued from previous page)

```
                <address>00:00:00:11:23:AE</address>
            </ethernet-source>
        </ethernet-match>
      <in-port>1</in-port>
    </match>
    <flow-name>vlan_flow</flow-name>
    <priority>2</priority>
</flow>
```

Push MPLS

```
<?xml version="1.0" encoding="UTF-8" standalone="no"?>
<flow
    xmlns="urn:opendaylight:flow:inventory">
    <flow-name>push-mpls-action</flow-name>
    <instructions>
        <instruction>
            <order>3</order>
            <apply-actions>
                <action>
                    <push-mpls-action>
                        <ethernet-type>34887</ethernet-type>
                    </push-mpls-action>
                    <order>0</order>
                </action>
                <action>
                    <set-field>
                        <protocol-match-fields>
                            <mpls-label>27</mpls-label>
                        </protocol-match-fields>
                    </set-field>
                    <order>1</order>
                </action>
                <action>
                    <output-action>
                        <output-node-connector>2</output-node-connector>
                    </output-action>
                    <order>2</order>
                </action>
            </apply-actions>
        </instruction>
    </instructions>
    <strict>false</strict>
    <id>100</id>
    <match>
        <ethernet-match>
            <ethernet-type>
                <type>2048</type>
            </ethernet-type>
        </ethernet-match>
        <in-port>1</in-port>
        <ipv4-destination>10.0.0.4/32</ipv4-destination>
    </match>
    <idle-timeout>0</idle-timeout>
```

(continues on next page)

(continued from previous page)

```
    <cookie_mask>255</cookie_mask>
    <cookie>401</cookie>
    <priority>8</priority>
    <hard-timeout>0</hard-timeout>
    <installHw>false</installHw>
    <table_id>0</table_id>
</flow>
```

Swap MPLS

- Note that ethernet-type MUST be 34887

```
<?xml version="1.0" encoding="UTF-8" standalone="no"?>
<flow
    xmlns="urn:opendaylight:flow:inventory">
    <flow-name>push-mpls-action</flow-name>
    <instructions>
        <instruction>
            <order>2</order>
            <apply-actions>
                <action>
                    <set-field>
                        <protocol-match-fields>
                            <mpls-label>37</mpls-label>
                        </protocol-match-fields>
                    </set-field>
                    <order>1</order>
                </action>
                <action>
                    <output-action>
                        <output-node-connector>2</output-node-connector>
                    </output-action>
                    <order>2</order>
                </action>
            </apply-actions>
        </instruction>
    </instructions>
    <strict>false</strict>
    <id>101</id>
    <match>
        <ethernet-match>
            <ethernet-type>
                <type>34887</type>
            </ethernet-type>
        </ethernet-match>
        <in-port>1</in-port>
        <protocol-match-fields>
            <mpls-label>27</mpls-label>
        </protocol-match-fields>
    </match>
    <idle-timeout>0</idle-timeout>
    <cookie_mask>255</cookie_mask>
    <cookie>401</cookie>
    <priority>8</priority>
    <hard-timeout>0</hard-timeout>
```

(continues on next page)

(continued from previous page)

```
    <installHw>false</installHw>
    <table_id>0</table_id>
</flow>
```

Pop MPLS

- Note that ethernet-type MUST be 34887
- Issue with OVS 2.1 OVS fix

```
<?xml version="1.0" encoding="UTF-8" standalone="no"?>
<flow
    xmlns="urn:opendaylight:flow:inventory">
    <flow-name>FooXf10</flow-name>
    <instructions>
        <instruction>
            <order>0</order>
            <apply-actions>
                <action>
                    <pop-mpls-action>
                        <ethernet-type>2048</ethernet-type>
                    </pop-mpls-action>
                    <order>1</order>
                </action>
                <action>
                    <output-action>
                        <output-node-connector>2</output-node-connector>
                        <max-length>60</max-length>
                    </output-action>
                    <order>2</order>
                </action>
            </apply-actions>
        </instruction>
    </instructions>
    <id>11</id>
    <strict>false</strict>
    <match>
        <ethernet-match>
            <ethernet-type>
                <type>34887</type>
            </ethernet-type>
        </ethernet-match>
        <in-port>1</in-port>
        <protocol-match-fields>
            <mpls-label>37</mpls-label>
        </protocol-match-fields>
    </match>
    <idle-timeout>0</idle-timeout>
    <cookie>889</cookie>
    <cookie_mask>255</cookie_mask>
    <installHw>false</installHw>
    <hard-timeout>0</hard-timeout>
    <priority>10</priority>
    <table_id>0</table_id>
</flow>
```

Learn

- Nicira extension defined in https://github.com/osrg/openvswitch/blob/master/include/openflow/nicira-ext.h
- Example section is - https://github.com/osrg/openvswitch/blob/master/include/openflow/nicira-ext.h#L788

```
<flow>
  <id>ICMP_Ingress258a5a5ad-08a8-4ff7-98f5-ef0b96ca3bb8</id>
  <hard-timeout>0</hard-timeout>
  <idle-timeout>0</idle-timeout>
  <match>
    <ethernet-match>
      <ethernet-type>
        <type>2048</type>
      </ethernet-type>
    </ethernet-match>
    <metadata>
      <metadata>2199023255552</metadata>
      <metadata-mask>23058419097020661 76</metadata-mask>
    </metadata>
    <ip-match>
      <ip-protocol>1</ip-protocol>
    </ip-match>
  </match>
  <cookie>110100480</cookie>
  <instructions>
    <instruction>
      <order>0</order>
      <apply-actions>
        <action>
          <order>1</order>
          <nx-resubmit
            xmlns="urn:opendaylight:openflowplugin:extension:nicira:action">
            <table>220</table>
          </nx-resubmit>
        </action>
        <action>
          <order>0</order>
          <nx-learn
            xmlns="urn:opendaylight:openflowplugin:extension:nicira:action">
            <idle-timeout>60</idle-timeout>
            <fin-idle-timeout>0</fin-idle-timeout>
            <hard-timeout>60</hard-timeout>
            <flags>0</flags>
            <table-id>41</table-id>
            <priority>61010</priority>
            <fin-hard-timeout>0</fin-hard-timeout>
            <flow-mods>
              <flow-mod-add-match-from-value>
                <src-ofs>0</src-ofs>
                <value>2048</value>
                <src-field>1538</src-field>
                <flow-mod-num-bits>16</flow-mod-num-bits>
              </flow-mod-add-match-from-value>
            </flow-mods>
            <flow-mods>
              <flow-mod-add-match-from-field>
                <src-ofs>0</src-ofs>
```

(continues on next page)

(continued from previous page)

```
                <dst-ofs>0</dst-ofs>
                <dst-field>4100</dst-field>
                <src-field>3588</src-field>
                <flow-mod-num-bits>32</flow-mod-num-bits>
              </flow-mod-add-match-from-field>
            </flow-mods>
            <flow-mods>
              <flow-mod-add-match-from-field>
                <src-ofs>0</src-ofs>
                <dst-ofs>0</dst-ofs>
                <dst-field>518</dst-field>
                <src-field>1030</src-field>
                <flow-mod-num-bits>48</flow-mod-num-bits>
              </flow-mod-add-match-from-field>
            </flow-mods>
            <flow-mods>
              <flow-mod-add-match-from-field>
                <src-ofs>0</src-ofs>
                <dst-ofs>0</dst-ofs>
                <dst-field>3073</dst-field>
                <src-field>3073</src-field>
                <flow-mod-num-bits>8</flow-mod-num-bits>
              </flow-mod-add-match-from-field>
            </flow-mods>
            <flow-mods>
              <flow-mod-copy-value-into-field>
                <dst-ofs>0</dst-ofs>
                <value>1</value>
                <dst-field>65540</dst-field>
                <flow-mod-num-bits>8</flow-mod-num-bits>
              </flow-mod-copy-value-into-field>
            </flow-mods>
            <cookie>110100480</cookie>
          </nx-learn>
        </action>
      </apply-actions>
    </instruction>
  </instructions>
  <installHw>true</installHw>
  <barrier>false</barrier>
  <strict>false</strict>
  <priority>61010</priority>
  <table_id>253</table_id>
  <flow-name>ACL</flow-name>
</flow>
```

Opendaylight OpenFlow Plugin: Troubleshooting

empty section

OpFlex agent-ovs User Guide

Introduction

agent-ovs is a policy agent that works with OVS to enforce a group-based policy networking model with locally attached virtual machines or containers. The policy agent is designed to work well with orchestration tools like OpenStack.

Agent Configuration

The agent configuration is handled using its config file which is by default found at "/etc/opflex-agent-ovs/opflex-agent-ovs.conf"

Here is an example configuration file that documents the available options:

```
{
    // Logging configuration
    // "log": {
    //    "level": "info"
    // },

    // Configuration related to the OpFlex protocol
    "opflex": {
        // The policy domain for this agent.
        "domain": "openstack",

        // The unique name in the policy domain for this agent.
        "name": "example-agent",

        // a list of peers to connect to, by hostname and port.  One
        // peer, or an anycast pseudo-peer, is sufficient to bootstrap
        // the connection without needing an exhaustive list of all
        // peers.
        "peers": [
            // EXAMPLE:
            {"hostname": "10.0.0.30", "port": 8009}
        ],

        "ssl": {
            // SSL mode.  Possible values:
            // disabled: communicate without encryption
            // encrypted: encrypt but do not verify peers
            // secure: encrypt and verify peer certificates
            "mode": "disabled",

            // The path to a directory containing trusted certificate
            // authority public certificates, or a file containing a
            // specific CA certificate.
            "ca-store": "/etc/ssl/certs/"
        },

        "inspector": {
            // Enable the MODB inspector service, which allows
            // inspecting the state of the managed object database.
        // Default: enabled
            "enabled": true,
```

(continues on next page)

(continued from previous page)

```
        // Listen on the specified socket for the inspector
    // Default /var/run/opflex-agent-ovs-inspect.sock
        "socket-name": "/var/run/opflex-agent-ovs-inspect.sock"
    }
},

// Endpoint sources provide metadata about local endpoints
"endpoint-sources": {
    // Filesystem path to monitor for endpoint information
    "filesystem": ["/var/lib/opflex-agent-ovs/endpoints"]
},

// Renderers enforce policy obtained via OpFlex.
"renderers": {
    // Stitched-mode renderer for interoperating with a
    // hardware fabric such as ACI
    // EXAMPLE:
    "stitched-mode": {
        "ovs-bridge-name": "br0",

        // Set encapsulation type.  Must set either vxlan or vlan.
        "encap": {
            // Encapsulate traffic with VXLAN.
            "vxlan" : {
                // The name of the tunnel interface in OVS
                "encap-iface": "br0_vxlan0",

                // The name of the interface whose IP should be used
                // as the source IP in encapsulated traffic.
                "uplink-iface": "eth0.4093",

                // The vlan tag, if any, used on the uplink interface.
                // Set to zero or omit if the uplink is untagged.
                "uplink-vlan": 4093,

                // The IP address used for the destination IP in
                // the encapsulated traffic.  This should be an
                // anycast IP address understood by the upstream
                // stiched-mode fabric.
                "remote-ip": "10.0.0.32",

                // UDP port number of the encapsulated traffic.
                "remote-port": 8472
            }

            // Encapsulate traffic with a locally-significant VLAN
            // tag
            // EXAMPLE:
            // "vlan" : {
            //     // The name of the uplink interface in OVS
            //     "encap-iface": "team0"
            // }
        },

        // Configure forwarding policy
        "forwarding": {
```

(continues on next page)

(continued from previous page)

```
                // Configure the virtual distributed router
                "virtual-router": {
                    // Enable virtual distributed router.  Set to true
                    // to enable or false to disable.  Default true.
                    "enabled": true,

                    // Override MAC address for virtual router.
                    // Default is "00:22:bd:f8:19:ff"
                    "mac": "00:22:bd:f8:19:ff",

                    // Configure IPv6-related settings for the virtual
                    // router
                    "ipv6" : {
                        // Send router advertisement messages in
                        // response to router solicitation requests as
                        // well as unsolicited advertisements.  This
                        // is not required in stitched mode since the
                        // hardware router will send them.
                        "router-advertisement": true
                    }
                },

                // Configure virtual distributed DHCP server
                "virtual-dhcp": {
                    // Enable virtual distributed DHCP server.  Set to
                    // true to enable or false to disable.  Default
                    // true.
                    "enabled": true,

                    // Override MAC address for virtual dhcp server.
                    // Default is "00:22:bd:f8:19:ff"
                    "mac": "00:22:bd:f8:19:ff"
                },

                "endpoint-advertisements": {
                    // Enable generation of periodic ARP/NDP
                    // advertisements for endpoints.  Default true.
                    "enabled": "true"
                }
            },

            // Location to store cached IDs for managing flow state
            "flowid-cache-dir": "/var/lib/opflex-agent-ovs/ids"
        }
    }
}
```

Endpoint Registration

The agent learns about endpoints using endpoint metadata files located by default in "/var/lib/opflex-agent-ovs/endpoints".

These are JSON-format files such as the (unusually complex) example below:

```
{
    "uuid": "83f18f0b-80f7-46e2-b06c-4d9487b0c754",
    "policy-space-name": "test",
    "endpoint-group-name": "group1",
    "interface-name": "veth0",
    "ip": [
        "10.0.0.1", "fd8f:69d8:c12c:ca62::1"
    ],
    "dhcp4": {
        "ip": "10.200.44.2",
        "prefix-len": 24,
        "routers": ["10.200.44.1"],
        "dns-servers": ["8.8.8.8", "8.8.4.4"],
        "domain": "example.com",
        "static-routes": [
            {
                "dest": "169.254.169.0",
                "dest-prefix": 24,
                "next-hop": "10.0.0.1"
            }
        ]
    },
    "dhcp6": {
        "dns-servers": ["2001:4860:4860::8888", "2001:4860:4860::8844"],
        "search-list": ["test1.example.com", "example.com"]
    },
    "ip-address-mapping": [
        {
           "uuid": "91c5b217-d244-432c-922d-533c6c036ab4",
           "floating-ip": "5.5.5.1",
           "mapped-ip": "10.0.0.1",
           "policy-space-name": "common",
           "endpoint-group-name": "nat-epg"
        },
        {
           "uuid": "22bfdc01-a390-4b6f-9b10-624d4ccb957b",
           "floating-ip": "fdf1:9f86:d1af:6cc9::1",
           "mapped-ip": "fd8f:69d8:c12c:ca62::1",
           "policy-space-name": "common",
           "endpoint-group-name": "nat-epg"
        }
    ],
    "mac": "00:00:00:00:00:01",
    "promiscuous-mode": false
}
```

The possible parameters for these files are:

uuid A globally unique ID for the endpoint

endpoint-group-name The name of the endpoint group for the endpoint

policy-space-name The name of the policy space for the endpoint group.

interface-name The name of the OVS interface to which the endpoint is attached

ip A list of strings contains either IPv4 or IPv6 addresses that the endpoint is allowed to use

mac The MAC address for the endpoint's interface.

promiscuous-mode Allow traffic from this VM to bypass default port security

dhcp4 A distributed DHCPv4 configuration block (see below)

dhcp6 A distributed DHCPv6 configuration block (see below)

ip-address-mapping A list of IP address mapping configuration blocks (see below)

DHCPv4 configuration blocks can contain the following parameters:

ip the IP address to return with DHCP. Must be one of the configured IPv4 addresses.

prefix the subnet prefix length

routers a list of default gateways for the endpoint

dns a list of DNS server addresses

domain The domain name parameter to send in the DHCP reply

static-routes A list of static route configuration blocks, which contains a "dest", "dest-prefix", and "next-hop" parameters to send as static routes to the end host

DHCPv6 configuration blocks can contain the following parameters:

dns A list of DNS servers for the endpoint

search-patch The DNS search path for the endpoint

IP address mapping configuration blocks can contain the following parameters:

uuid a globally unique ID for the virtual endpoint created by the mapping.

floating-ip Map using DNAT to this floating IPv4 or IPv6 address

mapped-ip the source IPv4 or IPv6 address; must be one of the IPs assigned to the endpoint.

endpoint-group-name The name of the endpoint group for the NATed IP

policy-space-name The name of the policy space for the NATed IP

Inspector

The Opflex inspector is a useful command-line tool that will allow you to inspect the state of the managed object database for the agent for debugging and diagnosis purposes.

The command is called "gbp_inspect" and takes the following arguments:

```
# gbp_inspect -h
Usage: ./gbp_inspect [options]
Allowed options:
  -h [ --help ]                         Print this help message
  --log arg                             Log to the specified file (default
                                        standard out)
  --level arg (=warning)                Use the specified log level (default
                                        info)
  --syslog                              Log to syslog instead of file or
                                        standard out
  --socket arg (=/usr/local/var/run/opflex-agent-ovs-inspect.sock)
                                        Connect to the specified UNIX domain
                                        socket (default /usr/local/var/run/opfl
                                        ex-agent-ovs-inspect.sock)
  -q [ --query ] arg                    Query for a specific object with
                                        subjectname,uri or all objects of a
                                        specific type with subjectname
```

(continues on next page)

(continued from previous page)

```
  -r [ --recursive ]                    Retrieve the whole subtree for each
                                        returned object
  -f [ --follow-refs ]                  Follow references in returned objects
  --load arg                            Load managed objects from the specified
                                        file into the MODB view
  -o [ --output ] arg                   Output the results to the specified
                                        file (default standard out)
  -t [ --type ] arg (=tree)             Specify the output format: tree, list,
                                        or dump (default tree)
  -p [ --props ]                        Include object properties in output
```

Here are some examples of the ways to use this tool.

You can get information about the running system using one or more queries, which consist of an object model class name and optionally the URI of a specific object. The simplest query is to get a single object, nonrecursively:

```
# gbp_inspect -q DmtreeRoot
--* DmtreeRoot,/
# gbp_inspect -q GbpEpGroup
--* GbpEpGroup,/PolicyUniverse/PolicySpace/common/GbpEpGroup/nat-epg/
--* GbpEpGroup,/PolicyUniverse/PolicySpace/test/GbpEpGroup/group1/
# gbp_inspect -q GbpEpGroup,/PolicyUniverse/PolicySpace/common/GbpEpGroup/nat-epg/
--* GbpEpGroup,/PolicyUniverse/PolicySpace/common/GbpEpGroup/nat-epg/
```

You can also display all the properties for each object:

```
# gbp_inspect -p -q GbpeL24Classifier
--* GbpeL24Classifier,/PolicyUniverse/PolicySpace/test/GbpeL24Classifier/classifier4/
     {
       connectionTracking : 1 (reflexive)
       dFromPort          : 80
       dToPort            : 80
       etherT             : 2048 (ipv4)
       name               : classifier4
       prot               : 6
     }
--* GbpeL24Classifier,/PolicyUniverse/PolicySpace/test/GbpeL24Classifier/classifier3/
     {
       etherT : 34525 (ipv6)
       name   : classifier3
       order  : 100
       prot   : 58
     }
--* GbpeL24Classifier,/PolicyUniverse/PolicySpace/test/GbpeL24Classifier/classifier2/
     {
       etherT : 2048 (ipv4)
       name   : classifier2
       order  : 101
       prot   : 1
     }
```

You can also request to get the all the children of an object you query for:

```
# gbp_inspect -r -q GbpEpGroup,/PolicyUniverse/PolicySpace/common/GbpEpGroup/nat-epg/
--* GbpEpGroup,/PolicyUniverse/PolicySpace/common/GbpEpGroup/nat-epg/
  |-* GbpeInstContext,/PolicyUniverse/PolicySpace/common/GbpEpGroup/nat-epg/
↪GbpeInstContext/
```

(continues on next page)

(continued from previous page)

```
  `-* GbpEpGroupToNetworkRSrc,/PolicyUniverse/PolicySpace/common/GbpEpGroup/nat-epg/
↪GbpEpGroupToNetworkRSrc/
```

You can also follow references found in any object downloads:

```
# gbp_inspect -fr -q GbpEpGroup,/PolicyUniverse/PolicySpace/common/GbpEpGroup/nat-epg/
--* GbpEpGroup,/PolicyUniverse/PolicySpace/common/GbpEpGroup/nat-epg/
  |-* GbpeInstContext,/PolicyUniverse/PolicySpace/common/GbpEpGroup/nat-epg/
↪GbpeInstContext/
  `-* GbpEpGroupToNetworkRSrc,/PolicyUniverse/PolicySpace/common/GbpEpGroup/nat-epg/
↪GbpEpGroupToNetworkRSrc/
--* GbpFloodDomain,/PolicyUniverse/PolicySpace/common/GbpFloodDomain/fd_ext/
  `-* GbpFloodDomainToNetworkRSrc,/PolicyUniverse/PolicySpace/common/GbpFloodDomain/
↪fd_ext/GbpFloodDomainToNetworkRSrc/
--* GbpBridgeDomain,/PolicyUniverse/PolicySpace/common/GbpBridgeDomain/bd_ext/
  `-* GbpBridgeDomainToNetworkRSrc,/PolicyUniverse/PolicySpace/common/GbpBridgeDomain/
↪bd_ext/GbpBridgeDomainToNetworkRSrc/
--* GbpRoutingDomain,/PolicyUniverse/PolicySpace/common/GbpRoutingDomain/rd_ext/
  |-* GbpRoutingDomainToIntSubnetsRSrc,/PolicyUniverse/PolicySpace/common/
↪GbpRoutingDomain/rd_ext/GbpRoutingDomainToIntSubnetsRSrc/122/%2fPolicyUniverse
↪%2fPolicySpace%2fcommon%2fGbpSubnets%2fsubnets_ext%2f/
  `-* GbpForwardingBehavioralGroupToSubnetsRSrc,/PolicyUniverse/PolicySpace/common/
↪GbpRoutingDomain/rd_ext/GbpForwardingBehavioralGroupToSubnetsRSrc/
--* GbpSubnets,/PolicyUniverse/PolicySpace/common/GbpSubnets/subnets_ext/
  |-* GbpSubnet,/PolicyUniverse/PolicySpace/common/GbpSubnets/subnets_ext/GbpSubnet/
↪subnet_ext4/
  `-* GbpSubnet,/PolicyUniverse/PolicySpace/common/GbpSubnets/subnets_ext/GbpSubnet/
↪subnet_ext6/
```

OVSDB User Guide

The OVSDB project implements the OVSDB protocol (RFC 7047), as well as plugins to support OVSDB Schemas, such as the Open_vSwitch database schema and the hardware_vtep database schema.

OVSDB Plugins

Overview and Architecture

There are currently two OVSDB Southbound plugins:

- odl-ovsdb-southbound: Implements the OVSDB Open_vSwitch database schema.
- odl-ovsdb-hwvtepsouthbound: Implements the OVSDB hardware_vtep database schema.

These plugins are normally installed and used automatically by higher level applications such as odl-ovsdb-openstack; however, they can also be installed separately and used via their REST APIs as is described in the following sections.

OVSDB Southbound Plugin

The OVSDB Southbound Plugin provides support for managing OVS hosts via an OVSDB model in the MD-SAL which maps to important tables and attributes present in the Open_vSwitch schema. The OVSDB Southbound Plugin is able to connect actively or passively to OVS hosts and operate as the OVSDB manager of the OVS host. Using the

OVSDB protocol it is able to manage the OVS database (OVSDB) on the OVS host as defined by the Open_vSwitch schema.

OVSDB YANG Model

The OVSDB Southbound Plugin provides a YANG model which is based on the abstract network topology model.

The details of the OVSDB YANG model are defined in the ovsdb.yang file.

The OVSDB YANG model defines three augmentations:

ovsdb-node-augmentation This augments the network-topology node and maps primarily to the Open_vSwitch table of the OVSDB schema. The ovsdb-node-augmentation is a representation of the OVS host. It contains the following attributes.

- **connection-info** - holds the local and remote IP address and TCP port numbers for the OpenDaylight to OVSDB node connections
- **db-version** - version of the OVSDB database
- **ovs-version** - version of OVS
- **list managed-node-entry** - a list of references to ovsdb-bridge-augmentation nodes, which are the OVS bridges managed by this OVSDB node
- **list datapath-type-entry** - a list of the datapath types supported by the OVSDB node (e.g. *system*, *netdev*) - depends on newer OVS versions
- **list interface-type-entry** - a list of the interface types supported by the OVSDB node (e.g. *internal*, *vxlan*, *gre*, *dpdk*, etc.) - depends on newer OVS verions
- **list openvswitch-external-ids** - a list of the key/value pairs in the Open_vSwitch table external_ids column
- **list openvswitch-other-config** - a list of the key/value pairs in the Open_vSwitch table other_config column
- **list managery-entry** - list of manager information entries and connection status
- **list qos-entries** - list of QoS entries present in the QoS table
- **list queues** - list of queue entries present in the queue table

ovsdb-bridge-augmentation This augments the network-topology node and maps to an specific bridge in the OVSDB bridge table of the associated OVSDB node. It contains the following attributes.

- **bridge-uuid** - UUID of the OVSDB bridge
- **bridge-name** - name of the OVSDB bridge
- **bridge-openflow-node-ref** - a reference (instance-identifier) of the OpenFlow node associated with this bridge
- **list protocol-entry** - the version of OpenFlow protocol to use with the OpenFlow controller
- **list controller-entry** - a list of controller-uuid and is-connected status of the OpenFlow controllers associated with this bridge
- **datapath-id** - the datapath ID associated with this bridge on the OVSDB node
- **datapath-type** - the datapath type of this bridge
- **fail-mode** - the OVSDB fail mode setting of this bridge
- **flow-node** - a reference to the flow node corresponding to this bridge

- **managed-by** - a reference to the ovsdb-node-augmentation (OVSDB node) that is managing this bridge
- **list bridge-external-ids** - a list of the key/value pairs in the bridge table external_ids column for this bridge
- **list bridge-other-configs** - a list of the key/value pairs in the bridge table other_config column for this bridge

ovsdb-termination-point-augmentation This augments the topology termination point model. The OVSDB Southbound Plugin uses this model to represent both the OVSDB port and OVSDB interface for a given port/interface in the OVSDB schema. It contains the following attributes.

- **port-uuid** - UUID of an OVSDB port row
- **interface-uuid** - UUID of an OVSDB interface row
- **name** - name of the port and interface
- **interface-type** - the interface type
- **list options** - a list of port options
- **ofport** - the OpenFlow port number of the interface
- **ofport_request** - the requested OpenFlow port number for the interface
- **vlan-tag** - the VLAN tag value
- **list trunks** - list of VLAN tag values for trunk mode
- **vlan-mode** - the VLAN mode (e.g. access, native-tagged, native-untagged, trunk)
- **list port-external-ids** - a list of the key/value pairs in the port table external_ids column for this port
- **list interface-external-ids** - a list of the key/value pairs in the interface table external_ids interface for this interface
- **list port-other-configs** - a list of the key/value pairs in the port table other_config column for this port
- **list interface-other-configs** - a list of the key/value pairs in the interface table other_config column for this interface
- **list inteface-lldp** - LLDP Auto Attach configuration for the interface
- **qos** - UUID of the QoS entry in the QoS table assigned to this port

Getting Started

To install the OVSDB Southbound Plugin, use the following command at the Karaf console:

```
feature:install odl-ovsdb-southbound-impl-ui
```

After installing the OVSDB Southbound Plugin, and before any OVSDB topology nodes have been created, the OVSDB topology will appear as follows in the configuration and operational MD-SAL.

HTTP GET:

```
http://<controller-ip>:8181/restconf/config/network-topology:network-topology/
↪topology/ovsdb:1/
 or
http://<controller-ip>:8181/restconf/operational/network-topology:network-topology/
↪topology/ovsdb:1/
```

Result Body:

```
{
  "topology": [
    {
      "topology-id": "ovsdb:1"
    }
  ]
}
```

Where

<controller-ip> is the IP address of the OpenDaylight controller

OpenDaylight as the OVSDB Manager

An OVS host is a system which is running the OVS software and is capable of being managed by an OVSDB manager. The OVSDB Southbound Plugin is capable of connecting to an OVS host and operating as an OVSDB manager. Depending on the configuration of the OVS host, the connection of OpenDaylight to the OVS host will be active or passive.

Active Connection to OVS Hosts

An active connection is when the OVSDB Southbound Plugin initiates the connection to an OVS host. This happens when the OVS host is configured to listen for the connection (i.e. the OVSDB Southbound Plugin is active the the OVS host is passive). The OVS host is configured with the following command:

```
sudo ovs-vsctl set-manager ptcp:6640
```

This configures the OVS host to listen on TCP port 6640.

The OVSDB Southbound Plugin can be configured via the configuration MD-SAL to actively connect to an OVS host.

HTTP PUT:

```
http://<controller-ip>:8181/restconf/config/network-topology:network-topology/
↪topology/ovsdb:1/node/ovsdb:%2F%2FHOST1
```

Body:

```
{
  "network-topology:node": [
    {
      "node-id": "ovsdb://HOST1",
      "connection-info": {
        "ovsdb:remote-port": "6640",
        "ovsdb:remote-ip": "<ovs-host-ip>"
      }
    }
  ]
}
```

Where

<ovs-host-ip> is the IP address of the OVS Host

Note that the configuration assigns a *node-id* of "ovsdb://HOST1" to the OVSDB node. This *node-id* will be used as the identifier for this OVSDB node in the MD-SAL.

Query the configuration MD-SAL for the OVSDB topology.

HTTP GET:

```
http://<controller-ip>:8181/restconf/config/network-topology:network-topology/
↪topology/ovsdb:1/
```

Result Body:

```
{
  "topology": [
    {
      "topology-id": "ovsdb:1",
      "node": [
        {
          "node-id": "ovsdb://HOST1",
          "ovsdb:connection-info": {
            "remote-ip": "<ovs-host-ip>",
            "remote-port": 6640
          }
        }
      ]
    }
  ]
}
```

As a result of the OVSDB node configuration being added to the configuration MD-SAL, the OVSDB Southbound Plugin will attempt to connect with the specified OVS host. If the connection is successful, the plugin will connect to the OVS host as an OVSDB manager, query the schemas and databases supported by the OVS host, and register to monitor changes made to the OVSDB tables on the OVS host. It will also set an external id key and value in the external-ids column of the Open_vSwtich table of the OVS host which identifies the MD-SAL instance identifier of the OVSDB node. This ensures that the OVSDB node will use the same *node-id* in both the configuration and operational MD-SAL.

```
"opendaylight-iid" = "instance identifier of OVSDB node in the MD-SAL"
```

When the OVS host sends the OVSDB Southbound Plugin the first update message after the monitoring has been established, the plugin will populate the operational MD-SAL with the information it receives from the OVS host.

Query the operational MD-SAL for the OVSDB topology.

HTTP GET:

```
http://<controller-ip>:8181/restconf/operational/network-topology:network-topology/
↪topology/ovsdb:1/
```

Result Body:

```
{
  "topology": [
    {
      "topology-id": "ovsdb:1",
      "node": [
        {
          "node-id": "ovsdb://HOST1",
          "ovsdb:openvswitch-external-ids": [
            {
              "external-id-key": "opendaylight-iid",
```

(continues on next page)

(continued from previous page)

```
            "external-id-value": "/network-topology:network-topology/network-
↪topology:topology[network-topology:topology-id='ovsdb:1']/network-
↪topology:node[network-topology:node-id='ovsdb://HOST1']"
          }
        ],
        "ovsdb:connection-info": {
          "local-ip": "<controller-ip>",
          "remote-port": 6640,
          "remote-ip": "<ovs-host-ip>",
          "local-port": 39042
        },
        "ovsdb:ovs-version": "2.3.1-git4750c96",
        "ovsdb:manager-entry": [
          {
            "target": "ptcp:6640",
            "connected": true,
            "number_of_connections": 1
          }
        ]
      }
    ]
  }
}
```

To disconnect an active connection, just delete the configuration MD-SAL entry.

HTTP DELETE:

```
http://<controller-ip>:8181/restconf/config/network-topology:network-topology/
↪topology/ovsdb:1/node/ovsdb:%2F%2FHOST1
```

Note in the above example, that / characters which are part of the *node-id* are specified in hexadecimal format as "%2F".

Passive Connection to OVS Hosts

A passive connection is when the OVS host initiates the connection to the OVSDB Southbound Plugin. This happens when the OVS host is configured to connect to the OVSDB Southbound Plugin. The OVS host is configured with the following command:

```
sudo ovs-vsctl set-manager tcp:<controller-ip>:6640
```

The OVSDB Southbound Plugin is configured to listen for OVSDB connections on TCP port 6640. This value can be changed by editing the "./karaf/target/assembly/etc/custom.properties" file and changing the value of the "ovsdb.listenPort" attribute.

When a passive connection is made, the OVSDB node will appear first in the operational MD-SAL. If the Open_vSwitch table does not contain an external-ids value of *opendaylight-iid*, then the *node-id* of the new OVSDB node will be created in the format:

```
"ovsdb://uuid/<actual UUID value>"
```

If there an *opendaylight-iid* value was already present in the external-ids column, then the instance identifier defined there will be used to create the *node-id* instead.

Query the operational MD-SAL for an OVSDB node after a passive connection.

HTTP GET:

```
http://<controller-ip>:8181/restconf/operational/network-topology:network-topology/
↪topology/ovsdb:1/
```

Result Body:

```
{
  "topology": [
    {
      "topology-id": "ovsdb:1",
      "node": [
        {
          "node-id": "ovsdb://uuid/163724f4-6a70-428a-a8a0-63b2a21f12dd",
          "ovsdb:openvswitch-external-ids": [
            {
              "external-id-key": "system-id",
              "external-id-value": "ecf160af-e78c-4f6b-a005-83a6baa5c979"
            }
          ],
          "ovsdb:connection-info": {
            "local-ip": "<controller-ip>",
            "remote-port": 46731,
            "remote-ip": "<ovs-host-ip>",
            "local-port": 6640
          },
          "ovsdb:ovs-version": "2.3.1-git4750c96",
          "ovsdb:manager-entry": [
            {
              "target": "tcp:10.11.21.7:6640",
              "connected": true,
              "number_of_connections": 1
            }
          ]
        }
      ]
    }
  ]
}
```

Take note of the *node-id* that was created in this case.

Manage Bridges

The OVSDB Southbound Plugin can be used to manage bridges on an OVS host.

This example shows how to add a bridge to the OVSDB node *ovsdb://HOST1*.

HTTP PUT:

```
http://<controller-ip>:8181/restconf/config/network-topology:network-topology/
↪topology/ovsdb:1/node/ovsdb:%2F%2FHOST1%2Fbridge%2Fbrtest
```

Body:

```
{
  "network-topology:node": [
    {
      "node-id": "ovsdb://HOST1/bridge/brtest",
      "ovsdb:bridge-name": "brtest",
      "ovsdb:protocol-entry": [
        {
          "protocol": "ovsdb:ovsdb-bridge-protocol-openflow-13"
        }
      ],
      "ovsdb:managed-by": "/network-topology:network-topology/network-
↪topology:topology[network-topology:topology-id='ovsdb:1']/network-
↪topology:node[network-topology:node-id='ovsdb://HOST1']"
    }
  ]
}
```

Notice that the *ovsdb:managed-by* attribute is specified in the command. This indicates the association of the new bridge node with its OVSDB node.

Bridges can be updated. In the following example, OpenDaylight is configured to be the OpenFlow controller for the bridge.

HTTP PUT:

```
http://<controller-ip>:8181/restconf/config/network-topology:network-topology/
↪topology/ovsdb:1/node/ovsdb:%2F%2FHOST1%2Fbridge%2Fbrtest
```

Body:

```
{
  "network-topology:node": [
        {
          "node-id": "ovsdb://HOST1/bridge/brtest",
             "ovsdb:bridge-name": "brtest",
              "ovsdb:controller-entry": [
                {
                  "target": "tcp:<controller-ip>:6653"
                }
              ],
             "ovsdb:managed-by": "/network-topology:network-topology/network-
↪topology:topology[network-topology:topology-id='ovsdb:1']/network-
↪topology:node[network-topology:node-id='ovsdb://HOST1']"
        }
    ]
}
```

If the OpenDaylight OpenFlow Plugin is installed, then checking on the OVS host will show that OpenDaylight has successfully connected as the controller for the bridge.

```
$ sudo ovs-vsctl show
    Manager "ptcp:6640"
        is_connected: true
    Bridge brtest
        Controller "tcp:<controller-ip>:6653"
            is_connected: true
        Port brtest
            Interface brtest
```

(continues on next page)

(continued from previous page)

```
                type: internal
    ovs_version: "2.3.1-git4750c96"
```

Query the operational MD-SAL to see how the bridge appears.

HTTP GET:

```
http://<controller-ip>:8181/restconf/operational/network-topology:network-topology/
↪topology/ovsdb:1/node/ovsdb:%2F%2FHOST1%2Fbridge%2Fbrtest/
```

Result Body:

```
{
  "node": [
    {
      "node-id": "ovsdb://HOST1/bridge/brtest",
      "ovsdb:bridge-name": "brtest",
      "ovsdb:datapath-type": "ovsdb:datapath-type-system",
      "ovsdb:datapath-id": "00:00:da:e9:0c:08:2d:45",
      "ovsdb:managed-by": "/network-topology:network-topology/network-
↪topology:topology[network-topology:topology-id='ovsdb:1']/network-
↪topology:node[network-topology:node-id='ovsdb://HOST1']",
      "ovsdb:bridge-external-ids": [
        {
          "bridge-external-id-key": "opendaylight-iid",
          "bridge-external-id-value": "/network-topology:network-topology/network-
↪topology:topology[network-topology:topology-id='ovsdb:1']/network-
↪topology:node[network-topology:node-id='ovsdb://HOST1/bridge/brtest']"
        }
      ],
      "ovsdb:protocol-entry": [
        {
          "protocol": "ovsdb:ovsdb-bridge-protocol-openflow-13"
        }
      ],
      "ovsdb:bridge-uuid": "080ce9da-101e-452d-94cd-ee8bef8a4b69",
      "ovsdb:controller-entry": [
        {
          "target": "tcp:10.11.21.7:6653",
          "is-connected": true,
          "controller-uuid": "c39b1262-0876-4613-8bfd-c67eec1a991b"
        }
      ],
      "termination-point": [
        {
          "tp-id": "brtest",
          "ovsdb:port-uuid": "c808ae8d-7af2-4323-83c1-e397696dc9c8",
          "ovsdb:ofport": 65534,
          "ovsdb:interface-type": "ovsdb:interface-type-internal",
          "ovsdb:interface-uuid": "49e9417f-4479-4ede-8faf-7c873b8c0413",
          "ovsdb:name": "brtest"
        }
      ]
    }
  ]
}
```

Notice that just like with the OVSDB node, an *opendaylight-iid* has been added to the external-ids column of the

bridge since it was created via the configuration MD-SAL.

A bridge node may be deleted as well.

HTTP DELETE:

```
http://<controller-ip>:8181/restconf/config/network-topology:network-topology/
↪topology/ovsdb:1/node/ovsdb:%2F%2FHOST1%2Fbridge%2Fbrtest
```

Manage Ports

Similarly, ports may be managed by the OVSDB Southbound Plugin.

This example illustrates how a port and various attributes may be created on a bridge.

HTTP PUT:

```
http://<controller-ip>:8181/restconf/config/network-topology:network-topology/
↪topology/ovsdb:1/node/ovsdb:%2F%2FHOST1%2Fbridge%2Fbrtest/termination-point/
↪testport/
```

Body:

```
{
  "network-topology:termination-point": [
    {
      "ovsdb:options": [
        {
          "ovsdb:option": "remote_ip",
          "ovsdb:value" : "10.10.14.11"
        }
      ],
      "ovsdb:name": "testport",
      "ovsdb:interface-type": "ovsdb:interface-type-vxlan",
      "tp-id": "testport",
      "vlan-tag": "1",
      "trunks": [
        {
          "trunk": "5"
        }
      ],
      "vlan-mode":"access"
    }
  ]
}
```

Ports can be updated - add another VLAN trunk.

HTTP PUT:

```
http://<controller-ip>:8181/restconf/config/network-topology:network-topology/
↪topology/ovsdb:1/node/ovsdb:%2F%2FHOST1%2Fbridge%2Fbrtest/termination-point/
↪testport/
```

Body:

```
{
  "network-topology:termination-point": [
```

(continues on next page)

(continued from previous page)

```
    {
      "ovsdb:name": "testport",
      "tp-id": "testport",
      "trunks": [
        {
          "trunk": "5"
        },
        {
          "trunk": "500"
        }
      ]
    }
  ]
}
```

Query the operational MD-SAL for the port.

HTTP GET:

```
http://<controller-ip>:8181/restconf/operational/network-topology:network-topology/
↪topology/ovsdb:1/node/ovsdb:%2F%2FHOST1%2Fbridge%2Fbrtest/termination-point/
↪testport/
```

Result Body:

```
{
  "termination-point": [
    {
      "tp-id": "testport",
      "ovsdb:port-uuid": "b1262110-2a4f-4442-b0df-84faf145488d",
      "ovsdb:options": [
        {
          "option": "remote_ip",
          "value": "10.10.14.11"
        }
      ],
      "ovsdb:port-external-ids": [
        {
          "external-id-key": "opendaylight-iid",
          "external-id-value": "/network-topology:network-topology/network-
↪topology:topology[network-topology:topology-id='ovsdb:1']/network-
↪topology:node[network-topology:node-id='ovsdb://HOST1/bridge/brtest']/network-
↪topology:termination-point[network-topology:tp-id='testport']"
        }
      ],
      "ovsdb:interface-type": "ovsdb:interface-type-vxlan",
      "ovsdb:trunks": [
        {
          "trunk": 5
        },
        {
          "trunk": 500
        }
      ],
      "ovsdb:vlan-mode": "access",
      "ovsdb:vlan-tag": 1,
      "ovsdb:interface-uuid": "7cec653b-f407-45a8-baec-7eb36b6791c9",
```

(continues on next page)

(continued from previous page)

```
      "ovsdb:name": "testport",
      "ovsdb:ofport": 1
    }
  ]
}
```

Remember that the OVSDB YANG model includes both OVSDB port and interface table attributes in the termination-point augmentation. Both kinds of attributes can be seen in the examples above. Again, note the creation of an *opendaylight-iid* value in the external-ids column of the port table.

Delete a port.

HTTP DELETE:

```
http://<controller-ip>:8181/restconf/config/network-topology:network-topology/
↪topology/ovsdb:1/node/ovsdb:%2F%2FHOST1%2Fbridge%2Fbrtest2/termination-point/
↪testport/
```

Overview of QoS and Queue

The OVSDB Southbound Plugin provides the capability of managing the QoS and Queue tables on an OVS host with OpenDaylight configured as the OVSDB manager.

QoS and Queue Tables in OVSDB

The OVSDB includes a QoS and Queue table. Unlike most of the other tables in the OVSDB, except the Open_vSwitch table, the QoS and Queue tables are "root set" tables, which means that entries, or rows, in these tables are not automatically deleted if they can not be reached directly or indirectly from the Open_vSwitch table. This means that QoS entries can exist and be managed independently of whether or not they are referenced in a Port entry. Similarly, Queue entries can be managed independently of whether or not they are referenced by a QoS entry.

Modelling of QoS and Queue Tables in OpenDaylight MD-SAL

Since the QoS and Queue tables are "root set" tables, they are modeled in the OpenDaylight MD-SAL as lists which are part of the attributes of the OVSDB node model.

The MD-SAL QoS and Queue models have an additonal identifier attribute per entry (e.g. "qos-id" or "queue-id") which is not present in the OVSDB schema. This identifier is used by the MD-SAL as a key for referencing the entry. If the entry is created originally from the configuration MD-SAL, then the value of the identifier is whatever is specified by the configuration. If the entry is created on the OVSDB node and received by OpenDaylight in an operational update, then the id will be created in the following format.

```
"queue-id": "queue://<UUID>"
"qos-id": "qos://<UUID>"
```

The UUID in the above identifiers is the actual UUID of the entry in the OVSDB database.

When the QoS or Queue entry is created by the configuration MD-SAL, the identifier will be configured as part of the external-ids column of the entry. This will ensure that the corresponding entry that is created in the operational MD-SAL uses the same identifier.

```
"queues-external-ids": [
  {
    "queues-external-id-key": "opendaylight-queue-id",
    "queues-external-id-value": "QUEUE-1"
  }
]
```

See more in the examples that follow in this section.

The QoS schema in OVSDB currently defines two types of QoS entries.

- linux-htb
- linux-hfsc

These QoS types are defined in the QoS model. Additional types will need to be added to the model in order to be supported. See the examples that folow for how the QoS type is specified in the model.

QoS entries can be configured with addtional attritubes such as "max-rate". These are configured via the *other-config* column of the QoS entry. Refer to OVSDB schema (in the reference section below) for all of the relevant attributes that can be configured. The examples in the rest of this section will demonstrate how the other-config column may be configured.

Similarly, the Queue entries may be configured with additional attributes via the other-config column.

Managing QoS and Queues via Configuration MD-SAL

This section will show some examples on how to manage QoS and Queue entries via the configuration MD-SAL. The examples will be illustrated by using RESTCONF (see QoS and Queue Postman Collection).

A pre-requisite for managing QoS and Queue entries is that the OVS host must be present in the configuration MD-SAL.

For the following examples, the following OVS host is configured.

HTTP POST:

```
http://<controller-ip>:8181/restconf/config/network-topology:network-topology/
↪topology/ovsdb:1/
```

Body:

```
{
  "node": [
    {
      "node-id": "ovsdb:HOST1",
      "connection-info": {
        "ovsdb:remote-ip": "<ovs-host-ip>",
        "ovsdb:remote-port": "<ovs-host-ovsdb-port>"
      }
    }
  ]
}
```

Where

- <*controller-ip*> is the IP address of the OpenDaylight controller
- <*ovs-host-ip*> is the IP address of the OVS host

- *<ovs-host-ovsdb-port>* is the TCP port of the OVSDB server on the OVS host (e.g. 6640)

This command creates an OVSDB node with the node-id "ovsdb:HOST1". This OVSDB node will be used in the following examples.

QoS and Queue entries can be created and managed without a port, but ultimately, QoS entries are associated with a port in order to use them. For the following examples a test bridge and port will be created.

Create the test bridge.

HTTP PUT

```
http://<controller-ip>:8181/restconf/config/network-topology:network-topology/
↪topology/ovsdb:1/node/ovsdb:HOST1%2Fbridge%2Fbr-test
```

Body:

```
{
  "network-topology:node": [
    {
      "node-id": "ovsdb:HOST1/bridge/br-test",
      "ovsdb:bridge-name": "br-test",
      "ovsdb:managed-by": "/network-topology:network-topology/network-
↪topology:topology[network-topology:topology-id='ovsdb:1']/network-
↪topology:node[network-topology:node-id='ovsdb:HOST1']"
    }
  ]
}
```

Create the test port (which is modeled as a termination point in the OpenDaylight MD-SAL).

HTTP PUT:

```
http://<controller-ip>:8181/restconf/config/network-topology:network-topology/
↪topology/ovsdb:1/node/ovsdb:HOST1%2Fbridge%2Fbr-test/termination-point/testport/
```

Body:

```
{
  "network-topology:termination-point": [
    {
      "ovsdb:name": "testport",
      "tp-id": "testport"
    }
  ]
}
```

If all of the previous steps were successful, a query of the operational MD-SAL should look something like the following results. This indicates that the configuration commands have been successfully instantiated on the OVS host.

HTTP GET:

```
http://<controller-ip>:8181/restconf/operational/network-topology:network-topology/
↪topology/ovsdb:1/node/ovsdb:HOST1%2Fbridge%2Fbr-test
```

Result Body:

```
{
  "node": [
    {
      "node-id": "ovsdb:HOST1/bridge/br-test",
      "ovsdb:bridge-name": "br-test",
      "ovsdb:datapath-type": "ovsdb:datapath-type-system",
      "ovsdb:managed-by": "/network-topology:network-topology/network-
↪topology:topology[network-topology:topology-id='ovsdb:1']/network-
↪topology:node[network-topology:node-id='ovsdb:HOST1']",
      "ovsdb:datapath-id": "00:00:8e:5d:22:3d:09:49",
      "ovsdb:bridge-external-ids": [
        {
          "bridge-external-id-key": "opendaylight-iid",
          "bridge-external-id-value": "/network-topology:network-topology/network-
↪topology:topology[network-topology:topology-id='ovsdb:1']/network-
↪topology:node[network-topology:node-id='ovsdb:HOST1/bridge/br-test']"
        }
      ],
      "ovsdb:bridge-uuid": "3d225d8d-d060-4909-93ef-6f4db58ef7cc",
      "termination-point": [
        {
          "tp-id": "br=-est",
          "ovsdb:port-uuid": "f85f7aa7-4956-40e4-9c94-e6ca2d5cd254",
          "ovsdb:ofport": 65534,
          "ovsdb:interface-type": "ovsdb:interface-type-internal",
          "ovsdb:interface-uuid": "29ff3692-6ed4-4ad7-a077-1edc277ecb1a",
          "ovsdb:name": "br-test"
        },
        {
          "tp-id": "testport",
          "ovsdb:port-uuid": "aa79a8e2-147f-403a-9fa9-6ee5ec276f08",
          "ovsdb:port-external-ids": [
            {
              "external-id-key": "opendaylight-iid",
              "external-id-value": "/network-topology:network-topology/network-
↪topology:topology[network-topology:topology-id='ovsdb:1']/network-
↪topology:node[network-topology:node-id='ovsdb:HOST1/bridge/br-test']/network-
↪topology:termination-point[network-topology:tp-id='testport']"
            }
          ],
          "ovsdb:interface-uuid": "e96f282e-882c-41dd-a870-80e6b29136ac",
          "ovsdb:name": "testport"
        }
      ]
    }
  ]
}
```

Create Queue

Create a new Queue in the configuration MD-SAL.

HTTP PUT:

```
http://<controller-ip>:8181/restconf/config/network-topology:network-topology/
↪topology/ovsdb:1/node/ovsdb:HOST1/ovsdb:queues/QUEUE-1/
```

Body:

```
{
  "ovsdb:queues": [
    {
      "queue-id": "QUEUE-1",
      "dscp": 25,
      "queues-other-config": [
        {
          "queue-other-config-key": "max-rate",
          "queue-other-config-value": "3600000"
        }
      ]
    }
  ]
}
```

Query Queue

Now query the operational MD-SAL for the Queue entry.

HTTP GET:

```
http://<controller-ip>:8181/restconf/operational/network-topology:network-topology/
↪topology/ovsdb:1/node/ovsdb:HOST1/ovsdb:queues/QUEUE-1/
```

Result Body:

```
{
  "ovsdb:queues": [
    {
      "queue-id": "QUEUE-1",
      "queues-other-config": [
        {
          "queue-other-config-key": "max-rate",
          "queue-other-config-value": "3600000"
        }
      ],
      "queues-external-ids": [
        {
          "queues-external-id-key": "opendaylight-queue-id",
          "queues-external-id-value": "QUEUE-1"
        }
      ],
      "queue-uuid": "83640357-3596-4877-9527-b472aa854d69",
      "dscp": 25
    }
  ]
}
```

Create QoS

Create a QoS entry. Note that the UUID of the Queue entry, obtained by querying the operational MD-SAL of the Queue entry, is specified in the queue-list of the QoS entry. Queue entries may be added to the QoS entry at the creation of the QoS entry, or by a subsequent update to the QoS entry.

HTTP PUT:

```
http://<controller-ip>:8181/restconf/config/network-topology:network-topology/
↪topology/ovsdb:1/node/ovsdb:HOST1/ovsdb:qos-entries/QOS-1/
```

Body:

```
{
  "ovsdb:qos-entries": [
    {
      "qos-id": "QOS-1",
      "qos-type": "ovsdb:qos-type-linux-htb",
      "qos-other-config": [
        {
          "other-config-key": "max-rate",
          "other-config-value": "4400000"
        }
      ],
      "queue-list": [
        {
          "queue-number": "0",
          "queue-uuid": "83640357-3596-4877-9527-b472aa854d69"
        }
      ]
    }
  ]
}
```

Query QoS

Query the operational MD-SAL for the QoS entry.

HTTP GET:

```
http://<controller-ip>:8181/restconf/operational/network-topology:network-topology/
↪topology/ovsdb:1/node/ovsdb:HOST1/ovsdb:qos-entries/QOS-1/
```

Result Body:

```
{
  "ovsdb:qos-entries": [
    {
      "qos-id": "QOS-1",
      "qos-other-config": [
        {
          "other-config-key": "max-rate",
          "other-config-value": "4400000"
        }
      ],
      "queue-list": [
        {
          "queue-number": 0,
          "queue-uuid": "83640357-3596-4877-9527-b472aa854d69"
        }
      ],
      "qos-type": "ovsdb:qos-type-linux-htb",
```

(continues on next page)

(continued from previous page)

```
      "qos-external-ids": [
        {
          "qos-external-id-key": "opendaylight-qos-id",
          "qos-external-id-value": "QOS-1"
        }
      ],
      "qos-uuid": "90ba9c60-3aac-499d-9be7-555f19a6bb31"
    }
  ]
}
```

Add QoS to a Port

Update the termination point entry to include the UUID of the QoS entry, obtained by querying the operational MD-SAL, to associate a QoS entry with a port.

HTTP PUT:

```
http://<controller-ip>:8181/restconf/config/network-topology:network-topology/
↪topology/ovsdb:1/node/ovsdb:HOST1%2Fbridge%2Fbr-test/termination-point/testport/
```

Body:

```
{
  "network-topology:termination-point": [
    {
      "ovsdb:name": "testport",
      "tp-id": "testport",
      "qos": "90ba9c60-3aac-499d-9be7-555f19a6bb31"
    }
  ]
}
```

Query the Port

Query the operational MD-SAL to see how the QoS entry appears in the termination point model.

HTTP GET:

```
http://<controller-ip>:8181/restconf/operational/network-topology:network-topology/
↪topology/ovsdb:1/node/ovsdb:HOST1%2Fbridge%2Fbr-test/termination-point/testport/
```

Result Body:

```
{
  "termination-point": [
    {
      "tp-id": "testport",
      "ovsdb:port-uuid": "aa79a8e2-147f-403a-9fa9-6ee5ec276f08",
      "ovsdb:port-external-ids": [
        {
          "external-id-key": "opendaylight-iid",
          "external-id-value": "/network-topology:network-topology/network-
↪topology:topology[network-topology:topology-id='ovsdb:1']/network-
↪topology:node[network-topology:node-id='ovsdb:HOST1/bridge/br-test']/network-
↪topology:termination-point[network-topology:tp-id='testport']"
```

(continues on next page)

(continued from previous page)

```
        }
      ],
      "ovsdb:qos": "90ba9c60-3aac-499d-9be7-555f19a6bb31",
      "ovsdb:interface-uuid": "e96f282e-882c-41dd-a870-80e6b29136ac",
      "ovsdb:name": "testport"
    }
  ]
}
```

Query the OVSDB Node

Query the operational MD-SAL for the OVS host to see how the QoS and Queue entries appear as lists in the OVS node model.

HTTP GET:

```
http://<controller-ip>:8181/restconf/operational/network-topology:network-topology/
↪topology/ovsdb:1/node/ovsdb:HOST1/
```

Result Body (edited to only show information relevant to the QoS and Queue entries):

```
{
  "node": [
    {
      "node-id": "ovsdb:HOST1",
      <content edited out>
      "ovsdb:queues": [
        {
          "queue-id": "QUEUE-1",
          "queues-other-config": [
            {
              "queue-other-config-key": "max-rate",
              "queue-other-config-value": "3600000"
            }
          ],
          "queues-external-ids": [
            {
              "queues-external-id-key": "opendaylight-queue-id",
              "queues-external-id-value": "QUEUE-1"
            }
          ],
          "queue-uuid": "83640357-3596-4877-9527-b472aa854d69",
          "dscp": 25
        }
      ],
      "ovsdb:qos-entries": [
        {
          "qos-id": "QOS-1",
          "qos-other-config": [
            {
              "other-config-key": "max-rate",
              "other-config-value": "4400000"
            }
          ],
          "queue-list": [
```

(continues on next page)

(continued from previous page)

```
            {
              "queue-number": 0,
              "queue-uuid": "83640357-3596-4877-9527-b472aa854d69"
            }
          ],
          "qos-type": "ovsdb:qos-type-linux-htb",
          "qos-external-ids": [
            {
              "qos-external-id-key": "opendaylight-qos-id",
              "qos-external-id-value": "QOS-1"
            }
          ],
          "qos-uuid": "90ba9c60-3aac-499d-9be7-555f19a6bb31"
        }
      ]
      <content edited out>
    }
  ]
}
```

Remove QoS from a Port

This example removes a QoS entry from the termination point and associated port. Note that this is a PUT command on the termination point with the QoS attribute absent. Other attributes of the termination point should be included in the body of the command so that they are not inadvertantly removed.

HTTP PUT:

```
http://<controller-ip>:8181/restconf/config/network-topology:network-topology/
↪topology/ovsdb:1/node/ovsdb:HOST1%2Fbridge%2Fbr-test/termination-point/testport/
```

Body:

```
{
  "network-topology:termination-point": [
    {
      "ovsdb:name": "testport",
      "tp-id": "testport"
    }
  ]
}
```

Remove a Queue from QoS

This example removes the specific Queue entry from the queue list in the QoS entry. The queue entry is specified by the queue number, which is "0" in this example.

HTTP DELETE:

```
http://<controller-ip>:8181/restconf/config/network-topology:network-topology/
↪topology/ovsdb:1/node/ovsdb:HOST1/ovsdb:qos-entries/QOS-1/queue-list/0/
```

Remove Queue

Once all references to a specific queue entry have been removed from QoS entries, the Queue itself can be removed.

HTTP DELETE:

```
http://<controller-ip>:8181/restconf/config/network-topology:network-topology/
↪topology/ovsdb:1/node/ovsdb:HOST1/ovsdb:queues/QUEUE-1/
```

Remove QoS

The QoS entry may be removed when it is no longer referenced by any ports.

HTTP DELETE:

```
http://<controller-ip>:8181/restconf/config/network-topology:network-topology/
↪topology/ovsdb:1/node/ovsdb:HOST1/ovsdb:qos-entries/QOS-1/
```

References

Openvswitch schema

OVSDB and Netvirt Postman Collection

OVSDB Hardware VTEP SouthBound Plugin

Overview

Hwvtepsouthbound plugin is used to configure a hardware VTEP which implements hardware ovsdb schema. This page will show how to use RESTConf API of hwvtepsouthbound. There are two ways to connect to ODL:

switch initiates connection..

Both will be introduced respectively.

User Initiates Connection

Prerequisite

Configure the hwvtep device/node to listen for the tcp connection in passive mode. In addition, management IP and tunnel source IP are also configured. After all this configuration is done, a physical switch is created automatically by the hwvtep node.

Connect to a hwvtep device/node

Send below Restconf request if you want to initiate the connection to a hwvtep node from the controller, where listening IP and port of hwvtep device/node are provided.

REST API: POST http://odl:8181/restconf/config/network-topology:network-topology/topology/hwvtep:1/

```
{
 "network-topology:node": [
       {
           "node-id": "hwvtep://192.168.1.115:6640",
           "hwvtep:connection-info":
           {
               "hwvtep:remote-port": 6640,
               "hwvtep:remote-ip": "192.168.1.115"
           }
       }
   ]
}
```

Please replace *odl* in the URL with the IP address of your OpenDaylight controller and change *192.168.1.115* to your hwvtep node IP.

NOTE: The format of node-id is fixed. It will be one of the two:

User initiates connection from ODL:

```
hwvtep://ip:port
```

Switch initiates connection:

```
hwvtep://uuid/<uuid of switch>
```

The reason for using UUID is that we can distinguish between multiple switches if they are behind a NAT.

After this request is completed successfully, we can get the physical switch from the operational data store.

REST API: GET http://odl:8181/restconf/operational/network-topology:network-topology/topology/hwvtep:1/node/hwvtep:%2F%2F192.168.1.115:6640

There is no body in this request.

The response of the request is:

```
{
   "node": [
         {
           "node-id": "hwvtep://192.168.1.115:6640",
           "hwvtep:switches": [
             {
               "switch-ref": "/network-topology:network-topology/network-
↪topology:topology[network-topology:topology-id='hwvtep:1']/network-
↪topology:node[network-topology:node-id='hwvtep://192.168.1.115:6640/physicalswitch/
↪br0']"
             }
           ],
           "hwvtep:connection-info": {
             "local-ip": "192.168.92.145",
             "local-port": 47802,
             "remote-port": 6640,
             "remote-ip": "192.168.1.115"
           }
         },
         {
           "node-id": "hwvtep://192.168.1.115:6640/physicalswitch/br0",
           "hwvtep:management-ips": [
```

(continues on next page)

(continued from previous page)

```
              {
                "management-ips-key": "192.168.1.115"
              }
            ],
            "hwvtep:physical-switch-uuid": "37eb5abd-a6a3-4aba-9952-a4d301bdf371",
            "hwvtep:managed-by": "/network-topology:network-topology/network-
↪topology:topology[network-topology:topology-id='hwvtep:1']/network-
↪topology:node[network-topology:node-id='hwvtep://192.168.1.115:6640']",
            "hwvtep:hwvtep-node-description": "",
            "hwvtep:tunnel-ips": [
              {
                "tunnel-ips-key": "192.168.1.115"
              }
            ],
            "hwvtep:hwvtep-node-name": "br0"
          }
        ]
}
```

If there is a physical switch which has already been created by manual configuration, we can get the node-id of the physical switch from this response, which is presented in "swith-ref". If the switch does not exist, we need to create the physical switch. Currently, most hwvtep devices do not support running multiple switches.

Create a physical switch

REST API: POST http://odl:8181/restconf/config/network-topology:network-topology/topology/hwvtep:1/

request body:

```
{
 "network-topology:node": [
        {
            "node-id": "hwvtep://192.168.1.115:6640/physicalswitch/br0",
            "hwvtep-node-name": "ps0",
            "hwvtep-node-description": "",
            "management-ips": [
              {
                "management-ips-key": "192.168.1.115"
              }
            ],
            "tunnel-ips": [
              {
                "tunnel-ips-key": "192.168.1.115"
              }
            ],
            "managed-by": "/network-topology:network-topology/network-
↪topology:topology[network-topology:topology-id='hwvtep:1']/network-
↪topology:node[network-topology:node-id='hwvtep://192.168.1.115:6640']"
        }
    ]
}
```

Note: "managed-by" must provided by user. We can get its value after the step *Connect to a hwvtep device/node* since the node-id of hwvtep device is provided by user. "managed-by" is a reference typed of instance identifier. Though the instance identifier is a little complicated for RestConf, the primary user of hwvtepsouthbound plugin will be provider-type code such as NetVirt and the instance identifier is much easier to write code for.

Create a logical switch

Creating a logical switch is effectively creating a logical network. For VxLAN, it is a tunnel network with the same VNI.

REST API: POST http://odl:8181/restconf/config/network-topology:network-topology/topology/hwvtep:1/node/hwvtep:%2F%2F192.168.1.115:6640

request body:

```
{
 "logical-switches": [
       {
           "hwvtep-node-name": "ls0",
           "hwvtep-node-description": "",
           "tunnel-key": "10000"
        }
   ]
}
```

Create a physical locator

After the VXLAN network is ready, we will add VTEPs to it. A VTEP is described by a physical locator.

REST API: POST http://odl:8181/restconf/config/network-topology:network-topology/topology/hwvtep:1/node/hwvtep:%2F%2F192.168.1.115:6640

request body:

```
{
 "termination-point": [
      {
          "tp-id": "vxlan_over_ipv4:192.168.0.116",
          "encapsulation-type": "encapsulation-type-vxlan-over-ipv4",
          "dst-ip": "192.168.0.116"
          }
     ]
}
```

The "tp-id" of locator is "{encapsualation-type}: {dst-ip}".

Note: As far as we know, the OVSDB database does not allow the insertion of a new locator alone. So, no locator is inserted after this request is sent. We will trigger off the creation until other entity refer to it, such as remote-mcast-macs.

Create a remote-mcast-macs entry

After adding a physical locator to a logical switch, we need to create a remote-mcast-macs entry to handle unknown traffic.

REST API: POST http://odl:8181/restconf/config/network-topology:network-topology/topology/hwvtep:1/node/hwvtep:%2F%2F192.168.1.115:6640

request body:

```
{
 "remote-mcast-macs": [
       {
           "mac-entry-key": "00:00:00:00:00:00",
           "logical-switch-ref": "/network-topology:network-topology/network-
↪topology:topology[network-topology:topology-id='hwvtep:1']/network-
↪topology:node[network-topology:node-id='hwvtep://192.168.1.115:6640']/
↪hwvtep:logical-switches[hwvtep:hwvtep-node-name='ls0']",
           "locator-set": [
                {
                      "locator-ref": "/network-topology:network-topology/network-
↪topology:topology[network-topology:topology-id='hwvtep:1']/network-
↪topology:node[network-topology:node-id='hwvtep://219.141.189.115:6640']/network-
↪topology:termination-point[network-topology:tp-id='vxlan_over_ipv4:192.168.0.116']"
                }
           ]
       }
   ]
}
```

The physical locator *vxlan_over_ipv4:192.168.0.116* is just created in "Create a physical locator". It should be noted that list "locator-set" is immutable, that is, we must provide a set of "locator-ref" as a whole.

Note: "00:00:00:00:00:00" stands for "unknown-dst" since the type of mac-entry-key is yang:mac and does not accept "unknown-dst".

Create a physical port

Now we add a physical port into the physical switch "hwvtep://192.168.1.115:6640/physicalswitch/br0". The port is attached with a physical server or an L2 network and with the vlan 100.

REST API: POST http://odl:8181/restconf/config/network-topology:network-topology/topology/hwvtep:1/node/hwvtep:%2F%2F192.168.1.115:6640%2Fphysicalswitch%2Fbr0

```
{
 "network-topology:termination-point": [
       {
           "tp-id": "port0",
           "hwvtep-node-name": "port0",
           "hwvtep-node-description": "",
           "vlan-bindings": [
               {
                 "vlan-id-key": "100",
                 "logical-switch-ref": "/network-topology:network-topology/network-
↪topology:topology[network-topology:topology-id='hwvtep:1']/network-
↪topology:node[network-topology:node-id='hwvtep://192.168.1.115:6640']/
↪hwvtep:logical-switches[hwvtep:hwvtep-node-name='ls0']"
               }
         ]
       }
   ]
}
```

At this point, we have completed the basic configuration.

Typically, hwvtep devices learn local MAC addresses automatically. But they also support getting MAC address entries from ODL.

Create a local-mcast-macs entry

It is similar to *Create a remote-mcast-macs entry*.

Create a remote-ucast-macs

REST API: POST http://odl:8181/restconf/config/network-topology:network-topology/topology/hwvtep:1/node/hwvtep:%2F%2F192.168.1.115:6640

```
request body:
```

```
{
 "remote-ucast-macs": [
       {
           "mac-entry-key": "11:11:11:11:11:11",
           "logical-switch-ref": "/network-topology:network-topology/network-
↪topology:topology[network-topology:topology-id='hwvtep:1']/network-
↪topology:node[network-topology:node-id='hwvtep://192.168.1.115:6640']/
↪hwvtep:logical-switches[hwvtep:hwvtep-node-name='ls0']",
           "ipaddr": "1.1.1.1",
           "locator-ref": "/network-topology:network-topology/network-
↪topology:topology[network-topology:topology-id='hwvtep:1']/network-
↪topology:node[network-topology:node-id='hwvtep://192.168.1.115:6640']/network-
↪topology:termination-point[network-topology:tp-id='vxlan_over_ipv4:192.168.0.116']"
       }
   ]
}
```

Create a local-ucast-macs entry

This is similar to *Create a remote-ucast-macs*.

Switch Initiates Connection

We do not need to connect to a hwvtep device/node when the switch initiates the connection. After switches connect to ODL successfully, we get the node-id's of switches by reading the operational data store. Once the node-id of a hwvtep device is received, the remaining steps are the same as when the user initiates the connection.

References

https://wiki.opendaylight.org/view/User_talk:Pzhang

PCEP User Guide

This guide contains information on how to use the OpenDaylight Path Computation Element Configuration Protocol (PCEP) plugin. The user should learn about PCEP basic concepts, supported capabilities, configuration and operations.

Contents

- *Overview*
- *Running PCEP*
- *Active Stateful PCE*
- *Test tools*
- *Troubleshooting*
- *References*

Overview

This section provides a high-level overview of the PCEP, SDN use-cases and OpenDaylight implementation.

Contents

- *Path Computation Element Communication Protocol*
- *PCEP in SDN*
- *OpenDaylight PCEP plugin*
  - *List of supported capabilities*

Path Computation Element Communication Protocol

The Path Computation Element (PCE) Communication Protocol (PCEP) is used for communication between a Path Computation Client (PCC) and a PCE in context of MPLS and GMPLS Traffic Engineering (TE) Label Switched Paths (LSPs). This interaction include path computation requests and computation replies. The PCE operates on a network graph, built from the (Traffic Engineering Database) TED, in order to compute paths based on the path computation request issued by the PCC. The path computation request includes the source and destination of the path and set of constrains to be applied during the computation. The PCE response contains the computed path or the computation failure reason. The PCEP operates on top the TCP, which provides reliable communication.

PCEP in SDN

The Path Computation Element perfectly fits into the centralized SDN controller architecture. The PCE's knowledge of the availability of network resources (i.e. TED) and active LSPs awareness (LSP-DB) allows to perform automated application-driven network operations:

- LSP Re-optimization
- Resource defragmentation
- Link failure restoration
- Auto-bandwidth adjustment
- Bandwidth scheduling
- Shared Risk Link Group (SRLG) diversity maintenance

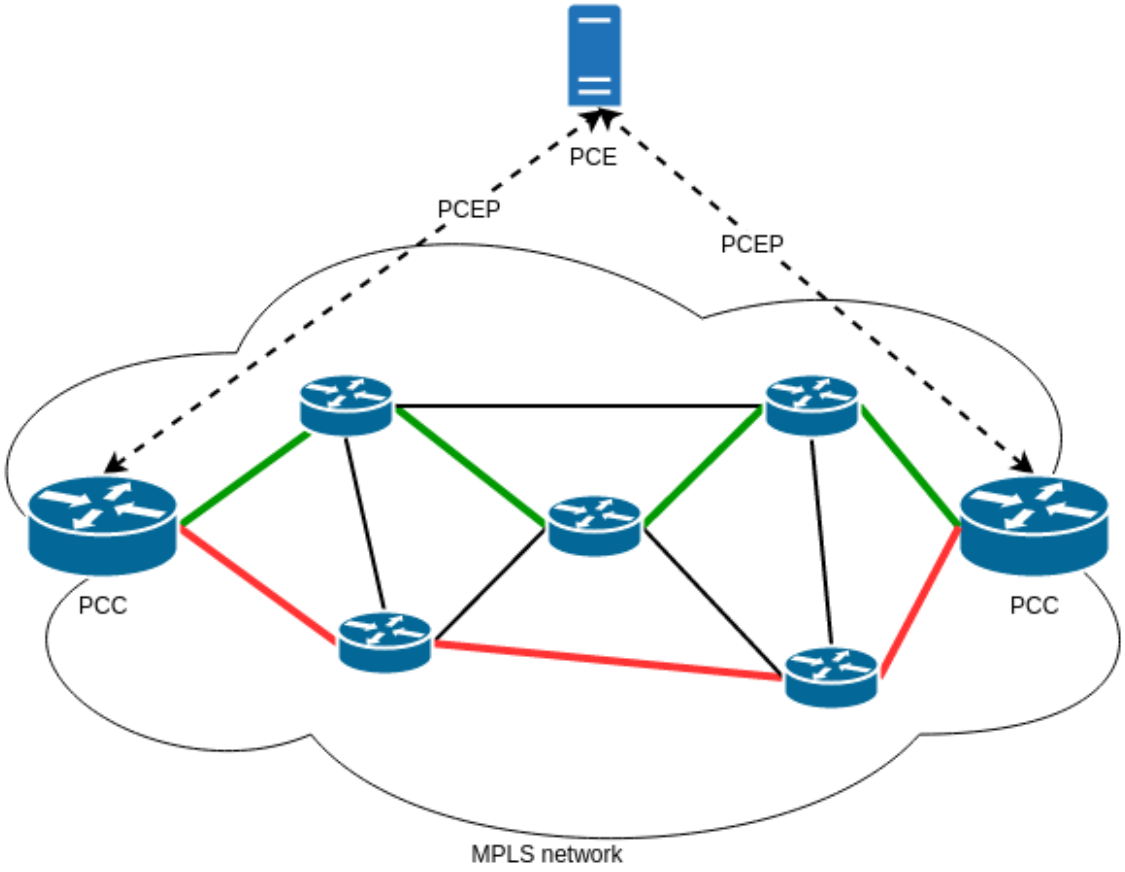

Fig. 80: PCE-based architecture.

OpenDaylight PCEP plugin

The OpenDaylight PCEP plugin provides all basic service units necessary to build-up a PCE-based controller. In addition, it offers LSP management functionality for Active Stateful PCE - the cornerstone for majority of PCE-enabled SDN solutions. It consists of the following components:

- Protocol library
- PCEP session handling
- Stateful PCE LSP-DB
- Active Stateful PCE LSP Operations

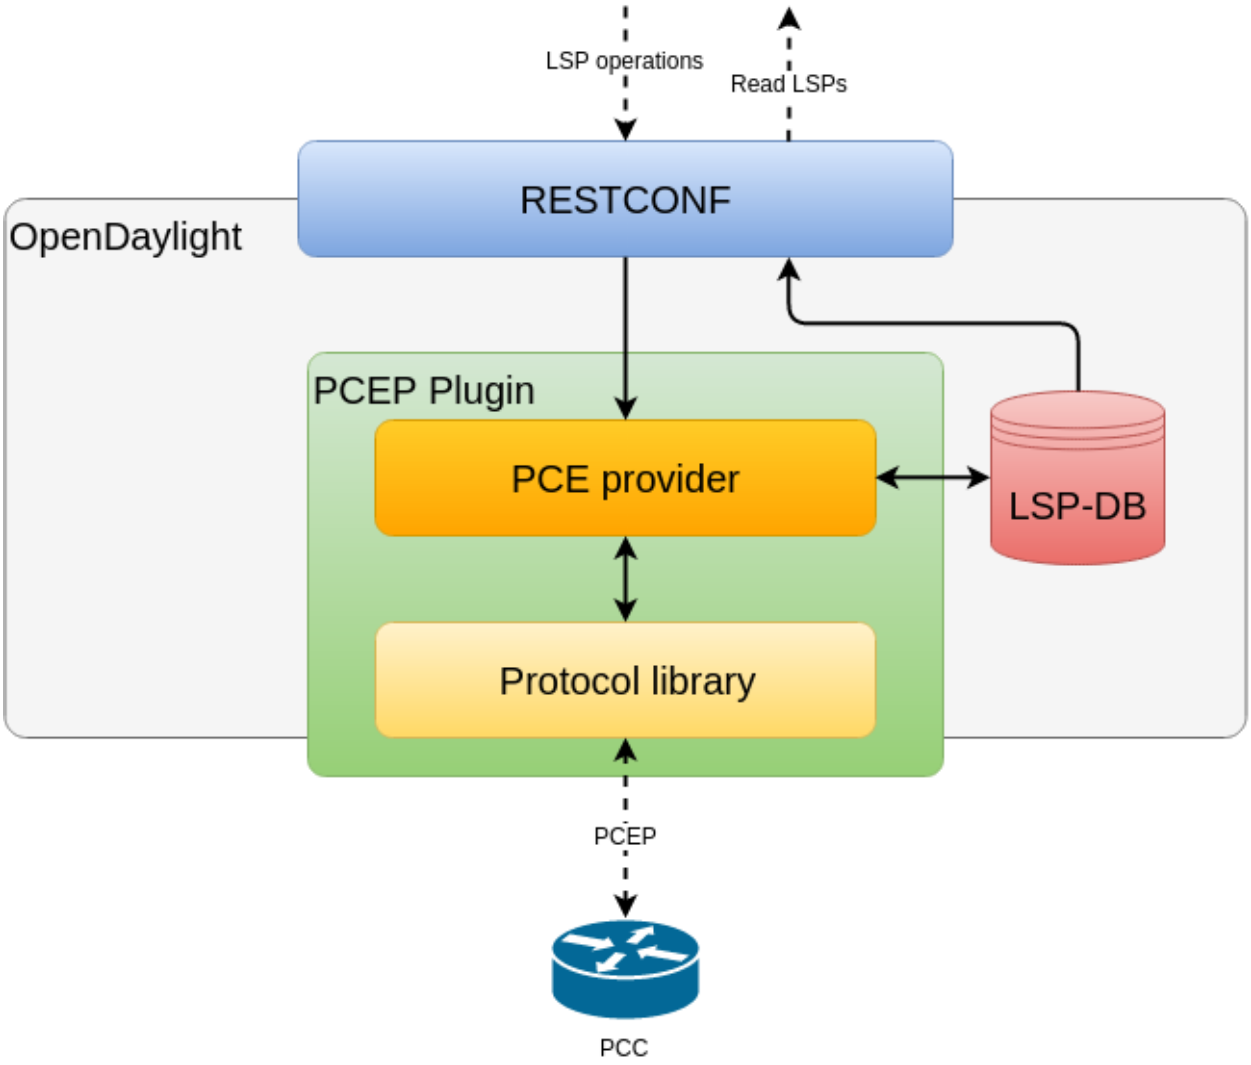

Fig. 81: OpenDaylight PCEP plugin overview.

Important: The PCEP plugin does not provide path computational functionality and does not build TED.

List of supported capabilities

- RFC5440 - Path Computation Element (PCE) Communication Protocol (PCEP)
- RFC5455 - Diffserv-Aware Class-Type Object for the Path Computation Element Communication Protocol
- RFC5520 - Preserving Topology Confidentiality in Inter-Domain Path Computation Using a Path-Key-Based Mechanism
- RFC5521 - Extensions to the Path Computation Element Communication Protocol (PCEP) for Route Exclusions
- RFC5541 - Encoding of Objective Functions in the Path Computation Element Communication Protocol (PCEP)
- RFC5557 - Path Computation Element Communication Protocol (PCEP) Requirements and Protocol Extensions in Support of Global Concurrent Optimization
- RFC5886 - A Set of Monitoring Tools for Path Computation Element (PCE)-Based Architecture
- RFC7470 - Conveying Vendor-Specific Constraints in the Path Computation Element Communication Protocol
- RFC7896 - Update to the Include Route Object (IRO) Specification in the Path Computation Element Communication Protocol (PCEP)
- draft-ietf-pce-stateful-pce - PCEP Extensions for Stateful PCE
  - draft-ietf-pce-pce-initiated-lsp - PCEP Extensions for PCE-initiated LSP Setup in a Stateful PCE Model
  - draft-ietf-pce-segment-routing - PCEP Extension for segment routing
  - draft-ietf-pce-lsp-setup-type - PCEP Extension for path setup type
  - draft-ietf-pce-stateful-sync-optimizations - Optimizations of Label Switched Path State Synchronization Procedures for a Stateful PCE
  - draft-sivabalan-pce-binding-label-sid - Carrying Binding Label/Segment-ID in PCE-based Networks
- draft-ietf-pce-pceps - Secure Transport for PCEP

Running PCEP

This section explains how to install PCEP plugin.

1. Install PCEP feature - `odl-bgpcep-pcep`. Also, for sake of this sample, it is required to install RESTCONF. In the Karaf console, type command:

```
feature:install odl-restconf odl-bgpcep-pcep
```

2. The PCEP plugin contains a default configuration, which is applied after the feature starts up. One instance of PCEP plugin is created (named *pcep-topology*), and its presence can be verified via REST:

   URL: `restconf/operational/network-topology:network-topology/topology/pcep-topology`

   Method: `GET`

   Response Body:

```
<topology xmlns="urn:TBD:params:xml:ns:yang:network-topology">
    <topology-id>pcep-topology</topology-id>
    <topology-types>
        <topology-pcep xmlns="urn:opendaylight:params:xml:ns:yang:topology:pcep">
↪</topology-pcep>
```

(continues on next page)

(continued from previous page)

```
    </topology-types>
</topology>
```

Active Stateful PCE

The PCEP extension for Stateful PCE brings a visibility of active LSPs to PCE, in order to optimize path computation, while considering individual LSPs and their interactions. This requires state synchronization mechanism between PCE and PCC. Moreover, Active Stateful PCE is capable to address LSP parameter changes to the PCC.

Contents

- *Configuration*
  - *MD5 authentication configuration*
- *LSP State Database*
  - *LSP-DB API*
  - *LSP Delegation*
  - *LSP Update*
- *PCE-initiated LSP Setup*
  - *Configuration*
  - *LSP Instantiation*
  - *LSP Deletion*
  - *PCE-initiated LSP Delegation*
- *Segment Routing*
  - *Configuration*
  - *LSP Operations for PCEP SR*
- *LSP State Synchronization Optimization Procedures*
  - *Configuration*
  - *State Synchronization Avoidance*
  - *Incremental State Synchronization*
  - *PCE-triggered Initial Synchronization*
  - *PCE-triggered Re-synchronization*

Configuration

This capability is enabled by default. No additional configuration is required.

MD5 authentication configuration

The OpenDaylight PCEP implementation is supporting TCP MD5 for authentication. Sample configuration below shows how to set authentication password for a particular PCC. It is required to install `odl-netconf-connector-ssh` feature first.

URL: `/restconf/config/network-topology:network-topology/topology/topology-netconf/node/controller-config/yang-ext:mount/config:modules/module/odl-pcep-topology-provider-cfg:pcep-topology-provider/pcep-topology`

Method: `PUT`

Content-Type: `application/xml`

Request Body:

```
<module xmlns="urn:opendaylight:params:xml:ns:yang:controller:config">
    <type xmlns:x="urn:opendaylight:params:xml:ns:yang:controller:pcep:topology:provider">x:pcep-topology-provider</type>
    <name>pcep-topology</name>
    <data-provider xmlns="urn:opendaylight:params:xml:ns:yang:controller:pcep:topology:provider">
        <type xmlns:x="urn:opendaylight:params:xml:ns:yang:controller:md:sal:binding">x:binding-async-data-broker</type>
        <name>pingpong-binding-data-broker</name>
    </data-provider>
    <dispatcher xmlns="urn:opendaylight:params:xml:ns:yang:controller:pcep:topology:provider">
        <type xmlns:x="urn:opendaylight:params:xml:ns:yang:controller:pcep">x:pcep-dispatcher</type>
        <name>global-pcep-dispatcher</name>
    </dispatcher>
    <rpc-registry xmlns="urn:opendaylight:params:xml:ns:yang:controller:pcep:topology:provider">
        <type xmlns:x="urn:opendaylight:params:xml:ns:yang:controller:md:sal:binding">x:binding-rpc-registry</type>
        <name>binding-rpc-broker</name>
    </rpc-registry>
    <scheduler xmlns="urn:opendaylight:params:xml:ns:yang:controller:pcep:topology:provider">
        <type xmlns:x="urn:opendaylight:params:xml:ns:yang:controller:programming:spi">x:instruction-scheduler</type>
        <name>global-instruction-scheduler</name>
    </scheduler>
    <stateful-plugin xmlns="urn:opendaylight:params:xml:ns:yang:controller:pcep:topology:provider">
        <type>pcep-topology-stateful</type>
        <name>stateful07</name>
    </stateful-plugin>
    <topology-id xmlns="urn:opendaylight:params:xml:ns:yang:controller:pcep:topology:provider">pcep-topology</topology-id>
    <client xmlns="urn:opendaylight:params:xml:ns:yang:controller:pcep:topology:provider">
        <address>43.43.43.43</address>
        <password>topsecret</password>
    </client>
</module>
```

@line 26: **address** - A PCC IP address.

@line 27: **password** - MD5 authentication phrase.

Warning: The PCE (*pcep-topology-provider*) configuration is going to be changed in Carbon release - moving to configuration datastore.

LSP State Database

The *LSP State Database* (LSP-DB) contains an information about all LSPs and their attributes. The LSP state is synchronized between the PCC and PCE. First, initial LSP state synchronization is performed once the session between PCC and PCE is established in order to learn PCC's LPSs. This step is a prerequisite to following LSPs manipulation operations.

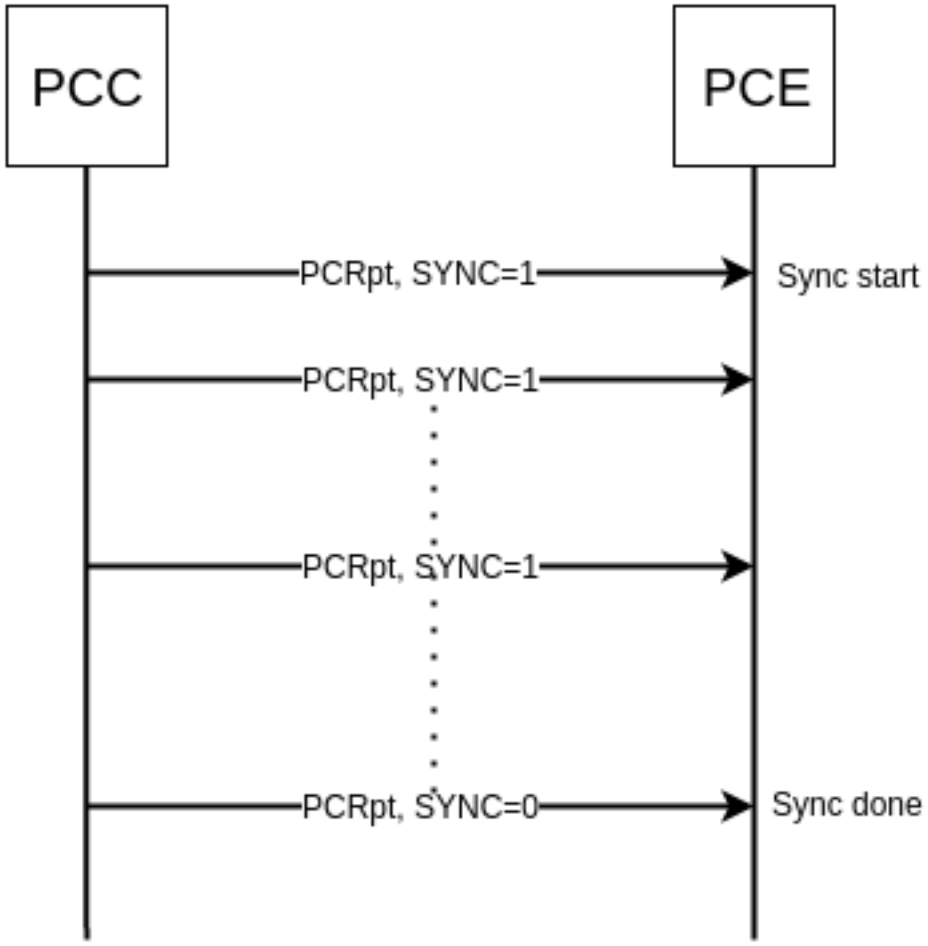

Fig. 82: LSP State Synchronization.

LSP-DB API

```
path-computation-client
   +--ro reported-lsp* [name]
      +--ro name        string
      +--ro path* [lsp-id]
      |  +--ro lsp-id                      rsvp:lsp-id
      |  +--ro ero
      |  |  +--ro processing-rule?   boolean
      |  |  +--ro ignore?            boolean
      |  |  +--ro subobject*
      |  |     +--ro loose         boolean
      |  |     +--ro (subobject-type)?
      |  |        +--:(as-number-case)
      |  |        |  +--ro as-number
      |  |        |     +--ro as-number    inet:as-number
      |  |        +--:(ip-prefix-case)
      |  |        |  +--ro ip-prefix
      |  |        |     +--ro ip-prefix    inet:ip-prefix
      |  |        +--:(label-case)
      |  |        |  +--ro label
      |  |        |     +--ro uni-directional             boolean
      |  |        |     +--ro (label-type)?
      |  |        |        +--:(type1-label-case)
      |  |        |        |  +--ro type1-label
      |  |        |        |     +--ro type1-label    uint32
      |  |        |        +--:(generalized-label-case)
      |  |        |        |  +--ro generalized-label
      |  |        |        |     +--ro generalized-label    binary
      |  |        |        +--:(waveband-switching-label-case)
      |  |        |           +--ro waveband-switching-label
      |  |        |              +--ro end-label      uint32
      |  |        |              +--ro start-label    uint32
      |  |        |              +--ro waveband-id    uint32
      |  |        +--:(srlg-case)
      |  |        |  +--ro srlg
      |  |        |     +--ro srlg-id    srlg-id
      |  |        +--:(unnumbered-case)
      |  |        |  +--ro unnumbered
      |  |        |     +--ro router-id       uint32
      |  |        |     +--ro interface-id    uint32
      |  |        +--:(exrs-case)
      |  |        |  +--ro exrs
      |  |        |     +--ro exrs*
      |  |        |        +--ro mandatory?    boolean
      |  |        |        +--ro attribute     enumeration
      |  |        |        +--ro (subobject-type)?
      |  |        |           +--:(as-number-case)
      |  |        |           |  +--ro as-number
      |  |        |           |     +--ro as-number    inet:as-number
      |  |        |           +--:(ip-prefix-case)
      |  |        |           |  +--ro ip-prefix
      |  |        |           |     +--ro ip-prefix    inet:ip-prefix
      |  |        |           +--:(label-case)
      |  |        |           |  +--ro label
      |  |        |           |     +--ro uni-directional             boolean
      |  |        |           |     +--ro (label-type)?
      |  |        |           |        +--:(type1-label-case)
      |  |        |           |        |  +--ro type1-label
```

(continues on next page)

(continued from previous page)

```
|  |  |        |           |        |     +--ro type1-label    uint32
|  |  |        |           |        +--:(generalized-label-case)
|  |  |        |           |        |  +--ro generalized-label
|  |  |        |           |        |     +--ro generalized-label    binary
|  |  |        |           |        +--:(waveband-switching-label-case)
|  |  |        |           |           +--ro waveband-switching-label
|  |  |        |           |              +--ro end-label      uint32
|  |  |        |           |              +--ro start-label    uint32
|  |  |        |           |              +--ro waveband-id    uint32
|  |  |        |           +--:(srlg-case)
|  |  |        |           |  +--ro srlg
|  |  |        |           |     +--ro srlg-id    srlg-id
|  |  |        |           +--:(unnumbered-case)
|  |  |        |              +--ro unnumbered
|  |  |        |                 +--ro router-id       uint32
|  |  |        |                 +--ro interface-id    uint32
|  |  |        +--:(path-key-case)
|  |  |           +--ro path-key
|  |  |              +--ro pce-id      pce-id
|  |  |              +--ro path-key    path-key
|  +--ro lspa
|  |  +--ro processing-rule?            boolean
|  |  +--ro ignore?                     boolean
|  |  +--ro hold-priority?              uint8
|  |  +--ro setup-priority?             uint8
|  |  +--ro local-protection-desired?   boolean
|  |  +--ro label-recording-desired?    boolean
|  |  +--ro se-style-desired?           boolean
|  |  +--ro session-name?               string
|  |  +--ro include-any?                attribute-filter
|  |  +--ro exclude-any?                attribute-filter
|  |  +--ro include-all?                attribute-filter
|  |  +--ro tlvs
|  |     +--ro vendor-information-tlv*
|  |        +--ro enterprise-number?   iana:enterprise-number
|  |        +--ro (enterprise-specific-information)?
|  +--ro bandwidth
|  |  +--ro processing-rule?   boolean
|  |  +--ro ignore?            boolean
|  |  +--ro bandwidth?         netc:bandwidth
|  +--ro reoptimization-bandwidth
|  |  +--ro processing-rule?   boolean
|  |  +--ro ignore?            boolean
|  |  +--ro bandwidth?         netc:bandwidth
|  +--ro metrics*
|  |  +--ro metric
|  |     +--ro processing-rule?   boolean
|  |     +--ro ignore?            boolean
|  |     +--ro metric-type        uint8
|  |     +--ro bound?             boolean
|  |     +--ro computed?          boolean
|  |     +--ro value?             ieee754:float32
|  +--ro iro
|  |  +--ro processing-rule?   boolean
|  |  +--ro ignore?            boolean
|  |  +--ro subobject*
|  |     +--ro loose         boolean
```

(continues on next page)

(continued from previous page)

```
      |  |     +--ro (subobject-type)?
      |  |        +--:(as-number-case)
      |  |        |  +--ro as-number
      |  |        |     +--ro as-number    inet:as-number
      |  |        +--:(ip-prefix-case)
      |  |        |  +--ro ip-prefix
      |  |        |     +--ro ip-prefix    inet:ip-prefix
      |  |        +--:(label-case)
      |  |        |  +--ro label
      |  |        |     +--ro uni-directional             boolean
      |  |        |     +--ro (label-type)?
      |  |        |        +--:(type1-label-case)
      |  |        |        |  +--ro type1-label
      |  |        |        |     +--ro type1-label    uint32
      |  |        |        +--:(generalized-label-case)
      |  |        |        |  +--ro generalized-label
      |  |        |        |     +--ro generalized-label    binary
      |  |        |        +--:(waveband-switching-label-case)
      |  |        |           +--ro waveband-switching-label
      |  |        |              +--ro end-label      uint32
      |  |        |              +--ro start-label    uint32
      |  |        |              +--ro waveband-id    uint32
      |  |        +--:(srlg-case)
      |  |        |  +--ro srlg
      |  |        |     +--ro srlg-id    srlg-id
      |  |        +--:(unnumbered-case)
      |  |        |  +--ro unnumbered
      |  |        |     +--ro router-id       uint32
      |  |        |     +--ro interface-id    uint32
      |  |        +--:(exrs-case)
      |  |        |  +--ro exrs
      |  |        |     +--ro exrs*
      |  |        |        +--ro mandatory?    boolean
      |  |        |        +--ro attribute     enumeration
      |  |        |        +--ro (subobject-type)?
      |  |        |           +--:(as-number-case)
      |  |        |           |  +--ro as-number
      |  |        |           |     +--ro as-number    inet:as-number
      |  |        |           +--:(ip-prefix-case)
      |  |        |           |  +--ro ip-prefix
      |  |        |           |     +--ro ip-prefix    inet:ip-prefix
      |  |        |           +--:(label-case)
      |  |        |           |  +--ro label
      |  |        |           |     +--ro uni-directional             boolean
      |  |        |           |     +--ro (label-type)?
      |  |        |           |        +--:(type1-label-case)
      |  |        |           |        |  +--ro type1-label
      |  |        |           |        |     +--ro type1-label    uint32
      |  |        |           |        +--:(generalized-label-case)
      |  |        |           |        |  +--ro generalized-label
      |  |        |           |        |     +--ro generalized-label    binary
      |  |        |           |        +--:(waveband-switching-label-case)
      |  |        |           |           +--ro waveband-switching-label
      |  |        |           |              +--ro end-label      uint32
      |  |        |           |              +--ro start-label    uint32
      |  |        |           |              +--ro waveband-id    uint32
      |  |        |           +--:(srlg-case)
```

(continues on next page)

(continued from previous page)

```
|  |  |        |           |  +--ro srlg
|  |  |        |           |     +--ro srlg-id    srlg-id
|  |  |        |           +--:(unnumbered-case)
|  |  |        |              +--ro unnumbered
|  |  |        |                 +--ro router-id       uint32
|  |  |        |                 +--ro interface-id    uint32
|  |  |        +--:(path-key-case)
|  |  |           +--ro path-key
|  |  |              +--ro pce-id      pce-id
|  |  |              +--ro path-key    path-key
|  +--ro rro
|  |  +--ro processing-rule?   boolean
|  |  +--ro ignore?            boolean
|  |  +--ro subobject*
|  |     +--ro protection-available?   boolean
|  |     +--ro protection-in-use?      boolean
|  |     +--ro (subobject-type)?
|  |        +--:(ip-prefix-case)
|  |        |  +--ro ip-prefix
|  |        |     +--ro ip-prefix    inet:ip-prefix
|  |        +--:(label-case)
|  |        |  +--ro label
|  |        |     +--ro uni-directional             boolean
|  |        |     +--ro (label-type)?
|  |        |     |  +--:(type1-label-case)
|  |        |     |  |  +--ro type1-label
|  |        |     |  |     +--ro type1-label    uint32
|  |        |     |  +--:(generalized-label-case)
|  |        |     |  |  +--ro generalized-label
|  |        |     |  |     +--ro generalized-label    binary
|  |        |     |  +--:(waveband-switching-label-case)
|  |        |     |     +--ro waveband-switching-label
|  |        |     |        +--ro end-label      uint32
|  |        |     |        +--ro start-label    uint32
|  |        |     |        +--ro waveband-id    uint32
|  |        |     +--ro global?                     boolean
|  |        +--:(unnumbered-case)
|  |        |  +--ro unnumbered
|  |        |     +--ro router-id       uint32
|  |        |     +--ro interface-id    uint32
|  |        +--:(path-key-case)
|  |           +--ro path-key
|  |              +--ro pce-id      pce-id
|  |              +--ro path-key    path-key
|  +--ro xro
|  |  +--ro processing-rule?   boolean
|  |  +--ro ignore?            boolean
|  |  +--ro flags              bits
|  |  +--ro subobject*
|  |     +--ro mandatory?    boolean
|  |     +--ro attribute     enumeration
|  |     +--ro (subobject-type)?
|  |        +--:(as-number-case)
|  |        |  +--ro as-number
|  |        |     +--ro as-number    inet:as-number
|  |        +--:(ip-prefix-case)
|  |        |  +--ro ip-prefix
```

(continues on next page)

(continued from previous page)

```
      |  |        |     +--ro ip-prefix    inet:ip-prefix
      |  |        +--:(label-case)
      |  |        |  +--ro label
      |  |        |     +--ro uni-directional             boolean
      |  |        |     +--ro (label-type)?
      |  |        |        +--:(type1-label-case)
      |  |        |        |  +--ro type1-label
      |  |        |        |     +--ro type1-label    uint32
      |  |        |        +--:(generalized-label-case)
      |  |        |        |  +--ro generalized-label
      |  |        |        |     +--ro generalized-label    binary
      |  |        |        +--:(waveband-switching-label-case)
      |  |        |           +--ro waveband-switching-label
      |  |        |              +--ro end-label      uint32
      |  |        |              +--ro start-label    uint32
      |  |        |              +--ro waveband-id    uint32
      |  |        +--:(srlg-case)
      |  |        |  +--ro srlg
      |  |        |     +--ro srlg-id    srlg-id
      |  |        +--:(unnumbered-case)
      |  |           +--ro unnumbered
      |  |              +--ro router-id       uint32
      |  |              +--ro interface-id    uint32
      |  +--ro of
      |  |  +--ro processing-rule?   boolean
      |  |  +--ro ignore?            boolean
      |  |  +--ro code               of-id
      |  |  +--ro tlvs
      |  |     +--ro vendor-information-tlv*
      |  |        +--ro enterprise-number?   iana:enterprise-number
      |  |        +--ro (enterprise-specific-information)?
      |  +--ro class-type
      |     +--ro processing-rule?   boolean
      |     +--ro ignore?            boolean
      |     +--ro class-type         class-type
      +--ro metadata
      +--ro lsp
      |  +--ro processing-rule?   boolean
      |  +--ro ignore?            boolean
      |  +--ro tlvs
      |  |  +--ro lsp-error-code
      |  |  |  +--ro error-code?   uint32
      |  |  +--ro lsp-identifiers
      |  |  |  +--ro lsp-id?      rsvp:lsp-id
      |  |  |  +--ro tunnel-id?   rsvp:tunnel-id
      |  |  |  +--ro (address-family)?
      |  |  |     +--:(ipv4-case)
      |  |  |     |  +--ro ipv4
      |  |  |     |     +--ro ipv4-tunnel-sender-address      inet:ipv4-address
      |  |  |     |     +--ro ipv4-extended-tunnel-id         rsvp:ipv4-extended-
↪tunnel-id
      |  |  |     |     +--ro ipv4-tunnel-endpoint-address    inet:ipv4-address
      |  |  |     +--:(ipv6-case)
      |  |  |        +--ro ipv6
      |  |  |           +--ro ipv6-tunnel-sender-address      inet:ipv6-address
      |  |  |           +--ro ipv6-extended-tunnel-id         rsvp:ipv6-extended-
↪tunnel-id
```

(continues on next page)

(continued from previous page)

```
      |  |  |           +--ro ipv6-tunnel-endpoint-address    inet:ipv6-address
      |  |  +--ro rsvp-error-spec
      |  |  |  +--ro (error-type)?
      |  |  |     +--:(rsvp-case)
      |  |  |     |  +--ro rsvp-error
      |  |  |     +--:(user-case)
      |  |  |        +--ro user-error
      |  |  +--ro symbolic-path-name
      |  |  |  +--ro path-name?   symbolic-path-name
      |  |  o--ro vs-tlv
      |  |  |  +--ro enterprise-number?   iana:enterprise-number
      |  |  |  +--ro (vendor-payload)?
      |  |  +--ro vendor-information-tlv*
      |  |  |  +--ro enterprise-number?   iana:enterprise-number
      |  |  |  +--ro (enterprise-specific-information)?
      |  |  +--ro path-binding
      |  |     x--ro binding-type?      uint8
      |  |     x--ro binding-value?     binary
      |  |     +--ro (binding-type-value)?
      |  |        +--:(mpls-label)
      |  |        |  +--ro mpls-label?        netc:mpls-label
      |  |        +--:(mpls-label-entry)
      |  |           +--ro label?             netc:mpls-label
      |  |           +--ro traffic-class?     uint8
      |  |           +--ro bottom-of-stack?   boolean
      |  |           +--ro time-to-live?      uint8
      |  +--ro plsp-id?           plsp-id
      |  +--ro delegate?          boolean
      |  +--ro sync?              boolean
      |  +--ro remove?            boolean
      |  +--ro administrative?    boolean
      |  +--ro operational?       operational-status
      +--ro path-setup-type
         +--ro pst?   uint8
```

The LSP-DB is accessible via RESTCONF. The PCC's LSPs are stored in the `pcep-topology` while the session is active. In a next example, there is one PCEP session with PCC identified by its IP address (*43.43.43.43*) and one reported LSP (*foo*).

URL: `/restconf/operational/network-topology:network-topology/topology/pcep-topology/node/pcc:%2F%2F43.43.43.43`

Method: `GET`

Response Body:

```
<node>
   <node-id>pcc://43.43.43.43</node-id>
   <path-computation-client>
      <ip-address>43.43.43.43</ip-address>
      <state-sync>synchronized</state-sync>
      <stateful-tlv>
         <stateful>
            <lsp-update-capability>true</lsp-update-capability>
         </stateful>
      </stateful-tlv>
```

(continues on next page)

(continued from previous page)

```
11      <reported-lsp>
12         <name>foo</name>
13         <lsp>
14            <operational>up</operational>
15            <sync>true</sync>
16            <plsp-id>1</plsp-id>
17            <create>false</create>
18            <administrative>true</administrative>
19            <remove>false</remove>
20            <delegate>true</delegate>
21            <tlvs>
22               <lsp-identifiers>
23                  <ipv4>
24                     <ipv4-tunnel-sender-address>43.43.43.43</ipv4-tunnel-sender-
   ↪address>
25                     <ipv4-tunnel-endpoint-address>39.39.39.39</ipv4-tunnel-endpoint-
   ↪address>
26                     <ipv4-extended-tunnel-id>39.39.39.39</ipv4-extended-tunnel-id>
27                  </ipv4>
28                  <tunnel-id>1</tunnel-id>
29                  <lsp-id>1</lsp-id>
30               </lsp-identifiers>
31               <symbolic-path-name>
32                  <path-name>Zm9v</path-name>
33               </symbolic-path-name>
34            </tlvs>
35         </lsp>
36         <ero>
37            <subobject>
38               <loose>false</loose>
39               <ip-prefix>
40                  <ip-prefix>201.20.160.40/32</ip-prefix>
41               </ip-prefix>
42            </subobject>
43            <subobject>
44               <loose>false</loose>
45               <ip-prefix>
46                  <ip-prefix>195.20.160.39/32</ip-prefix>
47               </ip-prefix>
48            </subobject>
49            <subobject>
50               <loose>false</loose>
51               <ip-prefix>
52                  <ip-prefix>39.39.39.39/32</ip-prefix>
53               </ip-prefix>
54            </subobject>
55         </ero>
56      </reported-lsp>
57   </path-computation-client>
58 </node>
```

@line 2: **node-id** The PCC identifier.

@line 4: **ip-address** IP address of the PCC.

@line 5: **state-sync** Synchronization status of the PCC's LSPs. The *synchronized* indicates the State Synchronization is done.

@line 8: **lsp-update-capability** - Indicates that PCC allows LSP modifications.

@line 12: **name** - Textual representation of LPS's name.

@line 14: **operational** - Represent operational status of the LSP:

- *down* - not active.
- *up* - signaled.
- *active* - up and carrying traffic.
- *going-down* - LSP is being torn down, resources are being released.
- *going-up* - LSP is being signaled.

@line 15: **sync** - The flag set by PCC during LSPs State Synchronization.

@line 16: **plsp-id** - A PCEP-specific identifier for the LSP. It is assigned by PCC and it is constant for a lifetime of a PCEP session.

@line 17: **create** - The *false* indicates that LSP is PCC-initiated.

@line 18: **administrative** - The flag indicates target operational status of the LSP.

@line 20: **delegate** - The delegate flag indicates that the PCC is delegating the LSP to the PCE.

@line 24: **ipv4-tunnel-sender-address** - Contains the sender node's IP address.

@line 25: **ipv4-tunnel-endpoint-address** - Contains the egress node's IP address.

@line 26: **ipv4-extended-tunnel-id** - The *Extended Tunnel ID* identifier.

@line 28: **tunnel-id** - The *Tunnel ID* identifier.

@line 29: **lsp-id** - The *LSP ID* identifier.

@line 32: **path-name** - The symbolic name for the LSP.

@line 36: **ero** - The *Explicit Route Object* is encoding the path of the TE LSP through the network.

LSP Delegation

The LSP control delegations is an mechanism, where PCC grants to a PCE the temporary right in order to modify LSP attributes. The PCC can revoke the delegation or the PCE may waive the delegation at any time. The LSP control is delegated to at most one PCE at the same time.

Following RPC example illustrates a request for the LSP delegation give up:

URL: `/restconf/operations/network-topology-pcep:update-lsp`

Method: `POST`

Content-Type: `application/xml`

Request Body:

```
<input>
   <node>pcc://43.43.43.43</node>
   <name>foo</name>
   <arguments>
      <lsp xmlns:stateful="urn:opendaylight:params:xml:ns:yang:pcep:ietf:stateful">
         <delegate>false</delegate>
```

(continues on next page)

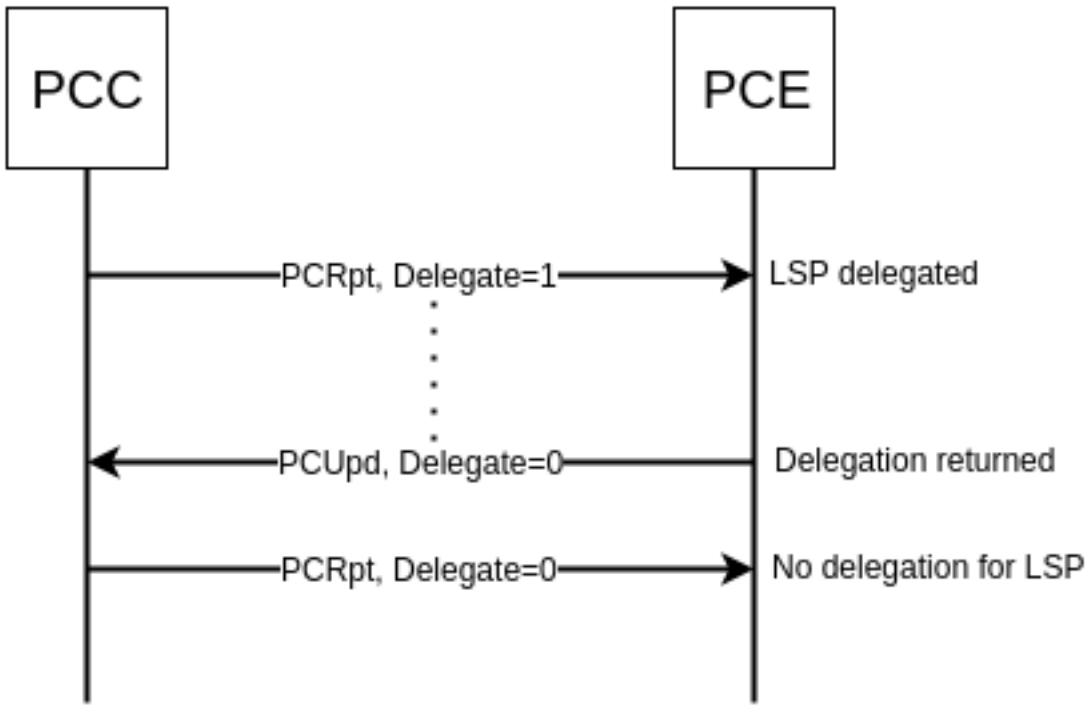

Fig. 83: Returning a Delegation.

(continued from previous page)

```
7           <administrative>true</administrative>
8           <tlvs>
9              <symbolic-path-name>
10                <path-name>Zm9v</path-name>
11             </symbolic-path-name>
12          </tlvs>
13       </lsp>
14    </arguments>
15    <network-topology-ref xmlns:topo="urn:TBD:params:xml:ns:yang:network-topology">/
 ↪topo:network-topology/topo:topology[topo:topology-id="pcep-topology"]</network-
 ↪topology-ref>
16 </input>
```

@line 2: **node** The PCC identifier.

@line 3: **name** The name of the LSP.

@line 6: **delegate** - Delegation flag set *false* in order to return the LSP delegation.

@line 10: **path-name** - The Symbolic Path Name TLV must be present when sending a request to give up the delegation.

LSP Update

The LSP Update Request is an operation where a PCE requests a PCC to update attributes of an LSP and to rebuild the LSP with updated attributes. In order to update LSP, the PCE must hold a LSP delegation. The LSP update is done in *make-before-break* fashion - first, new LSP is initiated and then the old LSP is torn down.

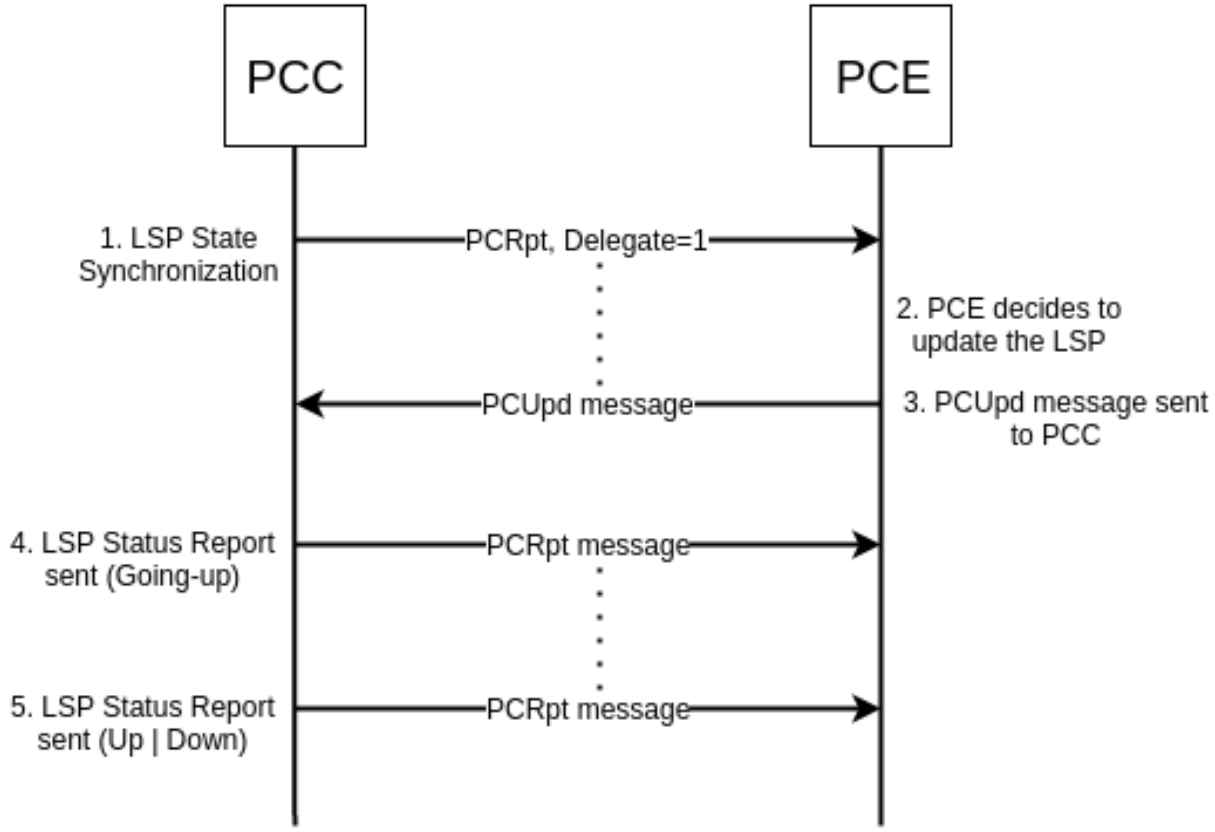

Fig. 84: Active Stateful PCE LSP Update.

Following RPC example shows a request for the LSP update:

URL: `/restconf/operations/network-topology-pcep:update-lsp`

Method: `POST`

Content-Type: `application/xml`

Request Body:

```
<input xmlns="urn:opendaylight:params:xml:ns:yang:topology:pcep">
   <node>pcc://43.43.43.43</node>
   <name>foo</name>
   <arguments>
      <lsp xmlns="urn:opendaylight:params:xml:ns:yang:pcep:ietf:stateful">
         <delegate>true</delegate>
         <administrative>true</administrative>
      </lsp>
      <ero>
         <subobject>
            <loose>false</loose>
            <ip-prefix>
               <ip-prefix>200.20.160.41/32</ip-prefix>
            </ip-prefix>
         </subobject>
         <subobject>
            <loose>false</loose>
            <ip-prefix>
               <ip-prefix>196.20.160.39/32</ip-prefix>
            </ip-prefix>
         </subobject>
         <subobject>
            <loose>false</loose>
            <ip-prefix>
               <ip-prefix>39.39.39.39/32</ip-prefix>
            </ip-prefix>
         </subobject>
      </ero>
   </arguments>
   <network-topology-ref xmlns:topo="urn:TBD:params:xml:ns:yang:network-topology">/
↪topo:network-topology/topo:topology[topo:topology-id="pcep-topology"]</network-
↪topology-ref>
</input>
```

@line 2: **node** The PCC identifier.

@line 3: **name** The name of the LSP to be updated.

@line 6: **delegate** - Delegation flag set *true* in order to keep the LSP control.

@line 7: **administrative** - Desired administrative status of the LSP is active.

@line 9: **ero** - This LSP attribute is changed.

PCE-initiated LSP Setup

The PCEP Extension for PCE-initiated LSP Setup allows PCE to request a creation and deletion of LSPs.

Configuration

This capability is enabled by default. No additional configuration is required.

LSP Instantiation

The PCE can request LSP creation. The LSP instantiation is done by sending an LSP Initiate Message to PCC. The PCC assign delegation to PCE which triggered creation. PCE-initiated LSPs are identified by *Create* flag.

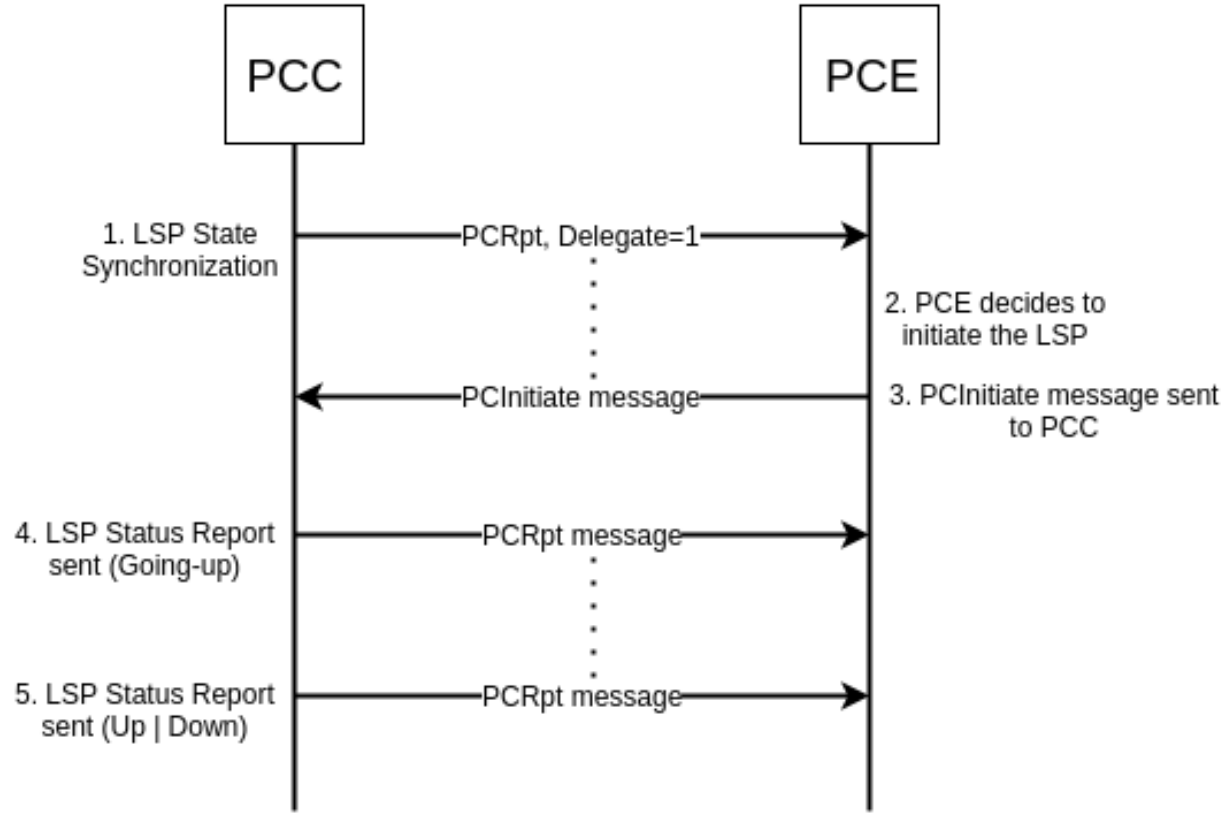

Fig. 85: LSP instantiation.

Following RPC example shows a request for the LSP initiation:

URL: `/restconf/operations/network-topology-pcep:add-lsp`

Method: `POST`

Content-Type: `application/xml`

Request Body:

```
<input xmlns="urn:opendaylight:params:xml:ns:yang:topology:pcep">
   <node>pcc://43.43.43.43</node>
   <name>update-tunel</name>
      <arguments>
         <lsp xmlns="urn:opendaylight:params:xml:ns:yang:pcep:ietf:stateful">
```

(continues on next page)

(continued from previous page)

```
6             <delegate>true</delegate>
7             <administrative>true</administrative>
8          </lsp>
9          <endpoints-obj>
10            <ipv4>
11               <source-ipv4-address>43.43.43.43</source-ipv4-address>
12               <destination-ipv4-address>39.39.39.39</destination-ipv4-address>
13            </ipv4>
14         </endpoints-obj>
15         <ero>
16            <subobject>
17               <loose>false</loose>
18               <ip-prefix>
19                  <ip-prefix>201.20.160.40/32</ip-prefix>
20               </ip-prefix>
21            </subobject>
22            <subobject>
23               <loose>false</loose>
24               <ip-prefix>
25                  <ip-prefix>195.20.160.39/32</ip-prefix>
26               </ip-prefix>
27            </subobject>
28            <subobject>
29               <loose>false</loose>
30               <ip-prefix>
31                  <ip-prefix>39.39.39.39/32</ip-prefix>
32               </ip-prefix>
33            </subobject>
34         </ero>
35      </arguments>
36   <network-topology-ref xmlns:topo="urn:TBD:params:xml:ns:yang:network-topology">/
 ↪topo:network-topology/topo:topology[topo:topology-id="pcep-topology"]</network-
 ↪topology-ref>
37 </input>
```

@line 2: **node** The PCC identifier.

@line 3: **name** The name of the LSP to be created.

@line 8: **endpoints-obj** - The *END-POINT* Object is mandatory for an instantiation request of an RSVP-signaled LSP. It contains source and destination addresses for provisioning the LSP.

@line 14: **ero** - The *ERO* object is mandatory for LSP initiation request.

LSP Deletion

The PCE may request a deletion of PCE-initiated LSPs. The PCE must be delegation holder for this particular LSP.

Following RPC example shows a request for the LSP deletion:

URL: `/restconf/operations/network-topology-pcep:remove-lsp`

Method: `POST`

Content-Type: `application/xml`

Request Body:

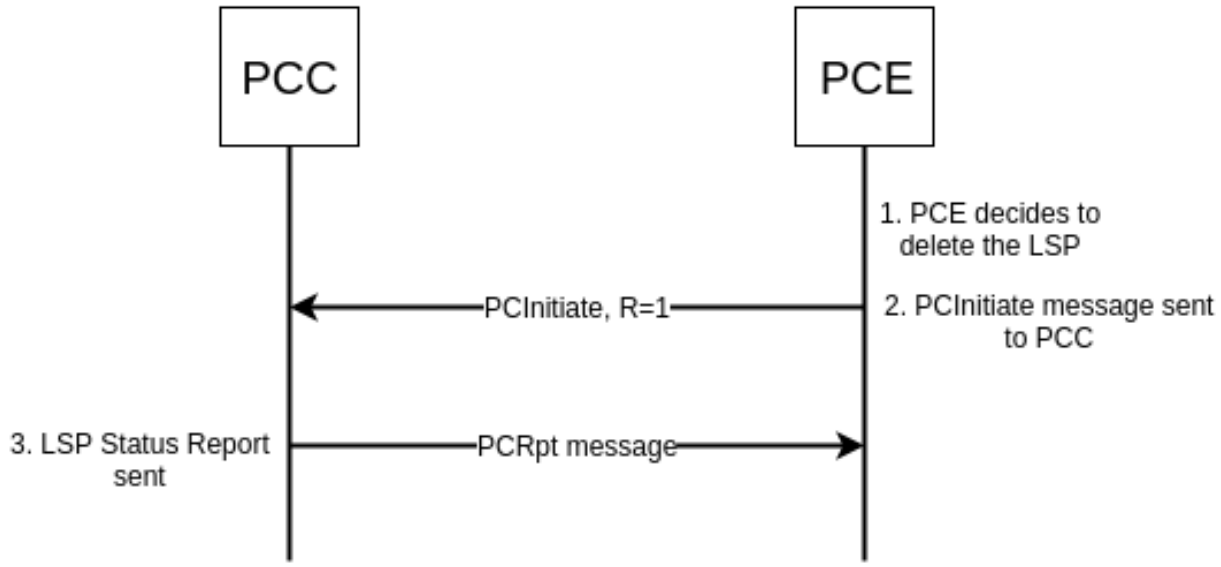

Fig. 86: LSP deletion.

```
1 <input xmlns="urn:opendaylight:params:xml:ns:yang:topology:pcep">
2    <node>pcc://43.43.43.43</node>
3    <name>update-tunel</name>
4    <network-topology-ref xmlns:topo="urn:TBD:params:xml:ns:yang:network-topology">/
  ↪topo:network-topology/topo:topology[topo:topology-id="pcep-topology"]</network-
  ↪topology-ref>
5 </input>
```

@line 2: **node** The PCC identifier.

@line 3: **name** The name of the LSP to be removed.

PCE-initiated LSP Delegation

The PCE-initiated LSP control is delegated to the PCE which requested the initiation. The PCC cannot revoke delegation of PCE-initiated LSP. When PCE returns delegation for such LSP or PCE fails, then the LSP become orphan and can be removed by a PCC after some time. The PCE may ask for a delegation of the orphan LSP.

Following RPC example illustrates a request for the LSP delegation:

URL: `/restconf/operations/network-topology-pcep:update-lsp`

Method: `POST`

Content-Type: `application/xml`

Request Body:

```
1 <input>
2    <node>pcc://43.43.43.43</node>
3    <name>update-tunel</name>
4    <arguments>
5       <lsp xmlns:stateful="urn:opendaylight:params:xml:ns:yang:pcep:ietf:stateful">
```

(continues on next page)

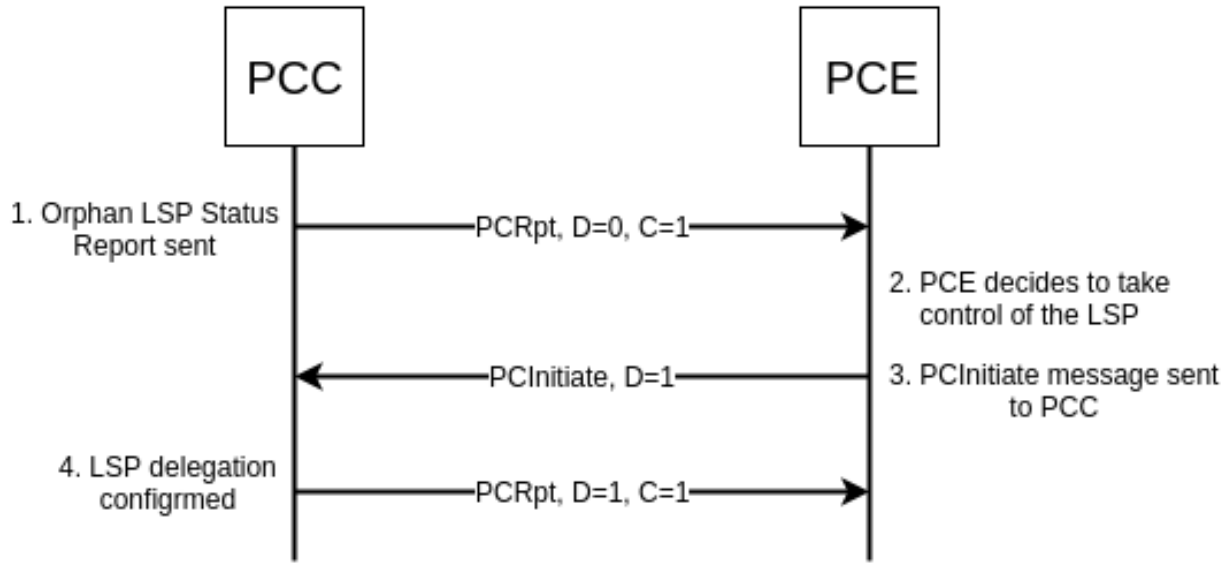

Fig. 87: Orphan PCE-initiated LSP - control taken by PCE.

(continued from previous page)

```
6           <delegate>true</delegate>
7           <administrative>true</administrative>
8           <tlvs>
9              <symbolic-path-name>
10                <path-name>dXBkYXRlLXR1bmVs</path-name>
11             </symbolic-path-name>
12          </tlvs>
13       </lsp>
14    </arguments>
15    <network-topology-ref xmlns:topo="urn:TBD:params:xml:ns:yang:network-topology">/
↪topo:network-topology/topo:topology[topo:topology-id="pcep-topology"]</network-
↪topology-ref>
16 </input>
```

@line 2: **node** The PCC identifier.

@line 3: **name** The name of the LSP.

@line 6: **delegate** - *Delegation* flag set *true* in order to take the LSP delegation.

@line 10: **path-name** - The *Symbolic Path Name* TLV must be present when sending a request to take a delegation.

Segment Routing

The PCEP Extensions for Segment Routing (SR) allow a stateful PCE to compute and initiate TE paths in SR networks. The SR path is defined as an order list of *segments*. Segment Routing architecture can be directly applied to the MPLS forwarding plane without changes. Segment Identifier (SID) is encoded as a MPLS label.

Configuration

This capability is enabled by default. In PCEP-SR draft version 6, SR Explicit Route Object/Record Route Object subobjects IANA code points change was proposed. In order to use the latest code points, a configuration should be

changed in following way:

URL: `/restconf/config/pcep-segment-routing-app-config:pcep-segment-routing-app-config`

Method: `PUT`

Content-Type: `application/xml`

Request Body:

```
<pcep-segment-routing-config xmlns=
↪"urn:opendaylight:params:xml:ns:yang:controller:pcep:segment-routing-app-config">
   <iana-sr-subobjects-type>true</iana-sr-subobjects-type>
</pcep-segment-routing-config>
```

LSP Operations for PCEP SR

The PCEP SR extension defines new ERO subobject - *SR-ERO subobject* capable of carrying a SID.

```
sr-ero-type
   +---- c-flag?                boolean
   +---- m-flag?                boolean
   +---- sid-type?              sid-type
   +---- sid?                   uint32
   +---- (nai)?
      +--:(ip-node-id)
      |  +---- ip-address             inet:ip-address
      +--:(ip-adjacency)
      |  +---- local-ip-address       inet:ip-address
      |  +---- remote-ip-address      inet:ip-address
      +--:(unnumbered-adjacency)
         +---- local-node-id          uint32
         +---- local-interface-id     uint32
         +---- remote-node-id         uint32
         +---- remote-interface-id    uint32
```

Following RPC example illustrates a request for the SR-TE LSP creation:

URL: `/restconf/operations/network-topology-pcep:add-lsp`

Method: `POST`

Content-Type: `application/xml`

Request Body:

```
<input xmlns="urn:opendaylight:params:xml:ns:yang:topology:pcep">
   <node>pcc://43.43.43.43</node>
   <name>sr-path</name>
   <arguments>
      <lsp xmlns="urn:opendaylight:params:xml:ns:yang:pcep:ietf:stateful">
         <delegate>true</delegate>
         <administrative>true</administrative>
      </lsp>
      <endpoints-obj>
         <ipv4>
            <source-ipv4-address>43.43.43.43</source-ipv4-address>
```

(continues on next page)

(continued from previous page)

```
12              <destination-ipv4-address>39.39.39.39</destination-ipv4-address>
13           </ipv4>
14        </endpoints-obj>
15        <path-setup-type xmlns="urn:opendaylight:params:xml:ns:yang:pcep:ietf:stateful">
16           <pst>1</pst>
17        </path-setup-type>
18        <ero>
19           <subobject>
20              <loose>false</loose>
21              <sid-type xmlns="urn:opendaylight:params:xml:ns:yang:pcep:segment:routing
    ">ipv4-node-id</sid-type>
22              <m-flag xmlns="urn:opendaylight:params:xml:ns:yang:pcep:segment:routing">
    true</m-flag>
23              <sid xmlns="urn:opendaylight:params:xml:ns:yang:pcep:segment:routing">
    24001</sid>
24              <ip-address xmlns=
    "urn:opendaylight:params:xml:ns:yang:pcep:segment:routing">39.39.39.39</ip-address>
25           </subobject>
26        </ero>
27     </arguments>
28     <network-topology-ref xmlns:topo="urn:TBD:params:xml:ns:yang:network-topology">/
    topo:network-topology/topo:topology[topo:topology-id="pcep-topology"]</network-
    topology-ref>
29  </input>
```

@line 16: **path-setup-type** - Set *1* for SR-TE LSP

@line 21: **ipv4-node-id** - The SR-ERO subobject represents *IPv4 Node ID* NAI.

@line 22: **m-flag** - The SID value represents an MPLS label.

@line 23: **sid** - The Segment Identifier.

Following RPC example illustrates a request for the SR-TE LSP update including modified path:

URL: `/restconf/operations/network-topology-pcep:update-lsp`

Method: `POST`

Content-Type: `application/xml`

Request Body:

```
1  <input xmlns="urn:opendaylight:params:xml:ns:yang:topology:pcep">
2     <node>pcc://43.43.43.43</node>
3     <name>update-tunnel</name>
4     <arguments>
5        <lsp xmlns="urn:opendaylight:params:xml:ns:yang:pcep:ietf:stateful">
6           <delegate>true</delegate>
7           <administrative>true</administrative>
8        </lsp>
9        <path-setup-type xmlns="urn:opendaylight:params:xml:ns:yang:pcep:ietf:stateful">
10          <pst>1</pst>
11       </path-setup-type>
12       <ero>
13          <subobject>
14             <loose>false</loose>
15             <sid-type xmlns="urn:opendaylight:params:xml:ns:yang:pcep:segment:routing
    ">ipv4-node-id</sid-type>
```

(continues on next page)

(continued from previous page)

```
            <m-flag xmlns="urn:opendaylight:params:xml:ns:yang:pcep:segment:routing">
↪true</m-flag>
            <sid xmlns="urn:opendaylight:params:xml:ns:yang:pcep:segment:routing">
↪24002</sid>
            <ip-address xmlns=
↪"urn:opendaylight:params:xml:ns:yang:pcep:segment:routing">200.20.160.41</ip-
↪address>
         </subobject>
         <subobject>
            <loose>false</loose>
            <sid-type xmlns="urn:opendaylight:params:xml:ns:yang:pcep:segment:routing
↪">ipv4-node-id</sid-type>
            <m-flag xmlns="urn:opendaylight:params:xml:ns:yang:pcep:segment:routing">
↪true</m-flag>
            <sid xmlns="urn:opendaylight:params:xml:ns:yang:pcep:segment:routing">
↪24001</sid>
            <ip-address xmlns=
↪"urn:opendaylight:params:xml:ns:yang:pcep:segment:routing">39.39.39.39</ip-address>
         </subobject>
      </ero>
   </arguments>
   <network-topology-ref xmlns:topo="urn:TBD:params:xml:ns:yang:network-topology">/
↪topo:network-topology/topo:topology[topo:topology-id="pcep-topology"]</network-
↪topology-ref>
</input>
```

LSP State Synchronization Optimization Procedures

This extension bring optimizations for state synchronization:

- State Synchronization Avoidance
- Incremental State Synchronization
- PCE-triggered Initial Synchronization
- PCE-triggered Re-synchronization

Configuration

This capability is enabled by default. No additional configuration is required.

State Synchronization Avoidance

The State Synchronization Avoidance procedure is intended to skip state synchronization if the state has survived and not changed during session restart.

Incremental State Synchronization

The Incremental State Synchronization procedure is intended to do incremental (delta) state synchronization when possible.

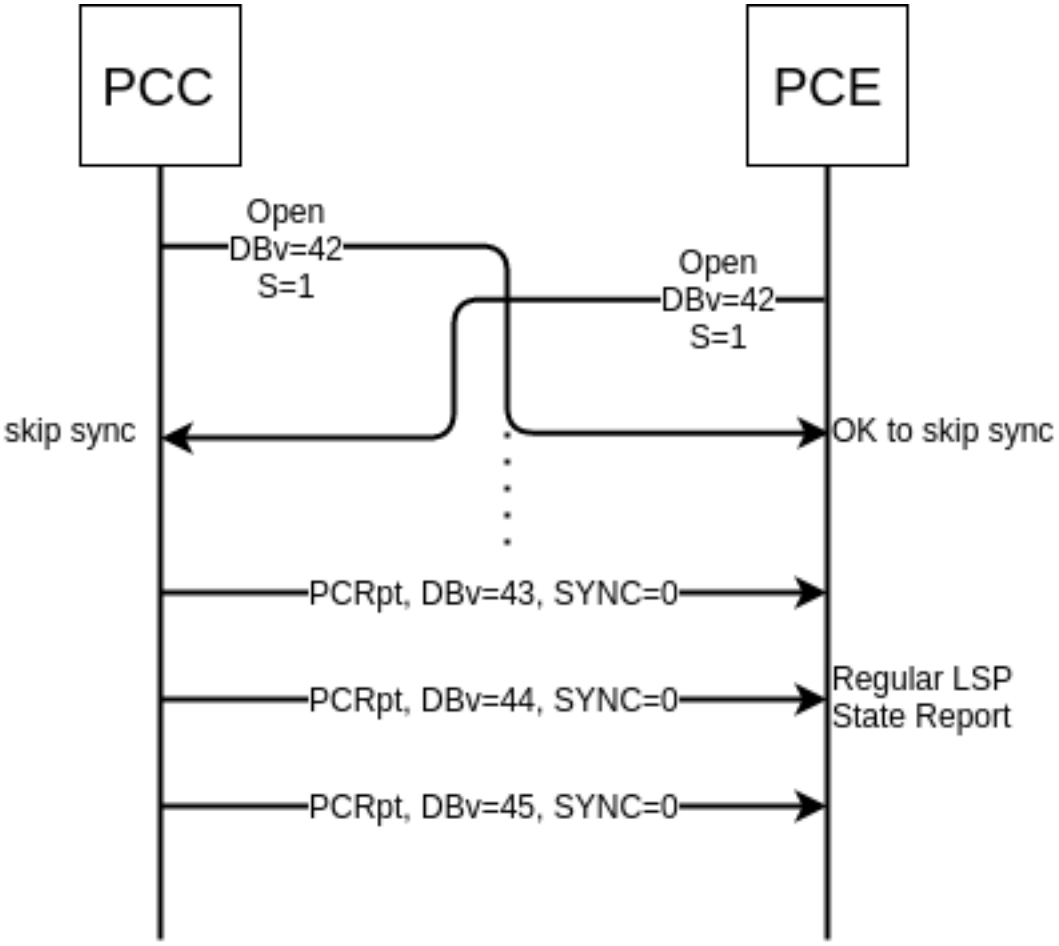

Fig. 88: State Synchronization Skipped.

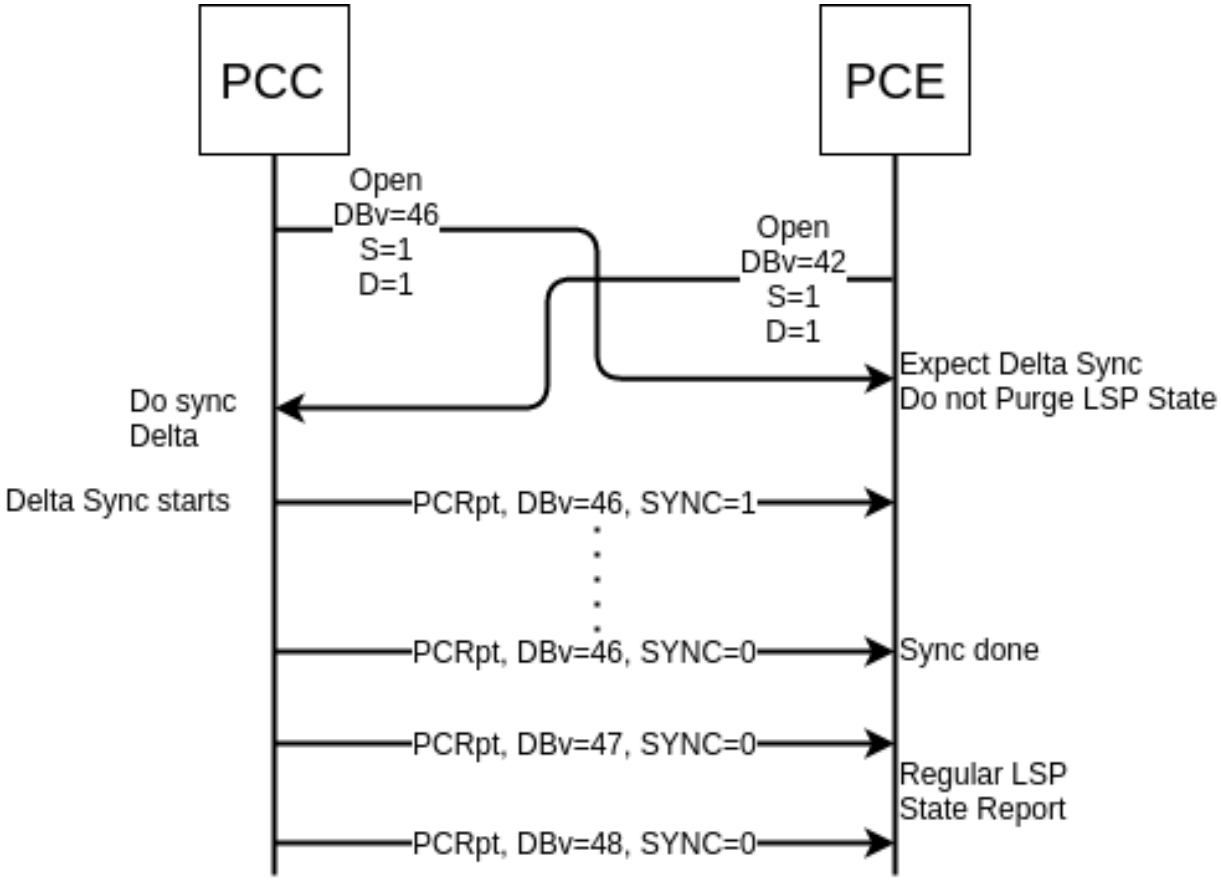

Fig. 89: Incremental Synchronization Procedure.

PCE-triggered Initial Synchronization

The PCE-triggered Initial Synchronization procedure is intended to do let PCE control the timing of the initial state synchronization.

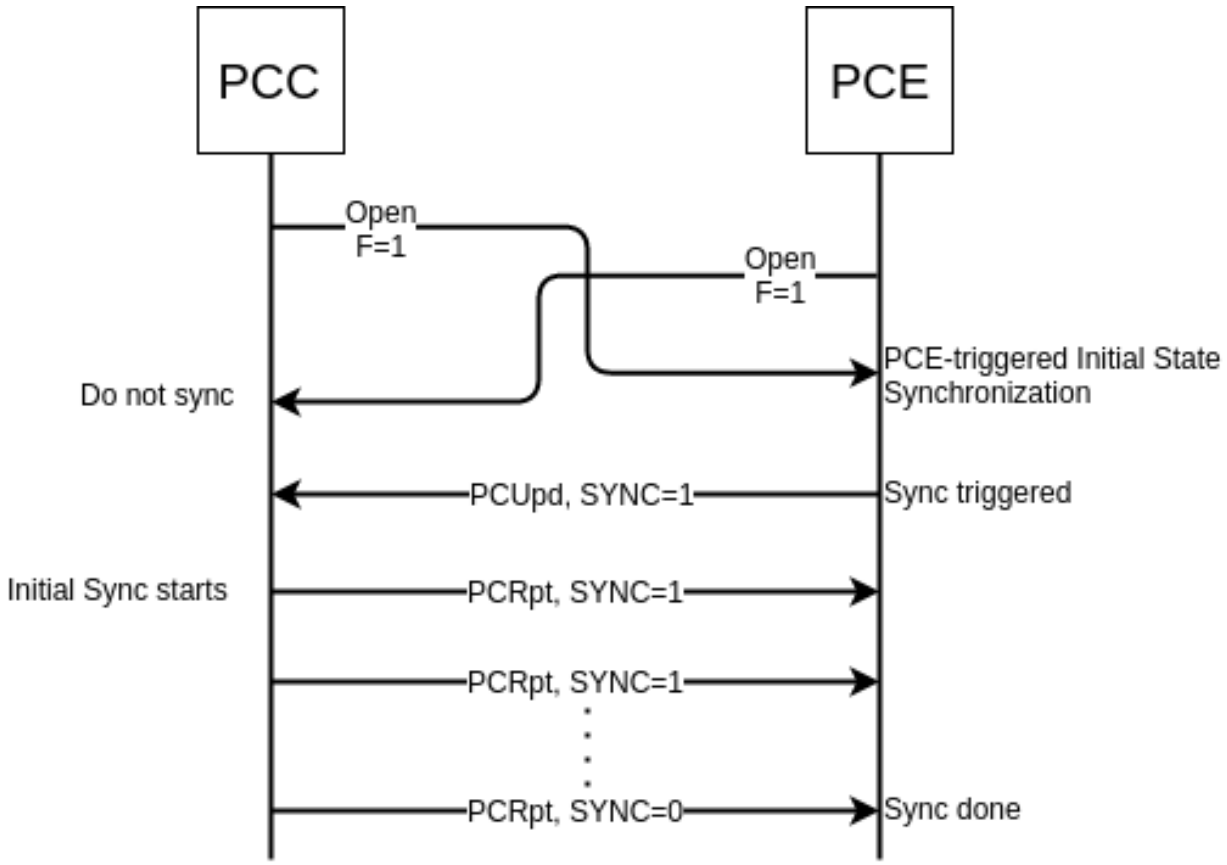

Fig. 90: PCE-triggered Initial State Synchronization Procedure.

Following RPC example illustrates a request for the initial synchronization:

URL: `/restconf/operations/network-topology-pcep:trigger-sync`

Method: `POST`

Content-Type: `application/xml`

Request Body:

```
<input xmlns="urn:opendaylight:params:xml:ns:yang:topology:pcep">
   <node>pcc://43.43.43.43</node>
   <network-topology-ref xmlns:topo="urn:TBD:params:xml:ns:yang:network-topology">/
↪topo:network-topology/topo:topology[topo:topology-id="pcep-topology"]</network-
↪topology-ref>
</input>
```

PCE-triggered Re-synchronization

The PCE-triggered Re-synchronization: To let PCE re-synchronize the state for sanity check.

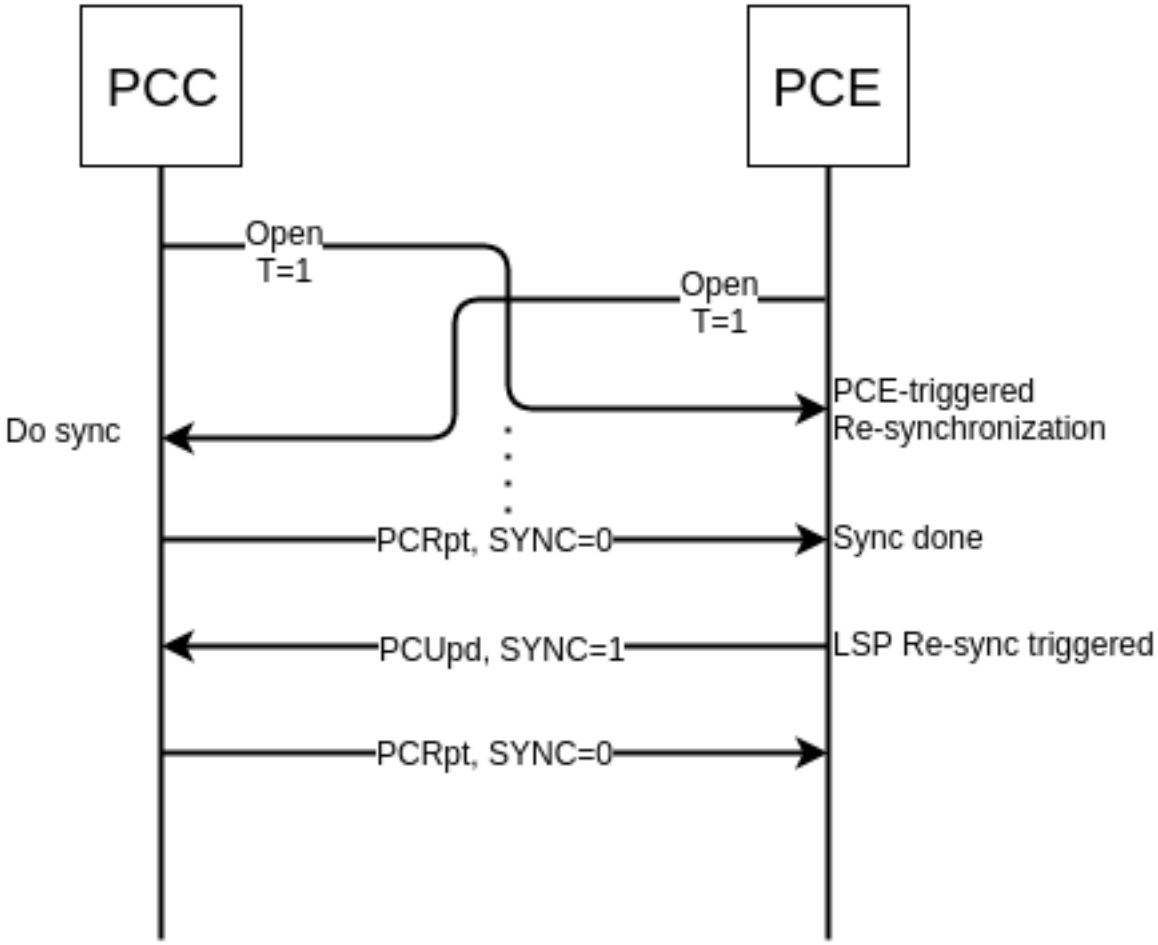

Fig. 91: PCE-triggered Re-synchronization Procedure.

Following RPC example illustrates a request for the LSP re-synchronization:

URL: `/restconf/operations/network-topology-pcep:trigger-sync`

Method: `POST`

Content-Type: `application/xml`

Request Body:

```
1 <input xmlns="urn:opendaylight:params:xml:ns:yang:topology:pcep">
2    <node>pcc://43.43.43.43</node>
3    <name>update-lsp</name>
4    <network-topology-ref xmlns:topo="urn:TBD:params:xml:ns:yang:network-topology">/
  ↪topo:network-topology/topo:topology[topo:topology-id="pcep-topology"]</network-
  ↪topology-ref>
5 </input>
```

@line 3: **name** - The LSP name. If this parameter is omitted, re-synchronization is requested for all PCC's LSPs.

Test tools

PCC Mock

The PCC Mock is a stand-alone Java application purposed to simulate a PCC(s). The simulator is capable to report sample LSPs, respond to delegation, LSP management operations and synchronization optimization procedures. This application is not part of the OpenDaylight Karaf distribution, however it can be downloaded from OpenDaylight's Nexus (use latest release version):

`https://nexus.opendaylight.org/content/repositories/opendaylight.release/org/opendaylight/bgpcep/pcep-pcc-mock`

Usage

The application can be run from command line:

```
java -jar pcep-pcc-mock-*-executable.jar
```

with optional input parameters:

```
--local-address <Address:Port> (optional, default 127.0.0.1)
   The first PCC IP address. If more PCCs are required, the IP address will be␣
↪incremented. Port number can be optionally specified.

--remote-address <Address1:Port1,Address2:Port2,Address3:Port3,...> (optional,␣
↪default 127.0.0.1:4189)
   The list of IP address for the PCE servers. Port number can be optionally␣
↪specified, otherwise default port number 4189 is used.

--pcc <N> (optional, default 1)
   Number of mocked PCC instances.

--lsp <N> (optional, default 1)
   Number of tunnels (LSPs) reported per PCC, might be zero.

--pcerr (optional flag)
   If the flag is present, response with PCErr, otherwise PCUpd.

--log-level <LEVEL> (optional, default INFO)
   Set logging level for pcc-mock.

-d, --deadtimer <0..255> (optional, default 120)
   DeadTimer value in seconds.

-ka, --keepalive <0.255> (optional, default 30)
   KeepAlive timer value in seconds.

--password <password> (optional)
   If the password is present, it is used in TCP MD5 signature, otherwise plain TCP␣
↪is used.

--reconnect <seconds> (optional)
   If the argument is present, the value in seconds, is used as a delay before each␣
↪new reconnect (initial connect or connection re-establishment) attempt.
   The number of reconnect attempts is unlimited. If the argument is omitted, pcc-
↪mock is not trying to reconnect.
```

(continues on next page)

(continued from previous page)

```
--redelegation-timeout <seconds> (optional, default 0)
   The timeout starts when LSP delegation is returned or PCE fails, stops when LSP is␣
↪re-delegated to PCE.
   When timeout expires, LSP delegation is revoked and held by PCC.

--state-timeout <seconds> (optional, default -1 (disabled))
   The timeout starts when LSP delegation is returned or PCE fails, stops when LSP is␣
↪re-delegated to PCE.
   When timeout expires, PCE-initiated LSP is removed.

--state-sync-avoidance <disconnect_after_x_seconds> <reconnect_after_x_seconds>
↪<dbVersion>
   Synchronization avoidance capability enabled.
      - disconnect_after_x_seconds: seconds that will pass until disconnections is␣
↪forced. If set to smaller number than 1, disconnection wont be performed.
      - reconnect_after_x_seconds: seconds that will pass between disconnection and␣
↪new connection attempt. Only happens if disconnection has been performed.
      - dbVersion: dbVersion used in new Open and must be always equal or bigger than␣
↪LSP. If equal than LSP skip synchronization will be performed,
        if not full synchronization will be performed taking in account new starting␣
↪dbVersion desired.
 --incremental-sync-procedure <disconnect_after_x_seconds> <reconnect_after_x_seconds>
↪ <dbVersion>
   Incremental synchronization capability enabled.
      - dbVersion: dbVersion used in new Open and must be always bigger than LSP.␣
↪Incremental synchronization will be performed taking in account new starting␣
↪dbVersion desired.

 --triggered-initial-sync
   PCE-triggered synchronization capability enabled. Can be combined combined with␣
↪state-sync-avoidance/incremental-sync-procedure.

 --triggered-re-sync
   PCE-triggered re-synchronization capability enabled.
```

Troubleshooting

This section offers advices in a case OpenDaylight PCEP plugin is not working as expected.

Contents

- *PCEP is not working...*
- *Bug reporting*

PCEP is not working...

- First of all, ensure that all required features are installed, local PCE and remote PCC configuration is correct.

  To list all installed features in OpenDaylight use the following command at the Karaf console:

```
feature:list -i
```

- Check OpenDaylight Karaf logs:

  From Karaf console:

```
log:tail
```

  or open log file: `data/log/karaf.log`

  Possibly, a reason/hint for a cause of the problem can be found there.

- Try to minimize effect of other OpenDaylight features, when searching for a reason of the problem.
- Try to set DEBUG severity level for PCEP logger via Karaf console commands, in order to collect more information:

```
log:set DEBUG org.opendaylight.protocol.pcep
```

```
log:set DEBUG org.opendaylight.bgpcep.pcep
```

Bug reporting

Before you report a bug, check BGPCEP Bugzilla to ensure same/similar bug is not already filed there.

Write an e-mail to bgpcep-users@lists.opendaylight.org and provide following information:

1. State OpenDaylight version
2. Describe your use-case and provide as much details related to PCEP as possible
3. Steps to reproduce
4. Attach Karaf log files, optionally packet captures, REST input/output

References

- A Path Computation Element (PCE)-Based Architecture
- Path Computation Element (PCE) Communication Protocol Generic Requirements
- Unanswered Questions in the Path Computation Element Architecture
- A PCE-Based Architecture for Application-Based Network Operations
- Framework for PCE-Based Inter-Layer MPLS and GMPLS Traffic Engineering
- Applicability of a Stateful Path Computation Element (PCE)

PacketCable User Guide

Overview

These components introduce a DOCSIS QoS Gates management using the PCMM protocol. The driver component is responsible for the PCMM/COPS/PDP functionality required to service requests from PacketCable Provider and FlowManager. Requests are transposed into PCMM Gate Control messages and transmitted via COPS to the CMTS. This plugin adheres to the PCMM/COPS/PDP functionality defined in the CableLabs specification. PacketCable solution is an MDSAL compliant component.

PacketCable Components

PacketCable is comprised of two OpenDaylight bundles:

Bundle	Description
odl-packetcable-policy-server	Plugin that provides PCMM model implementation based on CMTS structure and COPS protocol.
odl-packetcable-policy-model	The Model provided provides a direct mapping to the underlying QoS Gates of CMTS.

See the PacketCable YANG Models.

Installing PacketCable

To install PacketCable, run the following `feature:install` command from the Karaf CLI

```
feature:install odl-packetcable-policy-server-all odl-restconf odl-mdsal-apidocs
```

Explore and exercise the PacketCable REST API

To see the PacketCable APIs, browse to this URL: http://localhost:8181/apidoc/explorer/index.html

Replace localhost with the IP address or hostname where OpenDaylight is running if you are not running OpenDaylight locally on your machine.

Note: Prior to setting any PCMM gates, a CCAP must first be added.

Postman

Install the Chrome extension

Download and import sample packetcable collection

Postman Operations

PacketCable REST API Usage Examples

- CCAP "CONFIG" DATASTORE API STRUCTURE
  - Add and view CCAPConfigDatastore(add triggers also the CCAP COPS connection):

```
PUT http://localhost:8181/restconf/config/packetcable:ccaps/ccap/CMTS-1

{"ccap":[
   {"ccapId":"CMTS-1",
    "amId": {
          "am-tag": 51930,
          "am-type": 1
    },
```

(continues on next page)

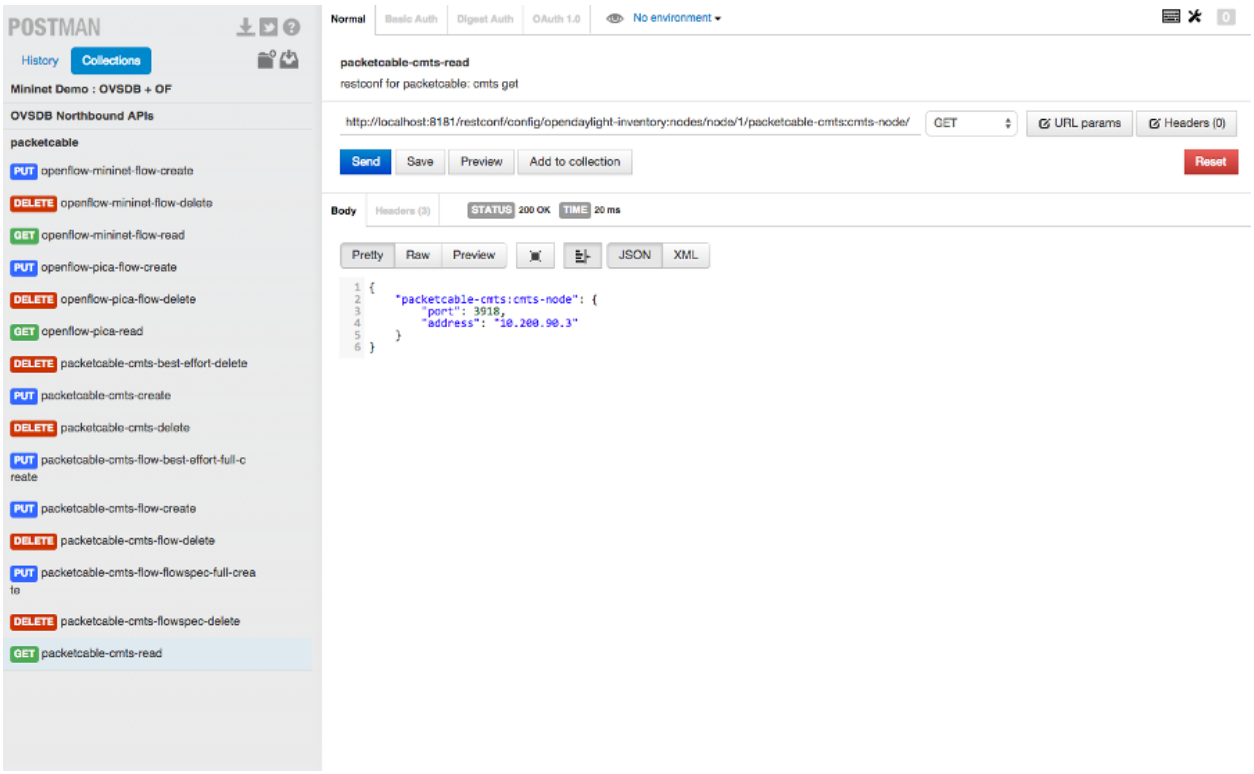

Fig. 92: Postman Operations

(continued from previous page)

```
    "connection": {
          "ipAddress": "10.20.30.40",
          "port":3918
    },"subscriber-subnets": [
          "2001:4978:030d:1000:0:0:0:0/52",
          "44.137.0.0/16"
    ],"upstream-scns": [
          "SCNA",
          "extrm_up"
    ],"downstream-scns": [
          "extrm_dn",
          "ipvideo_dn",
          "SCNC"
    ]}
]}

GET http://localhost:8181/restconf/config/packetcable:ccaps/ccap/CMTS-1
```

- CCAP OPERATIONAL STATUS - GET CCAP (COPS) CONNECTION STATUS
  - Shows the Operational Datastorecontents for the CCAP COPS connection.
  - The status is updated when the COPS connection is initiated or after an RPC poll:

```
GET http://localhost:8181/restconf/operational/packetcable:ccaps/ccap/CMTS-1/
Response: 200 OK
```

(continues on next page)

(continued from previous page)

```
{
  "ccap": [
        {
              "ccapId": "CMTS-1",
              "connection": {
                    "error": [
                          "E6-CTO: CCAP client is connected"
                    ],
                    "timestamp": "2016-03-23T14:15:54.129-05:00",
                    "connected": true
              }
        }
    ]
}
```

- CCAP OPERATIONAL STATUS - RPC CCAP POLL CONNECTION
  - A CCAP RPC poll returns the COPS connectivity status info and also triggers an Operational Datastore status update with the same data:

```
POST http://localhost:8181/restconf/operations/packetcable:ccap-poll-
↪connection
{
     "input": {
           "ccapId": "/packetcable:ccaps/packetcable:ccap[packetcable:ccapId=
↪'CMTS-1']"
     }
}
Response: 200 OK
{
"output": {
      "response": "CMTS-1: CCAP poll complete",
      "timestamp": "2016-03-23T14:15:54.131-05:00",
      "ccap": {
            "ccapId": "CMTS-1",
            "connection": {
                  "connection": {
                        "error": [
                              "CMTS-1: CCAP client is connected"
                        ],
                        "timestamp": "2016-03-23T14:15:54.129-05:00",
                        "connected": true
                  }
             }
        }
    }
}
```

- CCAP OPERATIONAL STATUS - RPC CCAP POLL CONNECTION (2) - CONNECTION DOWN:

```
POST http://localhost:8181/restconf/operations/packetcable:ccap-poll-connection
{
     "input": {
           "ccapId": "/packetcable:ccaps/packetcable:ccap[packetcable:ccapId=
↪'CMTS-1']"
     }
}
```

(continues on next page)

(continued from previous page)

```
Response: 200 OK
{
"output": {
      "response": "CMTS-1: CCAP poll complete",
      "timestamp": "2016-03-23T14:15:54.131-05:00",
      "ccap": {
            "ccapId": "CMTS-1",
            "connection": {
                  "error": [
                        "CMTS-1: CCAP client is disconnected with error: null",
                        "CMTS-1: CCAP Cops socket is closed"],
                  "timestamp": "2016-03-23T14:15:54.129-05:00",
                  "connected": false
             }
        }
    }
}
```

- CCAP OPERATIONAL STATUS - RPC CCAP SET CONNECTION
  - A CCAP RPC sets the CCAP COPS connection; possible values true or false meaning that the connection should be up or down.
  - RPC responds with the same info as RPC POLL CONNECTION, and also updates the Operational Datastore:

```
POST http://localhost:8181/restconf/operations/packetcable:ccap-set-connection
{
     "input": {
           "ccapId": "/packetcable:ccaps/packetcable:ccap[packetcable:ccapId=
↪'CMTS-1']",
            "connection": {
                  "connected": true
           }
     }
}
Response: 200 OK
{
       "output": {

              "response": "CMTS-1: CCAP set complete",
              "timestamp": "2016-03-23T17:47:29.446-05:00",
              "ccap": {
                     "ccapId": "CMTS-1",
                     "connection": {
                            "error": [
                                   "CMTS-1: CCAP client is connected",
                                   "CMTS-1: CCAP COPS socket is already open
↪"],
                            "timestamp": "2016-03-23T17:47:29.436-05:00",
                            "connected": true
                     }
              }
       }
}
```

- CCAP OPERATIONAL STATUS - RPC CCAP SET CONNECTION (2) - SHUTDOWN COPS CONNECTION:

```
POST http://localhost:8181/restconf/operations/packetcable:ccap-set-connection
{
     "input": {
           "ccapId": "/packetcable:ccaps/packetcable:ccap[packetcable:ccapId='E6-
↪CTO']",
            "connection": {
                  "connected": false
           }
     }
}
Response: 200 OK
{
       "output": {
              "response": "E6-CTO: CCAP set complete",
              "timestamp": "2016-03-23T17:47:29.446-05:00",
              "ccap": {
                     "ccapId": "E6-CTO",
                     "connection": {
                            "error": [
                                   "E60CTO: CCAP client is disconnected with␣
↪error: null"],
                            "timestamp": "2016-03-23T17:47:29.436-05:00",
                            "connected": false
                     }
              }
       }
}
```

Note: A "null" in the error information means that the CCAP connection has been disconnected as a result of a RPC SET.

- GATES "CONFIG" DATASTORE API STRUCTURE CHANGED
  - A CCAP RPC poll returns the gate status info, and also triggers a Operational Datastorestatus update.
  - At a minimum the appIdneeds to be included in the input, subscriberIdand gateIdare optional.
  - A gate status response is only included if the RPC request is done for a specific gate (subscriberIdand gateIdincluded in input).
  - Add and retrieve gates to/from the Config Datastore:

```
PUT http://localhost:8181/restconf/config/packetcable:qos/apps/app/cto-app/
↪subscribers/subscriber/44.137.0.12/gates/gate/gate88/

{
  "gate": [
    {
      "gateId": "gate88",
      "gate0spec": {
        "dscp-tos-overwrite": "0xa0",
        "dscp-tos-mask": "0xff"
      },
      "traffic-profile": {
        "service-class-name": "extrm_dn"
      },
      "classifiers": {
```

(continues on next page)

(continued from previous page)

```
        "classifier-container": [
          {
            "classifier-id": "1",
            "classifier": {
              "srcIp": "44.137.0.0",
              "dstIp": "44.137.0.11",
              "protocol": "0",
              "srcPort": "1234",
              "dstPort": "4321",
              "tos-byte": "0xa0",
              "tos-mask": "0xe0"
            }
          }
        ]
      }
    }
  ]
}

GET http://localhost:8181/restconf/config/packetcable:qos/apps/app/cto-app/
↪subscribers/subscriber/44.137.0.12/gates/gate/gate88/
```

- GATES SUPPORT MULTIPLE (UP TO FOUR) CLASSIFIERS
  - Please note that there is a classifier container now that can have up to four classifiers:

```
PUT http://localhost:8181/restconf/config/packetcable:qos/apps/app/cto-app/
↪subscribers/subscriber/44.137.0.12/gates/gate/gate88/
{ "gate":{
    "gateId": "gate44",
    "gate-spec": {
    "dscp-tos-overwrite": "0xa0",
              "dscp-tos-mask": "0xff" },
    "traffic-profile": {
              "service-class-name": "extrm_dn"},
    "classifiers":
              { "classifier-container":[
                        { "classifier-id": "1",
                                "ipv6-classifier": {
                                          "srcIp6":
↪"2001:4978:030d:1100:0:0:0:0/64",
                                                                    "dstIp6":
↪"2001:4978:030d:1000:0:0:0:0/64",
                                          "flow-label": "102",
                                          "tc-low": "0xa0",
                                          "tc-high": "0xc0",
                                          "tc-mask": "0xe0",
                                          "next-hdr": "256",
                                          "srcPort-start": "4321",
                                          "srcPort-end": "4322",
                                          "dstPort-start": "1234",
                                          "dstPort-end": "1235"
                        }},
                        { "classifier-id": "2",
                                 "ext-classifier" : {
                                           "srcIp": "44.137.0.12",
                                           "srcIpMask": "255.255.255.255",
```

(continues on next page)

(continued from previous page)

```
                                                    "dstIp": "10.10.10.0",
                                                    "dstIpMask": "255.255.255.0",
                                                    "tos-byte": "0xa0",
                                                    "tos-mask": "0xe0",
                                                    "protocol": "0",
                                                    "srcPort-start": "4321",
                                                    "srcPort-end": "4322",
                                                    "dstPort-start": "1234",
                                                    "dstPort-end": "1235"
                                            }
                                    }]
                            }
    }
}
```

- CCAP OPERATIONAL STATUS - GET GATE STATUS FROM OPERATIONAL DATASTORE
  - Shows the Operational Datastore contents for the gate.
  - The gate status is updated at the time when the gate is configured or during an RPC poll:

```
GET http://localhost:8181/restconf/operational/packetcable:qos/apps/app/cto-
↪app/subscribers/subscriber/44.137.0.12/gates/gate/gate88

Response: 200
{
    "gate":[{
            "gateId":"gate88",
            "cops-gate-usage-info": "0",
            "cops-gate-state": "Committed(4)/Other(-1)",
            "gatePath": "cto-app/44.137.0.12/gate88",
            "cops-gate-time-info": "0",
            "cops-gateId": "3e0800e9",
            "timestamp": "2016-03-24T10:30:18.763-05:00",
            "ccapId": "E6-CTO"
    }]
}
```

- CCAP OPERATIONAL STATUS - RPC GATE STATUS POLL
  - A CCAP RPC poll returns the gate status info and also triggers an Operational Datastore status update.
  - At a minimum, the appId needs to be included in the input; subscriberId and gateId are optional.
  - A gate status response is only included if the RPC request is done for a specific gate (subscriberId and gateId included in input):

```
POST http://localhost:8181/restconf/operations/packetcable:qos-poll-gates
{
     "input": {
            "appId": "/packetcable:apps/packetcable:apps[packetcable:appId=
↪'cto-app]",
            "subscriberId": "44.137.0.11",
            "gateId": "gate44"
     }
}
Response: 200 OK
{
```

(continues on next page)

(continued from previous page)


```
    "output": {
                "gate": {
                         "cops-gate-usage-info": "0",
                         "cops-gate-state": "Committed(4)/Other(-1)",
                         "gatePath": "ctoapp/44.137.0.12/gate88",
                         "cops-gate-time-info": "0",
                         "cops-gateId": "1ceb0001",
                         "error": [""],
                         "timestamp": "2016-03-24T13:22:59.900-05:00",
                         "ccapId": "E6-CTO"
                },
                "response": "cto-app/44.137.0.12/gate88: gate poll complete",
                "timestamp": "2016-03-24T13:22:59.906-05:00"
    }
}
```


– When multiple gates are polled (only appId or appId and subscriberId are provided), a generic response is returned and the Operational Datastore is updated in the background:


```
{  "output": {
       "gate": {},
       "response": "cto-app/: gate subtree poll in progress",
       "timestamp": "2016-03-24T13:25:30.471-05:00"
   }
}
```


Service Function Chaining

OpenDaylight Service Function Chaining (SFC) Overview

OpenDaylight Service Function Chaining (SFC) provides the ability to define an ordered list of a network services (e.g. firewalls, load balancers). These service are then "stitched" together in the network to create a service chain. This project provides the infrastructure (chaining logic, APIs) needed for ODL to provision a service chain in the network and an end-user application for defining such chains.

- ACE - Access Control Entry
- ACL - Access Control List
- SCF - Service Classifier Function
- SF - Service Function
- SFC - Service Function Chain
- SFF - Service Function Forwarder
- SFG - Service Function Group
- SFP - Service Function Path
- RSP - Rendered Service Path
- NSH - Network Service Header

SFC User Interface

Overview

SFC User Interface (SFC-UI) is based on Dlux project. It provides an easy way to create, read, update and delete configuration stored in Datastore. Moreover, it shows the status of all SFC features (e.g installed, uninstalled) and Karaf log messages as well.

SFC-UI Architecture

SFC-UI operates purely by using RESTCONF.

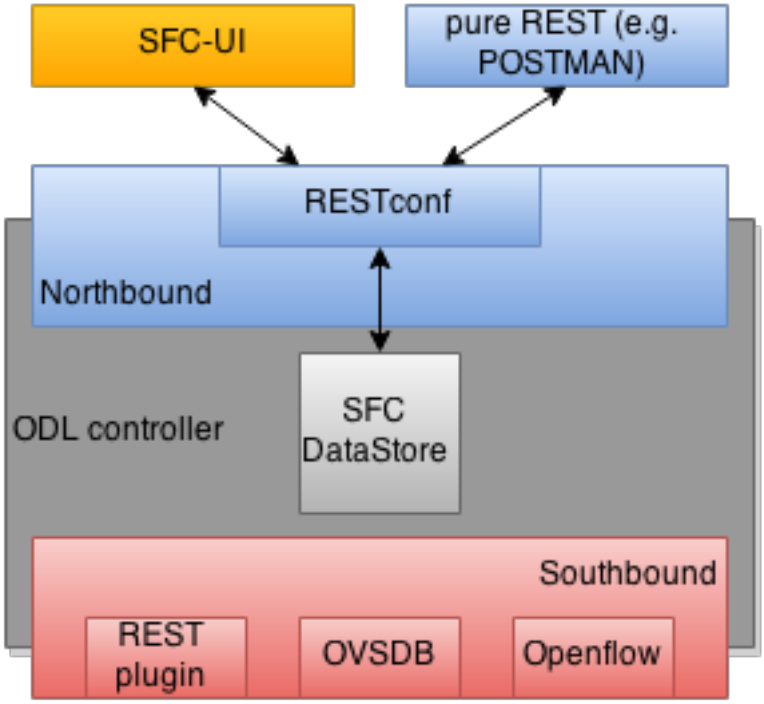

Fig. 93: SFC-UI integration into ODL

Configuring SFC-UI

1. Run ODL distribution (run karaf)
2. In karaf console execute: `feature:install odl-sfc-ui`
3. Visit SFC-UI on: `http://<odl_ip_address>:8181/sfc/index.html`

SFC Southbound REST Plugin

Overview

The Southbound REST Plugin is used to send configuration from DataStore down to network devices supporting a REST API (i.e. they have a configured REST URI). It supports POST/PUT/DELETE operations, which are triggered

accordingly by changes in the SFC data stores.

- Access Control List (ACL)
- Service Classifier Function (SCF)
- Service Function (SF)
- Service Function Group (SFG)
- Service Function Schedule Type (SFST)
- Service Function Forwader (SFF)
- Rendered Service Path (RSP)

Southbound REST Plugin Architecture

From the user perspective, the REST plugin is another SFC Southbound plugin used to communicate with network devices.

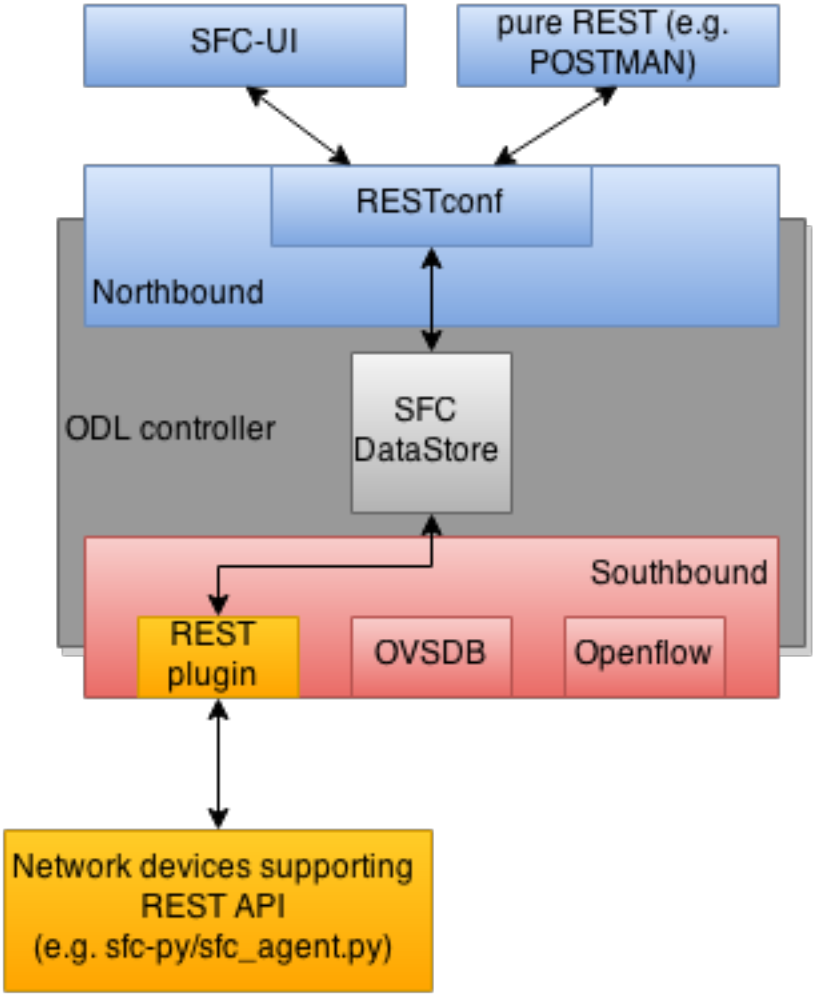

Fig. 94: Soutbound REST Plugin integration into ODL

Configuring Southbound REST Plugin

1. Run ODL distribution (run karaf)
2. In karaf console execute: `feature:install odl-sfc-sb-rest`
3. Configure REST URIs for SF/SFF through SFC User Interface or RESTCONF (required configuration steps can be found in the tutorial stated bellow)

Tutorial

Comprehensive tutorial on how to use the Southbound REST Plugin and how to control network devices with it can be found on: https://wiki.opendaylight.org/view/Service_Function_Chaining:Main#SFC_101

SFC-OVS integration

Overview

SFC-OVS provides integration of SFC with Open vSwitch (OVS) devices. Integration is realized through mapping of SFC objects (like SF, SFF, Classifier, etc.) to OVS objects (like Bridge, TerminationPoint=Port/Interface). The mapping takes care of automatic instantiation (setup) of corresponding object whenever its counterpart is created. For example, when a new SFF is created, the SFC-OVS plugin will create a new OVS bridge and when a new OVS Bridge is created, the SFC-OVS plugin will create a new SFF.

The feature is intended for SFC users willing to use Open vSwitch as underlying network infrastructure for deploying RSPs (Rendered Service Paths).

SFC-OVS Architecture

SFC-OVS uses the OVSDB MD-SAL Southbound API for getting/writing information from/to OVS devices. From the user perspective SFC-OVS acts as a layer between SFC DataStore and OVSDB.

Configuring SFC-OVS

1. Run ODL distribution (run karaf)
2. In karaf console execute: `feature:install odl-sfc-ovs`
3. Configure Open vSwitch to use ODL as a manager, using following command: `ovs-vsctl set-manager tcp:<odl_ip_address>:6640`

Tutorials

Verifying mapping from OVS to SFF

Overview

This tutorial shows the usual workflow when OVS configuration is transformed to corresponding SFC objects (in this case SFF).

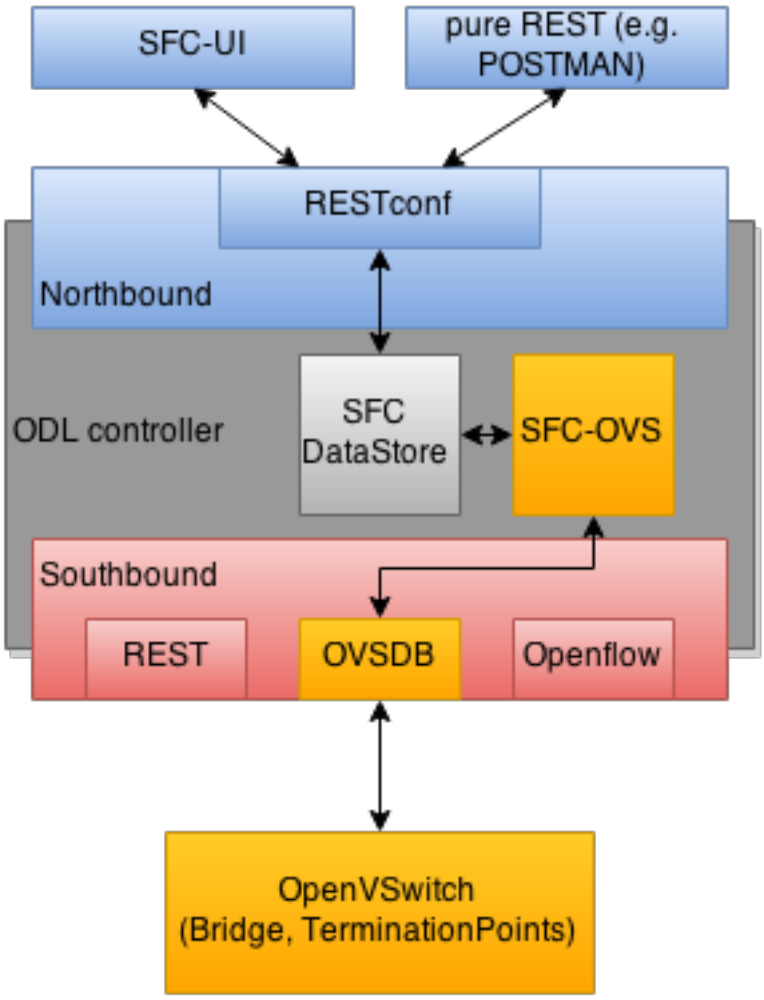

Fig. 95: SFC-OVS integration into ODL

Prerequisities

- Open vSwitch installed (ovs-vsctl command available in shell)
- SFC-OVS feature configured as stated above

Instructions

1. `ovs-vsctl set-manager tcp:<odl_ip_address>:6640`
2. `ovs-vsctl add-br br1`
3. `ovs-vsctl add-port br1 testPort`

Verification

1. visit SFC User Interface: `http://<odl_ip_address>:8181/sfc/index.html#/sfc/serviceforwarder`
2. use pure RESTCONF and send GET request to URL: `http://<odl_ip_address>:8181/restconf/config/service-function-forwarder:service-function-forwarders`

There should be SFF, which name will be ending with *br1* and the SFF should containt two DataPlane locators: *br1* and *testPort*.

Verifying mapping from SFF to OVS

Overview

This tutorial shows the usual workflow during creation of OVS Bridge with use of SFC APIs.

Prerequisities

- Open vSwitch installed (ovs-vsctl command available in shell)
- SFC-OVS feature configured as stated above

Instructions

1. In shell execute: `ovs-vsctl set-manager tcp:<odl_ip_address>:6640`
2. Send POST request to URL: `http://<odl_ip_address>:8181/restconf/operations/service-function-forwarder-ovs:create-ovs-bridge` Use Basic auth with credentials: "admin", "admin" and set `Content-Type: application/json`. The content of POST request should be following:

```
{
    "input":
    {
        "name": "br-test",
        "ovs-node": {
            "ip": "<Open_vSwitch_ip_address>"
```

(continues on next page)

(continued from previous page)

```
            }
        }
}
```

Open_vSwitch_ip_address is IP address of machine, where Open vSwitch is installed.

Verification

In shell execute: `ovs-vsctl show`. There should be Bridge with name *br-test* and one port/interface called *br-test*.

Also, corresponding SFF for this OVS Bridge should be configured, which can be verified through SFC User Interface or RESTCONF as stated in previous tutorial.

SFC Classifier User Guide

Overview

Description of classifier can be found in: https://datatracker.ietf.org/doc/draft-ietf-sfc-architecture/

There are two types of classifier:

1. OpenFlow Classifier
2. Iptables Classifier

OpenFlow Classifier

OpenFlow Classifier implements the classification criteria based on OpenFlow rules deployed into an OpenFlow switch. An Open vSwitch will take the role of a classifier and performs various encapsulations such NSH, VLAN, MPLS, etc. In the existing implementation, classifier can support NSH encapsulation. Matching information is based on ACL for MAC addresses, ports, protocol, IPv4 and IPv6. Supported protocols are TCP, UDP and SCTP. Actions information in the OF rules, shall be forwarding of the encapsulated packets with specific information related to the RSP.

Classifier Architecture

The OVSDB Southbound interface is used to create an instance of a bridge in a specific location (via IP address). This bridge contains the OpenFlow rules that perform the classification of the packets and react accordingly. The OpenFlow Southbound interface is used to translate the ACL information into OF rules within the Open vSwitch.

Note: in order to create the instance of the bridge that takes the role of a classifier, an "empty" SFF must be created.

Configuring Classifier

1. An empty SFF must be created in order to host the ACL that contains the classification information.
2. SFF data plane locator must be configured
3. Classifier interface must be mannually added to SFF bridge.

Administering or Managing Classifier

Classification information is based on MAC addresses, protocol, ports and IP. ACL gathers this information and is assigned to an RSP which turns to be a specific path for a Service Chain.

Iptables Classifier

Classifier manages everything from starting the packet listener to creation (and removal) of appropriate ip(6)tables rules and marking received packets accordingly. Its functionality is **available only on Linux** as it leverdges **NetfilterQueue**, which provides access to packets matched by an **iptables** rule. Classifier requires **root privileges** to be able to operate.

So far it is capable of processing ACL for MAC addresses, ports, IPv4 and IPv6. Supported protocols are TCP and UDP.

Classifier Architecture

Python code located in the project repository sfc-py/common/classifier.py.

Note: classifier assumes that Rendered Service Path (RSP) **already exists** in ODL when an ACL referencing it is obtained

---

1. sfc_agent receives an ACL and passes it for processing to the classifier
2. the RSP (its SFF locator) referenced by ACL is requested from ODL
3. if the RSP exists in the ODL then ACL based iptables rules for it are applied

After this process is over, every packet successfully matched to an iptables rule (i.e. successfully classified) will be NSH encapsulated and forwarded to a related SFF, which knows how to traverse the RSP.

Rules are created using appropriate iptables command. If the Access Control Entry (ACE) rule is MAC address related both iptables and ip6tabeles rules re issued. If ACE rule is IPv4 address related, only iptables rules are issued, same for IPv6.

Note: iptables **raw** table contains all created rules

Configuring Classifier

Classfier does't need any configuration.
Its only requirement is that the **second (2) Netfilter Queue** is not used by any other process and is **avalilable for the classifier**.

Administering or Managing Classifier

Classfier runs alongside sfc_agent, therefore the commad for starting it locally is:

```
sudo python3.4 sfc-py/sfc_agent.py --rest --odl-ip-port localhost:8181 --auto-sff-
↪name --nfq-class
```

SFC OpenFlow Renderer User Guide

Overview

The Service Function Chaining (SFC) OpenFlow Renderer (SFC OF Renderer) implements Service Chaining on OpenFlow switches. It listens for the creation of a Rendered Service Path (RSP), and once received it programs Service Function Forwarders (SFF) that are hosted on OpenFlow capable switches to steer packets through the service chain.

Common acronyms used in the following sections:

- SF - Service Function
- SFF - Service Function Forwarder
- SFC - Service Function Chain
- SFP - Service Function Path
- RSP - Rendered Service Path

SFC OpenFlow Renderer Architecture

The SFC OF Renderer is invoked after a RSP is created using an MD-SAL listener called `SfcOfRspDataListener`. Upon SFC OF Renderer initialization, the `SfcOfRspDataListener` registers itself to listen for RSP changes. When invoked, the `SfcOfRspDataListener` processes the RSP and calls the `SfcOfFlowProgrammerImpl` to create the necessary flows in the Service Function Forwarders configured in the RSP. Refer to the following diagram for more details.

SFC OpenFlow Switch Flow pipeline

The SFC OpenFlow Renderer uses the following tables for its Flow pipeline:

- Table 0, Classifier
- Table 1, Transport Ingress
- Table 2, Path Mapper
- Table 3, Path Mapper ACL
- Table 4, Next Hop
- Table 10, Transport Egress

The OpenFlow Table Pipeline is intended to be generic to work for all of the different encapsulations supported by SFC.

All of the tables are explained in detail in the following section.

The SFFs (SFF1 and SFF2), SFs (SF1), and topology used for the flow tables in the following sections are as described in the following diagram.

Classifier Table detailed

It is possible for the SFF to also act as a classifier. This table maps subscriber traffic to RSPs, and is explained in detail in the classifier documentation.

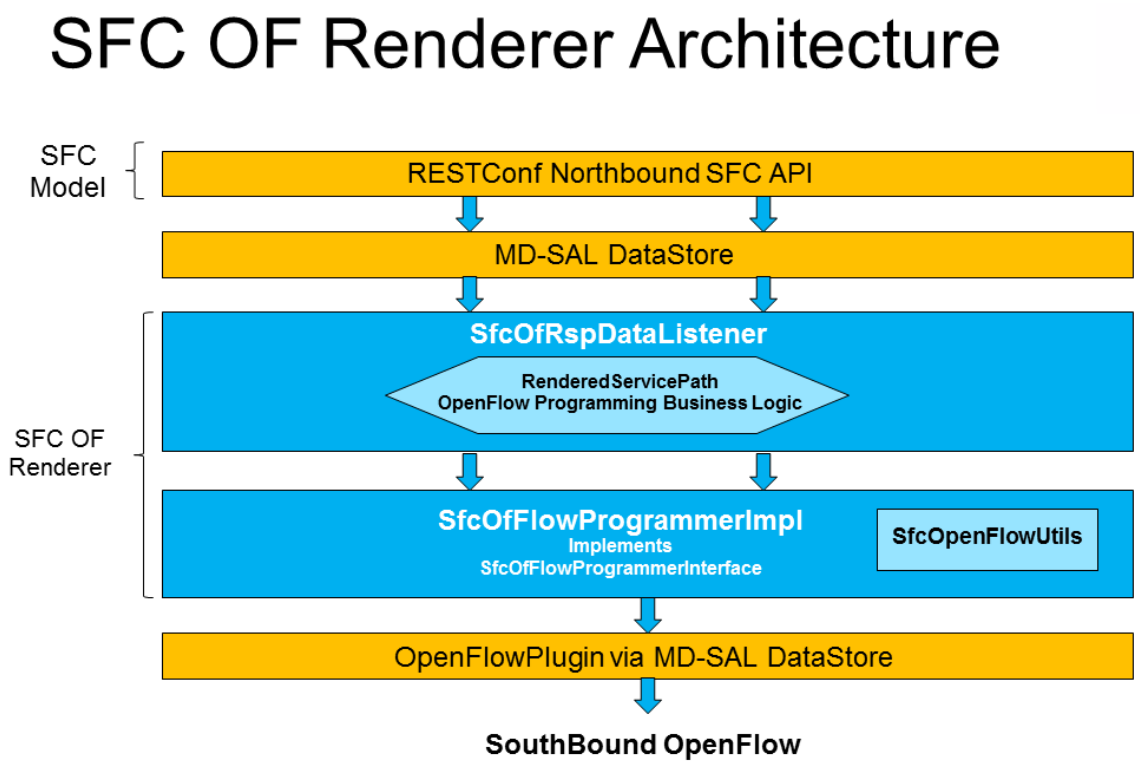

Fig. 96: SFC OpenFlow Renderer High Level Architecture

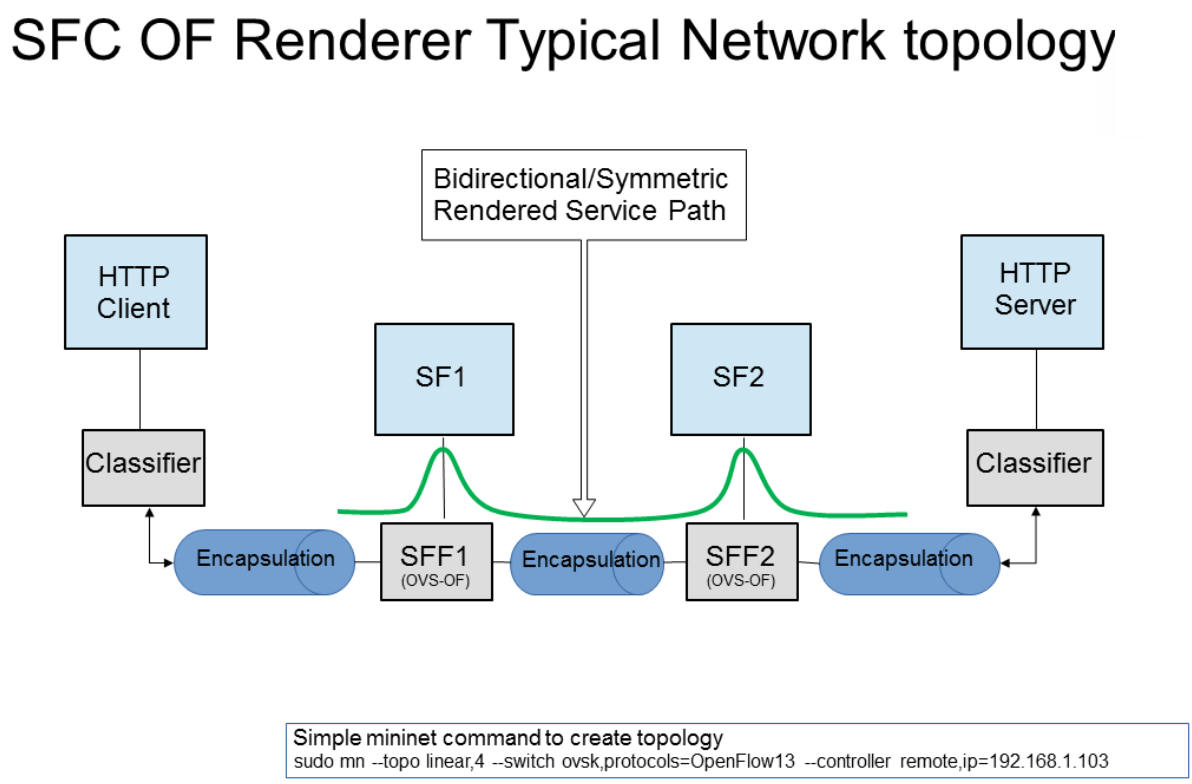

Fig. 97: SFC OpenFlow Renderer Typical Network Topology

If the SFF is not a classifier, then this table will just have a simple Goto Table 1 flow.

Transport Ingress Table detailed

The Transport Ingress table has an entry per expected tunnel transport type to be received in a particular SFF, as established in the SFC configuration.

Here are two example on SFF1: one where the RSP ingress tunnel is MPLS assuming VLAN is used for the SFF-SF, and the other where the RSP ingress tunnel is NSH GRE (UDP port 4789):

Priority	Match	Action
256	EtherType==0x8847 (MPLS unicast)	Goto Table 2
256	EtherType==0x8100 (VLAN)	Goto Table 2
256	EtherType==0x0800,udp,tp_dst==4789 (IP v4)	Goto Table 2
5	Match Any	Drop

Table: Table Transport Ingress

Path Mapper Table detailed

The Path Mapper table has an entry per expected tunnel transport info to be received in a particular SFF, as established in the SFC configuration. The tunnel transport info is used to determine the RSP Path ID, and is stored in the OpenFlow Metadata. This table is not used for NSH, since the RSP Path ID is stored in the NSH header.

For SF nodes that do not support NSH tunneling, the IP header DSCP field is used to store the RSP Path Id. The RSP Path Id is written to the DSCP field in the Transport Egress table for those packets sent to an SF.

Here is an example on SFF1, assuming the following details:

- VLAN ID 1000 is used for the SFF-SF
- The RSP Path 1 tunnel uses MPLS label 100 for ingress and 101 for egress
- The RSP Path 2 (symmetric downlink path) uses MPLS label 101 for ingress and 100 for egress

Priority	Match	Action
256	MPLS Label==100	RSP Path=1, Pop MPLS, Goto Table 4
256	MPLS Label==101	RSP Path=2, Pop MPLS, Goto Table 4
256	VLAN ID==1000, IP DSCP==1	RSP Path=1, Pop VLAN, Goto Table 4
256	VLAN ID==1000, IP DSCP==2	RSP Path=2, Pop VLAN, Goto Table 4
5	Match Any	Goto Table 3

Table: Table Path Mapper

Path Mapper ACL Table detailed

This table is only populated when PacketIn packets are received from the switch for TcpProxy type SFs. These flows are created with an inactivity timer of 60 seconds and will be automatically deleted upon expiration.

Next Hop Table detailed

The Next Hop table uses the RSP Path Id and appropriate packet fields to determine where to send the packet next. For NSH, only the NSP (Network Services Path, RSP ID) and NSI (Network Services Index, next hop) fields from the NSH header are needed to determine the VXLAN tunnel destination IP. For VLAN or MPLS, then the source MAC address is used to determine the destination MAC address.

Here are two examples on SFF1, assuming SFF1 is connected to SFF2. RSP Paths 1 and 2 are symmetric VLAN paths. RSP Paths 3 and 4 are symmetric NSH paths. RSP Path 1 ingress packets come from external to SFC, for which we don't have the source MAC address (MacSrc).

Priority	Match	Action
256	RSP Path==1, MacSrc==SF1	MacDst=SFF2, Goto Table 10
256	RSP Path==2, MacSrc==SF1	Goto Table 10
256	RSP Path==2, MacSrc==SFF2	MacDst=SF1, Goto Table 10
246	RSP Path==1	MacDst=SF1, Goto Table 10
256	nsp=3,nsi=255 (SFF Ingress RSP 3)	load:0xa000002→NXM_NX_TUN_IPV4_DST[], Goto Table 10
256	nsp=3,nsi=254 (SFF Ingress from SF, RSP 3)	load:0xa00000a→NXM_NX_TUN_IPV4_DST[], Goto Table 10
256	nsp=4,nsi=254 (SFF1 Ingress from SFF2)	load:0xa00000a→NXM_NX_TUN_IPV4_DST[], Goto Table 10
5	Match Any	Drop

Table: Table Next Hop

Transport Egress Table detailed

The Transport Egress table prepares egress tunnel information and sends the packets out.

Here are two examples on SFF1. RSP Paths 1 and 2 are symmetric MPLS paths that use VLAN for the SFF-SF. RSP Paths 3 and 4 are symmetric NSH paths. Since it is assumed that switches used for NSH will only have one VXLANport, the NSH packets are just sent back where they came from.

Priority	Match	Action
256	RSP Path==1, MacDst==SF1	Push VLAN ID 1000, Port=SF1
256	RSP Path==1, MacDst==SFF2	Push MPLS Label 101, Port=SFF2
256	RSP Path==2, MacDst==SF1	Push VLAN ID 1000, Port=SF1
246	RSP Path==2	Push MPLS Label 100, Port=Ingress
256	nsp=3,nsi=255 (SFF Ingress RSP 3)	IN_PORT
256	nsp=3,nsi=254 (SFF Ingress from SF, RSP 3)	IN_PORT
256	nsp=4,nsi=254 (SFF1 Ingress from SFF2)	IN_PORT
5	Match Any	Drop

Table: Table Transport Egress

Administering SFC OF Renderer

To use the SFC OpenFlow Renderer Karaf, at least the following Karaf features must be installed.

- odl-openflowplugin-nxm-extensions

- odl-openflowplugin-flow-services
- odl-sfc-provider
- odl-sfc-model
- odl-sfc-openflow-renderer
- odl-sfc-ui (optional)

The following command can be used to view all of the currently installed Karaf features:

```
opendaylight-user@root>feature:list -i
```

Or, pipe the command to a grep to see a subset of the currently installed Karaf features:

```
opendaylight-user@root>feature:list -i | grep sfc
```

To install a particular feature, use the Karaf `feature:install` command.

SFC OF Renderer Tutorial

Overview

In this tutorial, 2 different encapsulations will be shown: MPLS and NSH. The following Network Topology diagram is a logical view of the SFFs and SFs involved in creating the Service Chains.

Prerequisites

To use this example, SFF OpenFlow switches must be created and connected as illustrated above. Additionally, the SFs must be created and connected.

Target Environment

The target environment is not important, but this use-case was created and tested on Linux.

Instructions

The steps to use this tutorial are as follows. The referenced configuration in the steps is listed in the following sections.

There are numerous ways to send the configuration. In the following configuration chapters, the appropriate `curl` command is shown for each configuration to be sent, including the URL.

Steps to configure the SFC OF Renderer tutorial:

1. Send the `SF` RESTCONF configuration
2. Send the `SFF` RESTCONF configuration
3. Send the `SFC` RESTCONF configuration
4. Send the `SFP` RESTCONF configuration
5. Create the `RSP` with a RESTCONF RPC command

Once the configuration has been successfully created, query the Rendered Service Paths with either the SFC UI or via RESTCONF. Notice that the RSP is symmetrical, so the following 2 RSPs will be created:

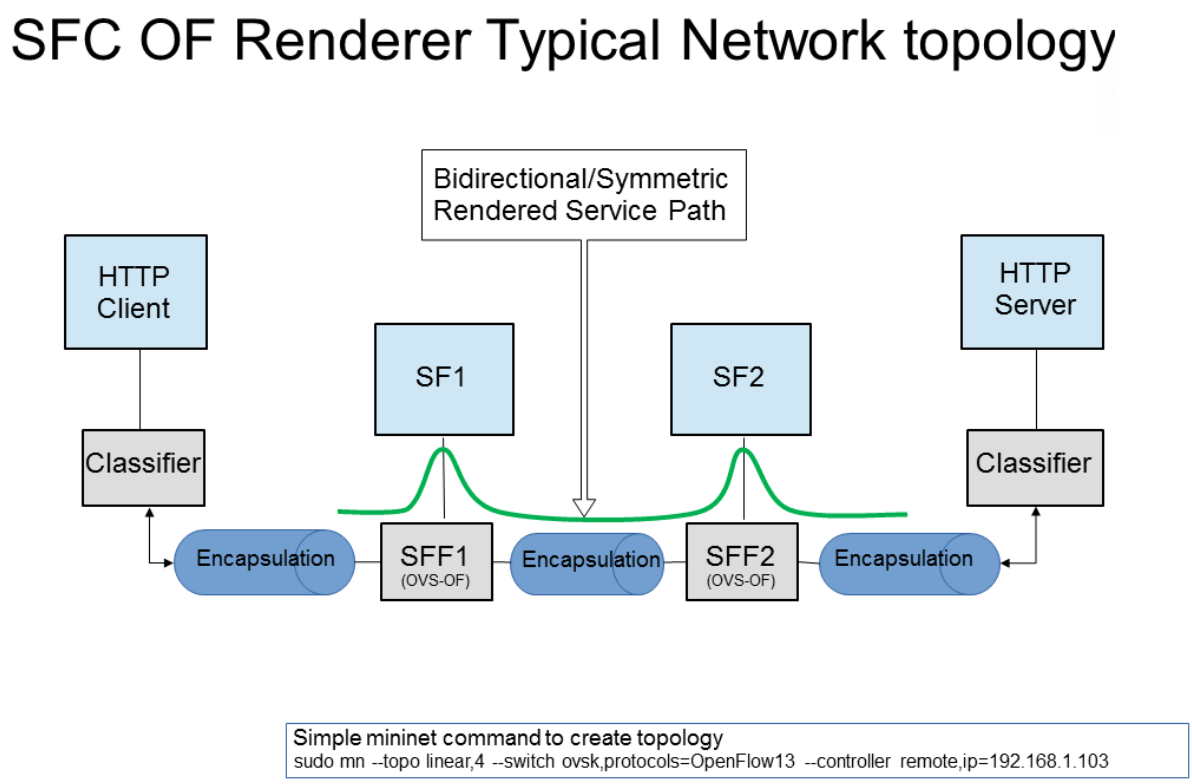

Fig. 98: SFC OpenFlow Renderer Typical Network Topology

- sfc-path1
- sfc-path1-Reverse

At this point the Service Chains have been created, and the OpenFlow Switches are programmed to steer traffic through the Service Chain. Traffic can now be injected from a client into the Service Chain. To debug problems, the OpenFlow tables can be dumped with the following commands, assuming SFF1 is called `s1` and SFF2 is called `s2`.

```
sudo ovs-ofctl -O OpenFlow13  dump-flows s1
```

```
sudo ovs-ofctl -O OpenFlow13  dump-flows s2
```

In all the following configuration sections, replace the `${JSON}` string with the appropriate JSON configuration. Also, change the `localhost` desintation in the URL accordingly.

SFC OF Renderer NSH Tutorial

The following configuration sections show how to create the different elements using NSH encapsulation.

NSH Service Function configuration

The Service Function configuration can be sent with the following command:

```
curl -i -H "Content-Type: application/json" -H "Cache-Control: no-cache" --data '$
↪{JSON}' -X PUT --user admin:admin http://localhost:8181/restconf/config/service-
↪function:service-functions/
```

SF configuration JSON.

```
{
 "service-functions": {
   "service-function": [
     {
       "name": "sf1",
       "type": "http-header-enrichment",
       "nsh-aware": true,
       "ip-mgmt-address": "10.0.0.2",
       "sf-data-plane-locator": [
         {
           "name": "sf1dpl",
           "ip": "10.0.0.10",
           "port": 4789,
           "transport": "service-locator:vxlan-gpe",
           "service-function-forwarder": "sff1"
         }
       ]
     },
     {
       "name": "sf2",
       "type": "firewall",
       "nsh-aware": true,
       "ip-mgmt-address": "10.0.0.3",
       "sf-data-plane-locator": [
         {
```

(continues on next page)

(continued from previous page)

```
            "name": "sf2dpl",
             "ip": "10.0.0.20",
             "port": 4789,
             "transport": "service-locator:vxlan-gpe",
            "service-function-forwarder": "sff2"
          }
        ]
      }
    ]
  }
}
```

NSH Service Function Forwarder configuration

The Service Function Forwarder configuration can be sent with the following command:

```
curl -i -H "Content-Type: application/json" -H "Cache-Control: no-cache" --data '$
↪{JSON}' -X PUT --user admin:admin http://localhost:8181/restconf/config/service-
↪function-forwarder:service-function-forwarders/
```

SFF configuration JSON.

```
{
 "service-function-forwarders": {
   "service-function-forwarder": [
     {
       "name": "sff1",
       "service-node": "openflow:2",
       "sff-data-plane-locator": [
         {
           "name": "sff1dpl",
           "data-plane-locator":
           {
               "ip": "10.0.0.1",
               "port": 4789,
               "transport": "service-locator:vxlan-gpe"
           }
         }
       ],
       "service-function-dictionary": [
         {
           "name": "sf1",
           "sff-sf-data-plane-locator":
           {
               "sf-dpl-name": "sf1dpl",
               "sff-dpl-name": "sff1dpl"
           }
         }
       ]
     },
     {
       "name": "sff2",
       "service-node": "openflow:3",
       "sff-data-plane-locator": [
```

(continues on next page)

(continued from previous page)

```
          {
            "name": "sff2dpl",
            "data-plane-locator":
            {
                "ip": "10.0.0.2",
                "port": 4789,
                "transport": "service-locator:vxlan-gpe"
            }
          }
        ],
        "service-function-dictionary": [
          {
            "name": "sf2",
            "sff-sf-data-plane-locator":
            {
                "sf-dpl-name": "sf2dpl",
                "sff-dpl-name": "sff2dpl"
            }
          }
        ]
      }
    ]
  }
}
```

NSH Service Function Chain configuration

The Service Function Chain configuration can be sent with the following command:

```
curl -i -H "Content-Type: application/json" -H "Cache-Control: no-cache" --data '$
↪{JSON}' -X PUT --user admin:admin http://localhost:8181/restconf/config/service-
↪function-chain:service-function-chains/
```

SFC configuration JSON.

```
{
 "service-function-chains": {
   "service-function-chain": [
     {
       "name": "sfc-chain1",
       "symmetric": true,
       "sfc-service-function": [
         {
           "name": "hdr-enrich-abstract1",
           "type": "http-header-enrichment"
         },
         {
           "name": "firewall-abstract1",
           "type": "firewall"
         }
       ]
     }
   ]
```

(continues on next page)

(continued from previous page)

```
 }
}
```

NSH Service Function Path configuration

The Service Function Path configuration can be sent with the following command:

```
curl -i -H "Content-Type: application/json" -H "Cache-Control: no-cache" --data '$
↪{JSON}' -X PUT --user admin:admin http://localhost:8181/restconf/config/service-
↪function-path:service-function-paths/
```

SFP configuration JSON.

```
{
  "service-function-paths": {
    "service-function-path": [
      {
        "name": "sfc-path1",
        "service-chain-name": "sfc-chain1",
        "transport-type": "service-locator:vxlan-gpe",
        "symmetric": true
      }
    ]
  }
}
```

NSH Rendered Service Path creation

```
curl -i -H "Content-Type: application/json" -H "Cache-Control: no-cache" --data '$
↪{JSON}' -X POST --user admin:admin http://localhost:8181/restconf/operations/
↪rendered-service-path:create-rendered-path/
```

RSP creation JSON.

```
{
 "input": {
     "name": "sfc-path1",
     "parent-service-function-path": "sfc-path1",
     "symmetric": true
 }
}
```

NSH Rendered Service Path removal

The following command can be used to remove a Rendered Service Path called `sfc-path1`:

```
curl -i -H "Content-Type: application/json" -H "Cache-Control: no-cache" --data '{
↪"input": {"name": "sfc-path1" } }' -X POST --user admin:admin http://localhost:8181/
↪restconf/operations/rendered-service-path:delete-rendered-path/
```

(continues on next page)

(continued from previous page)

NSH Rendered Service Path Query

The following command can be used to query all of the created Rendered Service Paths:

```
curl -H "Content-Type: application/json" -H "Cache-Control: no-cache" -X GET --user␣
↪admin:admin http://localhost:8181/restconf/operational/rendered-service-
↪path:rendered-service-paths/
```

SFC OF Renderer MPLS Tutorial

The following configuration sections show how to create the different elements using MPLS encapsulation.

MPLS Service Function configuration

The Service Function configuration can be sent with the following command:

```
curl -i -H "Content-Type: application/json" -H "Cache-Control: no-cache" --data '$
↪{JSON}' -X PUT --user admin:admin http://localhost:8181/restconf/config/service-
↪function:service-functions/
```

SF configuration JSON.

```
{
 "service-functions": {
   "service-function": [
     {
       "name": "sf1",
       "type": "http-header-enrichment",
       "nsh-aware": false,
       "ip-mgmt-address": "10.0.0.2",
       "sf-data-plane-locator": [
         {
           "name": "sf1-sff1",
           "mac": "00:00:08:01:02:01",
           "vlan-id": 1000,
           "transport": "service-locator:mac",
           "service-function-forwarder": "sff1"
         }
       ]
     },
     {
       "name": "sf2",
       "type": "firewall",
       "nsh-aware": false,
       "ip-mgmt-address": "10.0.0.3",
       "sf-data-plane-locator": [
         {
           "name": "sf2-sff2",
```

(continues on next page)

(continued from previous page)

```
            "mac": "00:00:08:01:03:01",
            "vlan-id": 2000,
            "transport": "service-locator:mac",
            "service-function-forwarder": "sff2"
          }
        ]
      }
    ]
  }
}
```

MPLS Service Function Forwarder configuration

The Service Function Forwarder configuration can be sent with the following command:

```
curl -i -H "Content-Type: application/json" -H "Cache-Control: no-cache" --data '$
↪{JSON}' -X PUT --user admin:admin http://localhost:8181/restconf/config/service-
↪function-forwarder:service-function-forwarders/
```

SFF configuration JSON.

```
{
 "service-function-forwarders": {
   "service-function-forwarder": [
     {
       "name": "sff1",
       "service-node": "openflow:2",
       "sff-data-plane-locator": [
         {
           "name": "ulSff1Ingress",
           "data-plane-locator":
           {
               "mpls-label": 100,
               "transport": "service-locator:mpls"
           },
           "service-function-forwarder-ofs:ofs-port":
           {
               "mac": "11:11:11:11:11:11",
               "port-id" : "1"
           }
         },
         {
           "name": "ulSff1ToSff2",
           "data-plane-locator":
           {
               "mpls-label": 101,
               "transport": "service-locator:mpls"
           },
           "service-function-forwarder-ofs:ofs-port":
           {
               "mac": "33:33:33:33:33:33",
               "port-id" : "2"
           }
         },
```

(continues on next page)

(continued from previous page)

```
        {
          "name": "toSf1",
          "data-plane-locator":
          {
              "mac": "22:22:22:22:22:22",
              "vlan-id": 1000,
              "transport": "service-locator:mac",
          },
          "service-function-forwarder-ofs:ofs-port":
          {
              "mac": "33:33:33:33:33:33",
              "port-id" : "3"
          }
        }
      ],
      "service-function-dictionary": [
        {
          "name": "sf1",
          "sff-sf-data-plane-locator":
          {
              "sf-dpl-name": "sf1-sff1",
              "sff-dpl-name": "toSf1"
          }
        }
      ]
    },
    {
      "name": "sff2",
      "service-node": "openflow:3",
      "sff-data-plane-locator": [
        {
          "name": "ulSff2Ingress",
          "data-plane-locator":
          {
              "mpls-label": 101,
              "transport": "service-locator:mpls"
          },
          "service-function-forwarder-ofs:ofs-port":
          {
              "mac": "44:44:44:44:44:44",
              "port-id" : "1"
          }
        },
        {
          "name": "ulSff2Egress",
          "data-plane-locator":
          {
              "mpls-label": 102,
              "transport": "service-locator:mpls"
          },
          "service-function-forwarder-ofs:ofs-port":
          {
              "mac": "66:66:66:66:66:66",
              "port-id" : "2"
          }
        },
        {
```

(continues on next page)

(continued from previous page)

```
            "name": "toSf2",
            "data-plane-locator":
            {
                "mac": "55:55:55:55:55:55",
                "vlan-id": 2000,
                "transport": "service-locator:mac"
            },
            "service-function-forwarder-ofs:ofs-port":
            {
                "port-id" : "3"
            }
          }
        ],
        "service-function-dictionary": [
          {
            "name": "sf2",
            "sff-sf-data-plane-locator":
            {
                "sf-dpl-name": "sf2-sff2",
                "sff-dpl-name": "toSf2"

            },
            "service-function-forwarder-ofs:ofs-port":
            {
                "port-id" : "3"
            }
          }
        ]
      }
    ]
 }
}
```

MPLS Service Function Chain configuration

The Service Function Chain configuration can be sent with the following command:

```
curl -i -H "Content-Type: application/json" -H "Cache-Control: no-cache" --data '$
↪{JSON}' -X PUT --user admin:admin http://localhost:8181/restconf/config/service-
↪function-chain:service-function-chains/
```

SFC configuration JSON.

```
{
 "service-function-chains": {
   "service-function-chain": [
     {
       "name": "sfc-chain1",
       "symmetric": true,
       "sfc-service-function": [
         {
           "name": "hdr-enrich-abstract1",
           "type": "http-header-enrichment"
         },
```

(continues on next page)

(continued from previous page)

```
          {
            "name": "firewall-abstract1",
            "type": "firewall"
          }
        ]
      }
    ]
  }
}
```

MPLS Service Function Path configuration

The Service Function Path configuration can be sent with the following command:

```
curl -i -H "Content-Type: application/json" -H "Cache-Control: no-cache" --data '$
↪{JSON}' -X PUT --user admin:admin http://localhost:8181/restconf/config/service-
↪function-path:service-function-paths/
```

SFP configuration JSON.

```
{
  "service-function-paths": {
    "service-function-path": [
      {
        "name": "sfc-path1",
        "service-chain-name": "sfc-chain1",
        "transport-type": "service-locator:mpls",
        "symmetric": true
      }
    ]
  }
}
```

MPLS Rendered Service Path creation

```
curl -i -H "Content-Type: application/json" -H "Cache-Control: no-cache" --data '$
↪{JSON}' -X POST --user admin:admin http://localhost:8181/restconf/operations/
↪rendered-service-path:create-rendered-path/
```

RSP creation JSON.

```
{
 "input": {
     "name": "sfc-path1",
     "parent-service-function-path": "sfc-path1",
     "symmetric": true
 }
}
```

MPLS Rendered Service Path removal

The following command can be used to remove a Rendered Service Path called `sfc-path1`:

```
curl -i -H "Content-Type: application/json" -H "Cache-Control: no-cache" --data '{
↪"input": {"name": "sfc-path1" } }' -X POST --user admin:admin http://localhost:8181/
↪restconf/operations/rendered-service-path:delete-rendered-path/
```

MPLS Rendered Service Path Query

The following command can be used to query all of the created Rendered Service Paths:

```
curl -H "Content-Type: application/json" -H "Cache-Control: no-cache" -X GET --user␣
↪admin:admin http://localhost:8181/restconf/operational/rendered-service-
↪path:rendered-service-paths/
```

SFC IOS XE Renderer User Guide

Overview

The early Service Function Chaining (SFC) renderer for IOS-XE devices (SFC IOS-XE renderer) implements Service Chaining functionality on IOS-XE capable switches. It listens for the creation of a Rendered Service Path (RSP) and sets up Service Function Forwarders (SFF) that are hosted on IOS-XE switches to steer traffic through the service chain.

Common acronyms used in the following sections:

- SF - Service Function
- SFF - Service Function Forwarder
- SFC - Service Function Chain
- SP - Service Path
- SFP - Service Function Path
- RSP - Rendered Service Path
- LSF - Local Service Forwarder
- RSF - Remote Service Forwarder

SFC IOS-XE Renderer Architecture

When the SFC IOS-XE renderer is initialized, all required listeners are registered to handle incoming data. It involves CSR/IOS-XE `NodeListener` which stores data about all configurable devices including their mountpoints (used here as databrokers), `ServiceFunctionListener`, `ServiceForwarderListener` (see mapping) and `RenderedPathListener` used to listen for RSP changes. When the SFC IOS-XE renderer is invoked, `RenderedPathListener` calls the `IosXeRspProcessor` which processes the RSP change and creates all necessary Service Paths and Remote Service Forwarders (if necessary) on IOS-XE devices.

Service Path details

Each Service Path is defined by index (represented by NSP) and contains service path entries. Each entry has appropriate service index (NSI) and definition of next hop. Next hop can be Service Function, different Service Function Forwarder or definition of end of chain - terminate. After terminating, the packet is sent to destination. If a SFF is defined as a next hop, it has to be present on device in the form of Remote Service Forwarder. RSFs are also created during RSP processing.

Example of Service Path:

```
service-chain service-path 200
   service-index 255 service-function firewall-1
   service-index 254 service-function dpi-1
   service-index 253 terminate
```

Mapping to IOS-XE SFC entities

Renderer contains mappers for SFs and SFFs. IOS-XE capable device is using its own definition of Service Functions and Service Function Forwarders according to appropriate .yang file. `ServiceFunctionListener` serves as a listener for SF changes. If SF appears in datastore, listener extracts its management ip address and looks into cached IOS-XE nodes. If some of available nodes match, Service function is mapped in `IosXeServiceFunctionMapper` to be understandable by IOS-XE device and it's written into device's config. `ServiceForwarderListener` is used in a similar way. All SFFs with suitable management ip address it mapped in `IosXeServiceForwarderMapper`. Remapped SFFs are configured as a Local Service Forwarders. It is not possible to directly create Remote Service Forwarder using IOS-XE renderer. RSF is created only during RSP processing.

Administering SFC IOS-XE renderer

To use the SFC IOS-XE Renderer Karaf, at least the following Karaf features must be installed:

- odl-aaa-shiro
- odl-sfc-model
- odl-sfc-provider
- odl-restconf
- odl-netconf-topology
- odl-sfc-ios-xe-renderer

SFC IOS-XE renderer Tutorial

Overview

This tutorial is a simple example how to create Service Path on IOS-XE capable device using IOS-XE renderer

Preconditions

To connect to IOS-XE device, it is necessary to use several modified yang models and override device's ones. All .yang files are in the `Yang/netconf` folder in the `sfc-ios-xe-renderer module` in the SFC project. These

files have to be copied to the `cache/schema` directory, before Karaf is started. After that, custom capabilities have to be sent to network-topology:

```
PUT ./config/network-topology:network-topology/topology/topology-netconf/node/<device-
↪name>

<node xmlns="urn:TBD:params:xml:ns:yang:network-topology">
  <node-id>device-name</node-id>
  <host xmlns="urn:opendaylight:netconf-node-topology">device-ip</host>
  <port xmlns="urn:opendaylight:netconf-node-topology">2022</port>
  <username xmlns="urn:opendaylight:netconf-node-topology">login</username>
  <password xmlns="urn:opendaylight:netconf-node-topology">password</password>
  <tcp-only xmlns="urn:opendaylight:netconf-node-topology">false</tcp-only>
  <keepalive-delay xmlns="urn:opendaylight:netconf-node-topology">0</keepalive-delay>
  <yang-module-capabilities xmlns="urn:opendaylight:netconf-node-topology">
     <override>true</override>
     <capability xmlns="urn:opendaylight:netconf-node-topology">
        urn:ietf:params:xml:ns:yang:ietf-inet-types?module=ietf-inet-types&amp;
↪revision=2013-07-15
     </capability>
     <capability xmlns="urn:opendaylight:netconf-node-topology">
        urn:ietf:params:xml:ns:yang:ietf-yang-types?module=ietf-yang-types&amp;
↪revision=2013-07-15
     </capability>
     <capability xmlns="urn:opendaylight:netconf-node-topology">
        urn:ios?module=ned&amp;revision=2016-03-08
     </capability>
     <capability xmlns="urn:opendaylight:netconf-node-topology">
        http://tail-f.com/yang/common?module=tailf-common&amp;revision=2015-05-22
     </capability>
     <capability xmlns="urn:opendaylight:netconf-node-topology">
        http://tail-f.com/yang/common?module=tailf-meta-extensions&amp;revision=2013-
↪11-07
     </capability>
     <capability xmlns="urn:opendaylight:netconf-node-topology">
        http://tail-f.com/yang/common?module=tailf-cli-extensions&amp;revision=2015-
↪03-19
     </capability>
  </yang-module-capabilities>
</node>
```

Note: The device name in the URL and in the XML must match.

Instructions

When the IOS-XE renderer is installed, all NETCONF nodes in topology-netconf are processed and all capable nodes with accessible mountpoints are cached. The first step is to create LSF on node.

`Service Function Forwarder configuration`

```
PUT ./config/service-function-forwarder:service-function-forwarders

{
    "service-function-forwarders": {
        "service-function-forwarder": [
```

(continues on next page)

(continued from previous page)

```
            {
                "name": "CSR1Kv-2",
                "ip-mgmt-address": "172.25.73.23",
                "sff-data-plane-locator": [
                    {
                        "name": "CSR1Kv-2-dpl",
                        "data-plane-locator": {
                            "transport": "service-locator:vxlan-gpe",
                            "port": 6633,
                            "ip": "10.99.150.10"
                        }
                    }
                ]
            }
        ]
    }
}
```

If the IOS-XE node with appropriate management IP exists, this configuration is mapped and LSF is created on the device. The same approach is used for Service Functions.

```
PUT ./config/service-function:service-functions

{
    "service-functions": {
        "service-function": [
            {
                "name": "Firewall",
                "ip-mgmt-address": "172.25.73.23",
                "type": "service-function-type: firewall",
                "nsh-aware": true,
                "sf-data-plane-locator": [
                    {
                        "name": "firewall-dpl",
                        "port": 6633,
                        "ip": "12.1.1.2",
                        "transport": "service-locator:gre",
                        "service-function-forwarder": "CSR1Kv-2"
                    }
                ]
            },
            {
                "name": "Dpi",
                "ip-mgmt-address": "172.25.73.23",
                "type":"service-function-type: dpi",
                "nsh-aware": true,
                "sf-data-plane-locator": [
                    {
                        "name": "dpi-dpl",
                        "port": 6633,
                        "ip": "12.1.1.1",
                        "transport": "service-locator:gre",
                        "service-function-forwarder": "CSR1Kv-2"
                    }
                ]
            },
```

(continues on next page)

(continued from previous page)

```
            {
                "name": "Qos",
                "ip-mgmt-address": "172.25.73.23",
                "type":"service-function-type: qos",
                "nsh-aware": true,
                "sf-data-plane-locator": [
                    {
                        "name": "qos-dpl",
                        "port": 6633,
                        "ip": "12.1.1.4",
                        "transport": "service-locator:gre",
                        "service-function-forwarder": "CSR1Kv-2"
                    }
                ]
            }
        ]
    }
}
```

All these SFs are configured on the same device as the LSF. The next step is to prepare Service Function Chain. SFC is symmetric.

```
PUT ./config/service-function-chain:service-function-chains/

{
    "service-function-chains": {
        "service-function-chain": [
            {
                "name": "CSR3XSF",
                "symmetric": "true",
                "sfc-service-function": [
                    {
                        "name": "Firewall",
                        "type": "service-function-type: firewall"
                    },
                    {
                        "name": "Dpi",
                        "type": "service-function-type: dpi"
                    },
                    {
                        "name": "Qos",
                        "type": "service-function-type: qos"
                    }
                ]
            }
        ]
    }
}
```

Service Function Path:

```
PUT ./config/service-function-path:service-function-paths/

{
    "service-function-paths": {
        "service-function-path": [
```

(continues on next page)

(continued from previous page)

```
            {
                "name": "CSR3XSF-Path",
                "service-chain-name": "CSR3XSF",
                "starting-index": 255,
                "symmetric": "true"
            }
        ]
    }
}
```

Without a classifier, there is possibility to POST RSP directly.

```
POST ./operations/rendered-service-path:create-rendered-path

{
  "input": {
      "name": "CSR3XSF-Path-RSP",
      "parent-service-function-path": "CSR3XSF-Path",
      "symmetric": true
  }
}
```

The resulting configuration:

```
!
service-chain service-function-forwarder local
  ip address 10.99.150.10
!
service-chain service-function firewall
ip address 12.1.1.2
  encapsulation gre enhanced divert
!
service-chain service-function dpi
ip address 12.1.1.1
  encapsulation gre enhanced divert
!
service-chain service-function qos
ip address 12.1.1.4
  encapsulation gre enhanced divert
!
service-chain service-path 1
  service-index 255 service-function firewall
  service-index 254 service-function dpi
  service-index 253 service-function qos
  service-index 252 terminate
!
service-chain service-path 2
  service-index 255 service-function qos
  service-index 254 service-function dpi
  service-index 253 service-function firewall
  service-index 252 terminate
!
```

Service Path 1 is direct, Service Path 2 is reversed. Path numbers may vary.

Service Function Scheduling Algorithms

Overview

When creating the Rendered Service Path, the origin SFC controller chose the first available service function from a list of service function names. This may result in many issues such as overloaded service functions and a longer service path as SFC has no means to understand the status of service functions and network topology. The service function selection framework supports at least four algorithms (Random, Round Robin, Load Balancing and Shortest Path) to select the most appropriate service function when instantiating the Rendered Service Path. In addition, it is an extensible framework that allows 3rd party selection algorithm to be plugged in.

Architecture

The following figure illustrates the service function selection framework and algorithms.

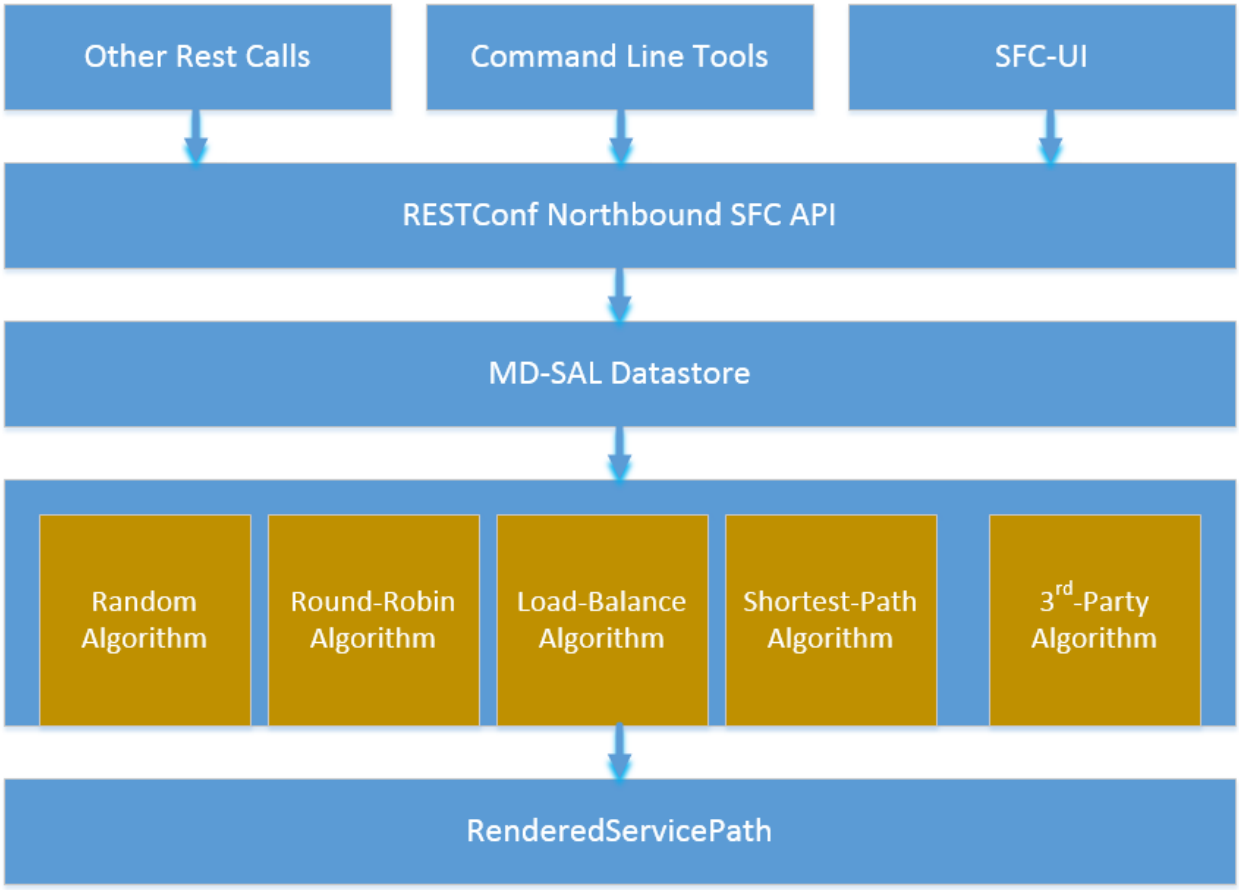

Fig. 99: SF Selection Architecture

A user has three different ways to select one service function selection algorithm:

1. Integrated RESTCONF Calls. OpenStack and/or other administration system could provide plugins to call the APIs to select one scheduling algorithm.
2. Command line tools. Command line tools such as curl or browser plugins such as POSTMAN (for Google Chrome) and RESTClient (for Mozilla Firefox) could select schedule algorithm by making RESTCONF calls.

3. SFC-UI. Now the SFC-UI provides an option for choosing a selection algorithm when creating a Rendered Service Path.

The RESTCONF northbound SFC API provides GUI/RESTCONF interactions for choosing the service function selection algorithm. MD-SAL data store provides all supported service function selection algorithms, and provides APIs to enable one of the provided service function selection algorithms. Once a service function selection algorithm is enabled, the service function selection algorithm will work when creating a Rendered Service Path.

Select SFs with Scheduler

Administrator could use both the following ways to select one of the selection algorithm when creating a Rendered Service Path.

- Command line tools. Command line tools includes Linux commands curl or even browser plugins such as POSTMAN(for Google Chrome) or RESTClient(for Mozilla Firefox). In this case, the following JSON content is needed at the moment: Service_function_schudule_type.json

```
{
  "service-function-scheduler-types": {
    "service-function-scheduler-type": [
      {
        "name": "random",
        "type": "service-function-scheduler-type:random",
        "enabled": false
      },
      {
        "name": "roundrobin",
        "type": "service-function-scheduler-type:round-robin",
        "enabled": true
      },
      {
        "name": "loadbalance",
        "type": "service-function-scheduler-type:load-balance",
        "enabled": false
      },
      {
        "name": "shortestpath",
        "type": "service-function-scheduler-type:shortest-path",
        "enabled": false
      }
    ]
  }
}
```

If using the Linux curl command, it could be:

```
curl -i -H "Content-Type: application/json" -H "Cache-Control: no-cache" --data '$
↪${Service_function_schudule_type.json}'
-X PUT --user admin:admin http://localhost:8181/restconf/config/service-function-
↪scheduler-type:service-function-scheduler-types/
```

Here is also a snapshot for using the RESTClient plugin:

- SFC-UI.SFC-UI provides a drop down menu for service function selection algorithm. Here is a snapshot for the user interaction from SFC-UI when creating a Rendered Service Path.

Note: Some service function selection algorithms in the drop list are not implemented yet. Only the first three

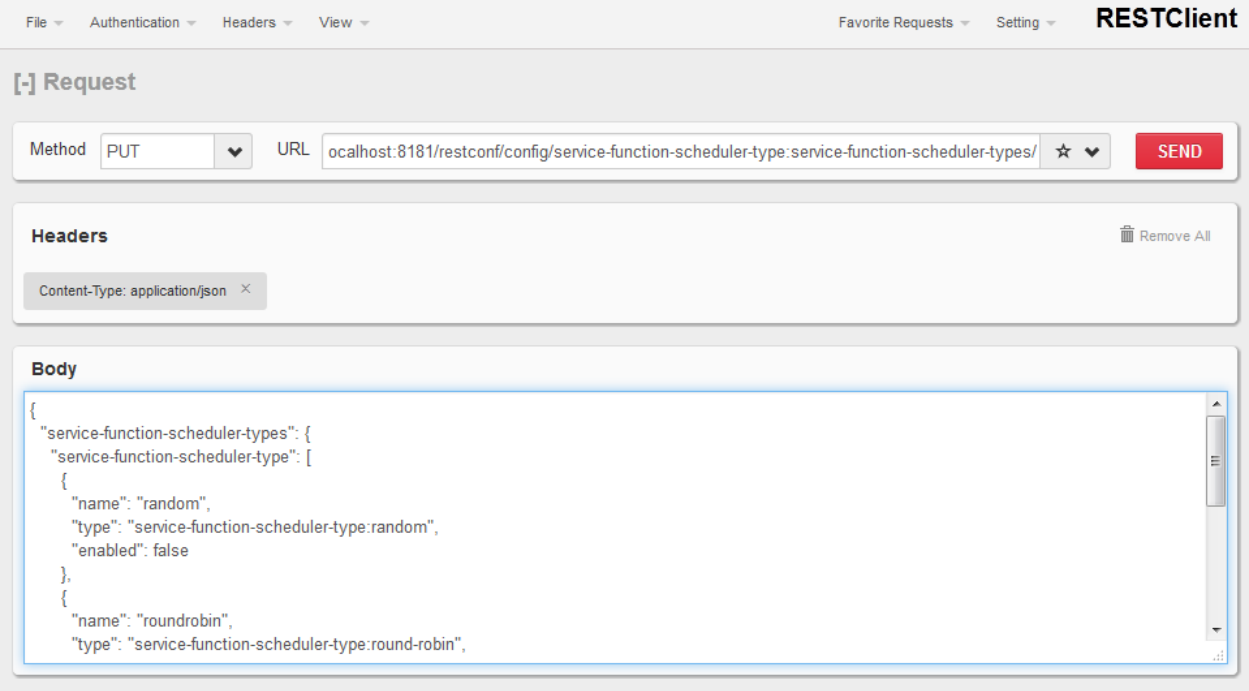

Fig. 100: Mozilla Firefox RESTClient

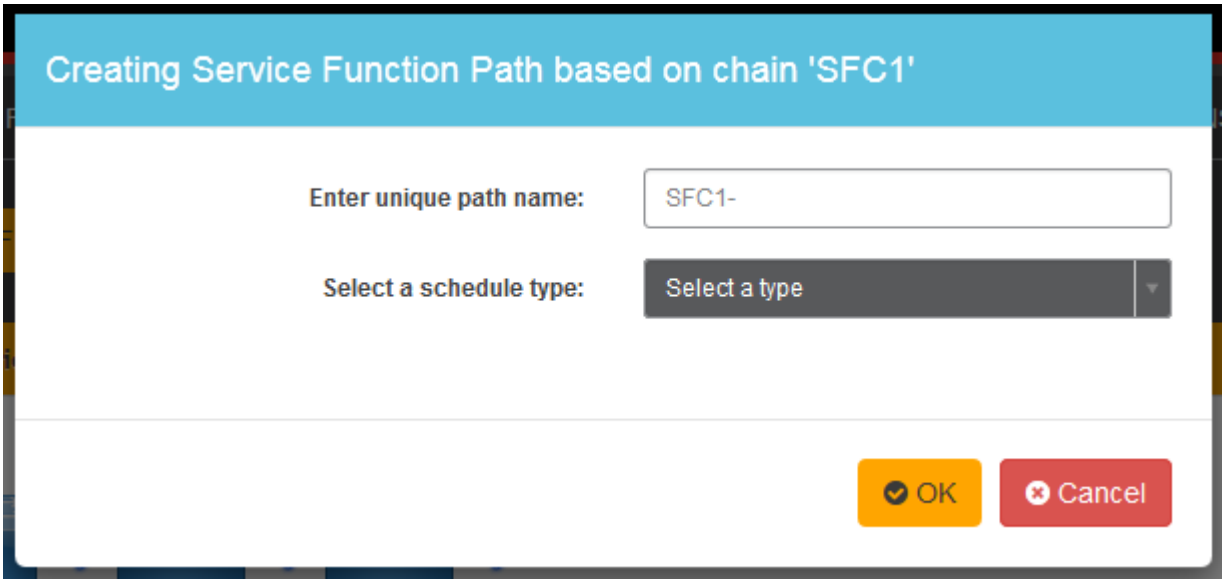

Fig. 101: Karaf Web UI

algorithms are committed at the moment.

Random

Select Service Function from the name list randomly.

Overview

The Random algorithm is used to select one Service Function from the name list which it gets from the Service Function Type randomly.

Prerequisites

- Service Function information are stored in datastore.
- Either no algorithm or the Random algorithm is selected.

Target Environment

The Random algorithm will work either no algorithm type is selected or the Random algorithm is selected.

Instructions

Once the plugins are installed into Karaf successfully, a user can use his favorite method to select the Random scheduling algorithm type. There are no special instructions for using the Random algorithm.

Round Robin

Select Service Function from the name list in Round Robin manner.

Overview

The Round Robin algorithm is used to select one Service Function from the name list which it gets from the Service Function Type in a Round Robin manner, this will balance workloads to all Service Functions. However, this method cannot help all Service Functions load the same workload because it's flow-based Round Robin.

Prerequisites

- Service Function information are stored in datastore.
- Round Robin algorithm is selected

Target Environment

The Round Robin algorithm will work one the Round Robin algorithm is selected.

Instructions

Once the plugins are installed into Karaf successfully, a user can use his favorite method to select the Round Robin scheduling algorithm type. There are no special instructions for using the Round Robin algorithm.

Load Balance Algorithm

Select appropriate Service Function by actual CPU utilization.

Overview

The Load Balance Algorithm is used to select appropriate Service Function by actual CPU utilization of service functions. The CPU utilization of service function obtained from monitoring information reported via NETCONF.

Prerequisites

- CPU-utilization for Service Function.
- NETCONF server.
- NETCONF client.
- Each VM has a NETCONF server and it could work with NETCONF client well.

Instructions

Set up VMs as Service Functions. enable NETCONF server in VMs. Ensure that you specify them separately. For example:

1. Set up 4 VMs include 2 SFs' type are Firewall, Others are Napt44. Name them as firewall-1, firewall-2, napt44-1, napt44-2 as Service Function. The four VMs can run either the same server or different servers.
2. Install NETCONF server on every VM and enable it. More information on NETCONF can be found on the OpenDaylight wiki here: https://wiki.opendaylight.org/view/OpenDaylight_Controller:Config:Examples:Netconf:Manual_netopeer_installation
3. Get Monitoring data from NETCONF server. These monitoring data should be get from the NETCONF server which is running in VMs. The following static XML data is an example:

static XML data like this:

```
<?xml version="1.0" encoding="UTF-8"?>
<service-function-description-monitor-report>
  <SF-description>
    <number-of-dataports>2</number-of-dataports>
    <capabilities>
      <supported-packet-rate>5</supported-packet-rate>
      <supported-bandwidth>10</supported-bandwidth>
      <supported-ACL-number>2000</supported-ACL-number>
      <RIB-size>200</RIB-size>
      <FIB-size>100</FIB-size>
      <ports-bandwidth>
        <port-bandwidth>
          <port-id>1</port-id>
```

(continues on next page)

(continued from previous page)

```
          <ipaddress>10.0.0.1</ipaddress>
          <macaddress>00:1e:67:a2:5f:f4</macaddress>
          <supported-bandwidth>20</supported-bandwidth>
        </port-bandwidth>
        <port-bandwidth>
          <port-id>2</port-id>
          <ipaddress>10.0.0.2</ipaddress>
          <macaddress>01:1e:67:a2:5f:f6</macaddress>
          <supported-bandwidth>10</supported-bandwidth>
        </port-bandwidth>
      </ports-bandwidth>
    </capabilities>
  </SF-description>
  <SF-monitoring-info>
    <liveness>true</liveness>
    <resource-utilization>
        <packet-rate-utilization>10</packet-rate-utilization>
        <bandwidth-utilization>15</bandwidth-utilization>
        <CPU-utilization>12</CPU-utilization>
        <memory-utilization>17</memory-utilization>
        <available-memory>8</available-memory>
        <RIB-utilization>20</RIB-utilization>
        <FIB-utilization>25</FIB-utilization>
        <power-utilization>30</power-utilization>
        <SF-ports-bandwidth-utilization>
          <port-bandwidth-utilization>
            <port-id>1</port-id>
            <bandwidth-utilization>20</bandwidth-utilization>
          </port-bandwidth-utilization>
          <port-bandwidth-utilization>
            <port-id>2</port-id>
            <bandwidth-utilization>30</bandwidth-utilization>
          </port-bandwidth-utilization>
        </SF-ports-bandwidth-utilization>
    </resource-utilization>
  </SF-monitoring-info>
</service-function-description-monitor-report>
```

1. Unzip SFC release tarball.
2. Run SFC: ${sfc}/bin/karaf. More information on Service Function Chaining can be found on the OpenDaylight SFC's wiki page: https://wiki.opendaylight.org/view/Service_Function_Chaining:Main

1. Deploy the SFC2 (firewall-abstract2napt44-abstract2) and click button to Create Rendered Service Path in SFC UI (http://localhost:8181/sfc/index.html).
2. Verify the Rendered Service Path to ensure the CPU utilization of the selected hop is the minimum one among all the service functions with same type. The correct RSP is firewall-1napt44-2

Shortest Path Algorithm

Select appropriate Service Function by Dijkstra's algorithm. Dijkstra's algorithm is an algorithm for finding the shortest paths between nodes in a graph.

Overview

The Shortest Path Algorithm is used to select appropriate Service Function by actual topology.

Prerequisites

- Depolyed topology (include SFFs, SFs and their links).
- Dijkstra's algorithm. More information on Dijkstra's algorithm can be found on the wiki here: http://en.wikipedia.org/wiki/Dijkstra%27s_algorithm

Instructions

1. Unzip SFC release tarball.
2. Run SFC: ${sfc}/bin/karaf.
3. Depoly SFFs and SFs. import the service-function-forwarders.json and service-functions.json in UI (http://localhost:8181/sfc/index.html#/sfc/config)

service-function-forwarders.json:

```
{
  "service-function-forwarders": {
    "service-function-forwarder": [
      {
        "name": "SFF-br1",
        "service-node": "OVSDB-test01",
        "rest-uri": "http://localhost:5001",
        "sff-data-plane-locator": [
          {
            "name": "eth0",
            "service-function-forwarder-ovs:ovs-bridge": {
              "uuid": "4c3778e4-840d-47f4-b45e-0988e514d26c",
              "bridge-name": "br-tun"
            },
            "data-plane-locator": {
              "port": 5000,
              "ip": "192.168.1.1",
              "transport": "service-locator:vxlan-gpe"
            }
          }
        ],
        "service-function-dictionary": [
          {
            "sff-sf-data-plane-locator": {
              "port": 10001,
              "ip": "10.3.1.103"
            },
            "name": "napt44-1",
            "type": "service-function-type:napt44"
          },
          {
            "sff-sf-data-plane-locator": {
              "port": 10003,
              "ip": "10.3.1.102"
```

(continues on next page)

(continued from previous page)


```
            },
            "name": "firewall-1",
            "type": "service-function-type:firewall"
          }
        ],
        "connected-sff-dictionary": [
          {
            "name": "SFF-br3"
          }
        ]
      },
      {
        "name": "SFF-br2",
        "service-node": "OVSDB-test01",
        "rest-uri": "http://localhost:5002",
        "sff-data-plane-locator": [
          {
            "name": "eth0",
            "service-function-forwarder-ovs:ovs-bridge": {
              "uuid": "fd4d849f-5140-48cd-bc60-6ad1f5fc0a1",
              "bridge-name": "br-tun"
            },
            "data-plane-locator": {
              "port": 5000,
              "ip": "192.168.1.2",
              "transport": "service-locator:vxlan-gpe"
            }
          }
        ],
        "service-function-dictionary": [
          {
            "sff-sf-data-plane-locator": {
              "port": 10002,
              "ip": "10.3.1.103"
            },
            "name": "napt44-2",
            "type": "service-function-type:napt44"
          },
          {
            "sff-sf-data-plane-locator": {
              "port": 10004,
              "ip": "10.3.1.101"
            },
            "name": "firewall-2",
            "type": "service-function-type:firewall"
          }
        ],
        "connected-sff-dictionary": [
          {
            "name": "SFF-br3"
          }
        ]
      },
      {
        "name": "SFF-br3",
        "service-node": "OVSDB-test01",
        "rest-uri": "http://localhost:5005",
```


(continues on next page)

(continued from previous page)

```
        "sff-data-plane-locator": [
          {
            "name": "eth0",
            "service-function-forwarder-ovs:ovs-bridge": {
              "uuid": "fd4d849f-5140-48cd-bc60-6ad1f5fc0a4",
              "bridge-name": "br-tun"
            },
            "data-plane-locator": {
              "port": 5000,
              "ip": "192.168.1.2",
              "transport": "service-locator:vxlan-gpe"
            }
          }
        ],
        "service-function-dictionary": [
          {
            "sff-sf-data-plane-locator": {
              "port": 10005,
              "ip": "10.3.1.104"
            },
            "name": "test-server",
            "type": "service-function-type:dpi"
          },
          {
            "sff-sf-data-plane-locator": {
              "port": 10006,
              "ip": "10.3.1.102"
            },
            "name": "test-client",
            "type": "service-function-type:dpi"
          }
        ],
        "connected-sff-dictionary": [
          {
            "name": "SFF-br1"
          },
          {
            "name": "SFF-br2"
          }
        ]
      }
    ]
  }
}
```

service-functions.json:

```
{
  "service-functions": {
    "service-function": [
      {
        "rest-uri": "http://localhost:10001",
        "ip-mgmt-address": "10.3.1.103",
        "sf-data-plane-locator": [
          {
            "name": "preferred",
            "port": 10001,
```

(continues on next page)

(continued from previous page)

```
            "ip": "10.3.1.103",
            "service-function-forwarder": "SFF-br1"
          }
        ],
        "name": "napt44-1",
        "type": "service-function-type:napt44",
        "nsh-aware": true
      },
      {
        "rest-uri": "http://localhost:10002",
        "ip-mgmt-address": "10.3.1.103",
        "sf-data-plane-locator": [
          {
            "name": "master",
            "port": 10002,
            "ip": "10.3.1.103",
            "service-function-forwarder": "SFF-br2"
          }
        ],
        "name": "napt44-2",
        "type": "service-function-type:napt44",
        "nsh-aware": true
      },
      {
        "rest-uri": "http://localhost:10003",
        "ip-mgmt-address": "10.3.1.103",
        "sf-data-plane-locator": [
          {
            "name": "1",
            "port": 10003,
            "ip": "10.3.1.102",
            "service-function-forwarder": "SFF-br1"
          }
        ],
        "name": "firewall-1",
        "type": "service-function-type:firewall",
        "nsh-aware": true
      },
      {
        "rest-uri": "http://localhost:10004",
        "ip-mgmt-address": "10.3.1.103",
        "sf-data-plane-locator": [
          {
            "name": "2",
            "port": 10004,
            "ip": "10.3.1.101",
            "service-function-forwarder": "SFF-br2"
          }
        ],
        "name": "firewall-2",
        "type": "service-function-type:firewall",
        "nsh-aware": true
      },
      {
        "rest-uri": "http://localhost:10005",
        "ip-mgmt-address": "10.3.1.103",
        "sf-data-plane-locator": [
```

(continues on next page)

(continued from previous page)

```
          {
            "name": "3",
            "port": 10005,
            "ip": "10.3.1.104",
            "service-function-forwarder": "SFF-br3"
          }
        ],
        "name": "test-server",
        "type": "service-function-type:dpi",
        "nsh-aware": true
      },
      {
        "rest-uri": "http://localhost:10006",
        "ip-mgmt-address": "10.3.1.103",
        "sf-data-plane-locator": [
          {
            "name": "4",
            "port": 10006,
            "ip": "10.3.1.102",
            "service-function-forwarder": "SFF-br3"
          }
        ],
        "name": "test-client",
        "type": "service-function-type:dpi",
        "nsh-aware": true
      }
    ]
  }
}
```

The depolyed topology like this:

```
          +----+           +----+          +----+
          |sff1|+----------|sff3|----------+|sff2|
          +----+           +----+          +----+
            |                                |
     +--------------+                 +--------------+
     |              |                 |              |
+----------+   +--------+         +----------+   +--------+
|firewall-1|   |napt44-1|         |firewall-2|   |napt44-2|
+----------+   +--------+         +----------+   +--------+
```

- Deploy the SFC2(firewall-abstract2napt44-abstract2), select "Shortest Path" as schedule type and click button to Create Rendered Service Path in SFC UI (http://localhost:8181/sfc/index.html).
- Verify the Rendered Service Path to ensure the selected hops are linked in one SFF. The correct RSP is firewall-1napt44-1 or firewall-2napt44-2. The first SF type is Firewall in Service Function Chain. So the algorithm will select first Hop randomly among all the SFs type is Firewall. Assume the first selected SF is firewall-2. All the path from firewall-1 to SF which type is Napt44 are list:
  - Path1: firewall-2 → sff2 → napt44-2
  - Path2: firewall-2 → sff2 → sff3 → sff1 → napt44-1 The shortest path is Path1, so the selected next hop is napt44-2.

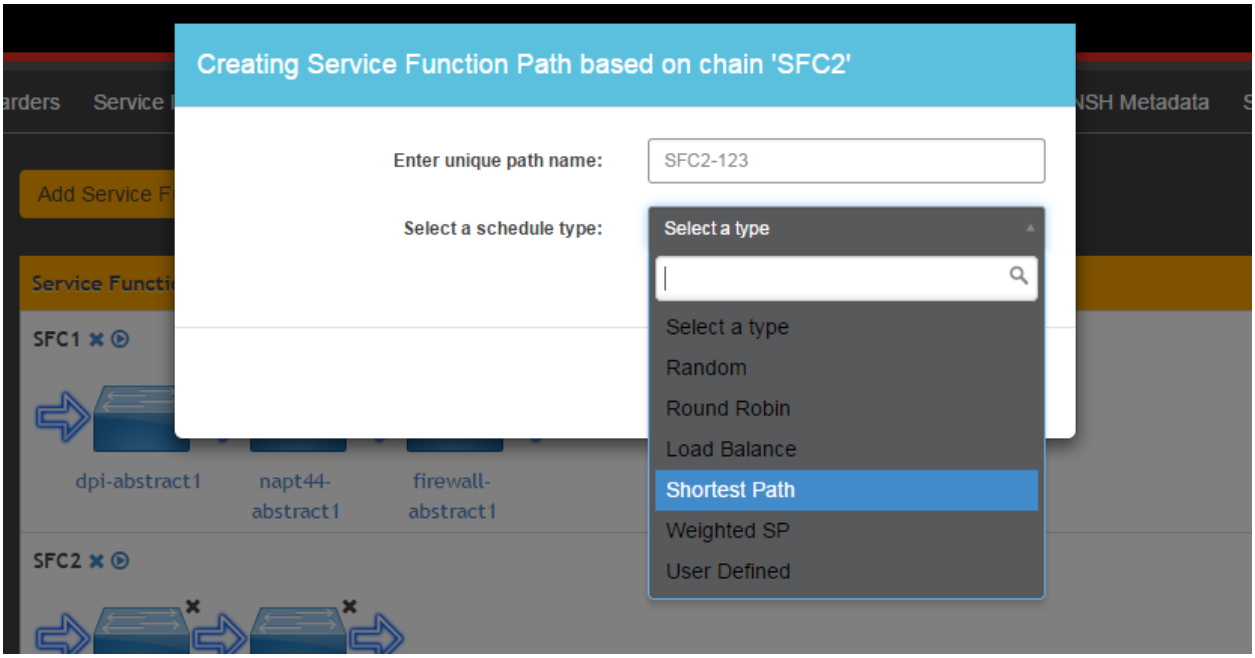

Fig. 102: select schedule type

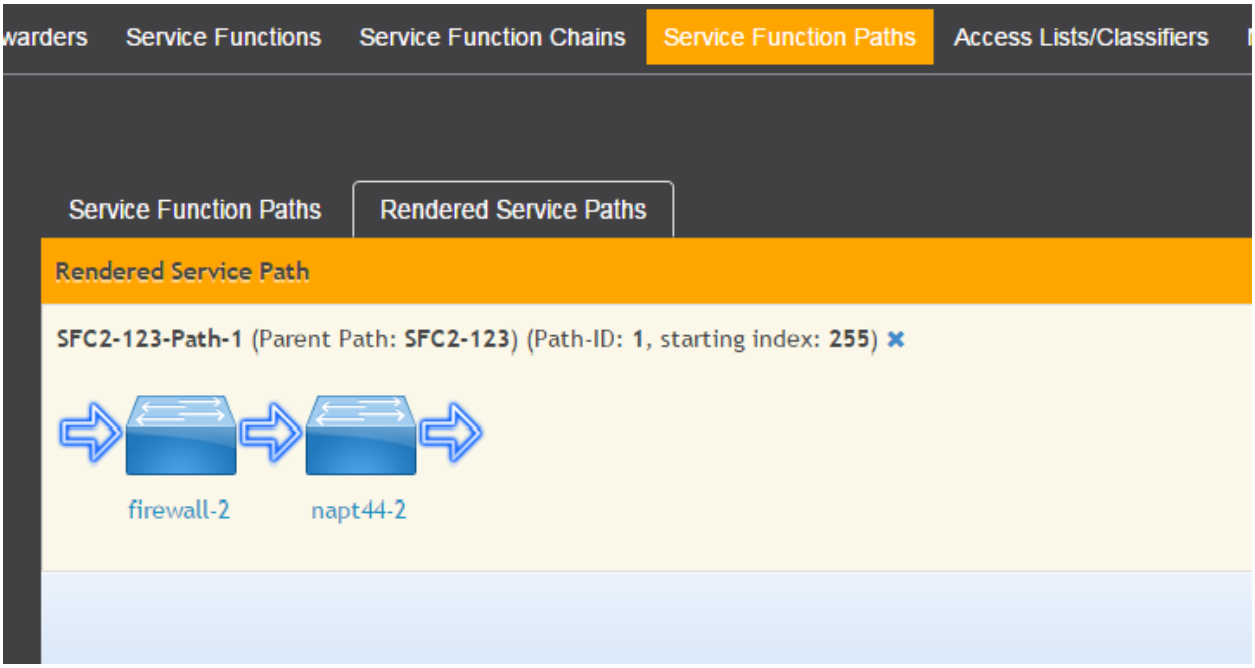

Fig. 103: rendered service path

Service Function Load Balancing User Guide

Overview

SFC Load-Balancing feature implements load balancing of Service Functions, rather than a one-to-one mapping between Service-Function-Forwarder and Service-Function.

Load Balancing Architecture

Service Function Groups (SFG) can replace Service Functions (SF) in the Rendered Path model. A Service Path can only be defined using SFGs or SFs, but not a combination of both.

Relevant objects in the YANG model are as follows:

1. Service-Function-Group-Algorithm:

```
Service-Function-Group-Algorithms {
    Service-Function-Group-Algorithm {
        String name
        String type
    }
}
```

```
Available types: ALL, SELECT, INDIRECT, FAST_FAILURE
```

2. Service-Function-Group:

```
Service-Function-Groups {
    Service-Function-Group {
        String name
        String serviceFunctionGroupAlgorithmName
        String type
        String groupId
        Service-Function-Group-Element {
            String service-function-name
            int index
        }
    }
}
```

3. ServiceFunctionHop: holds a reference to a name of SFG (or SF)

Tutorials

This tutorial will explain how to create a simple SFC configuration, with SFG instead of SF. In this example, the SFG will include two existing SF.

Setup SFC

For general SFC setup and scenarios, please see the SFC wiki page: https://wiki.opendaylight.org/view/Service_Function_Chaining:Main#SFC_101

Create an algorithm

POST - http://127.0.0.1:8181/restconf/config/service-function-group-algorithm:service-function-group-algorithms

```
{
    "service-function-group-algorithm": [
      {
        "name": "alg1"
        "type": "ALL"
      }
   ]
}
```

(Header "content-type": application/json)

Verify: get all algorithms

GET - http://127.0.0.1:8181/restconf/config/service-function-group-algorithm:service-function-group-algorithms

In order to delete all algorithms: DELETE - http://127.0.0.1:8181/restconf/config/service-function-group-algorithm:service-function-group-algorithms

Create a group

POST - http://127.0.0.1:8181/restconf/config/service-function-group:service-function-groups

```
{
    "service-function-group": [
    {
        "rest-uri": "http://localhost:10002",
        "ip-mgmt-address": "10.3.1.103",
        "algorithm": "alg1",
        "name": "SFG1",
        "type": "service-function-type:napt44",
        "sfc-service-function": [
            {
                "name":"napt44-104"
            },
            {
                "name":"napt44-103-1"
            }
        ]
      }
    ]
}
```

Verify: get all SFG's

GET - http://127.0.0.1:8181/restconf/config/service-function-group:service-function-groups

SFC Proof of Transit User Guide

Overview

Early Service Function Chaining (SFC) Proof of Transit (SFC Proof of Transit) implements Service Chaining Proof of Transit functionality on capable switches. After the creation of an Rendered Service Path (RSP), a user can configure to enable SFC proof of transit on the selected RSP to effect the proof of transit.

Common acronyms used in the following sections:

- SF - Service Function
- SFF - Service Function Forwarder
- SFC - Service Function Chain
- SFP - Service Function Path
- RSP - Rendered Service Path
- SFCPOT - Service Function Chain Proof of Transit

SFC Proof of Transit Architecture

When SFC Proof of Transit is initialized, all required listeners are registered to handle incoming data. It involves `SfcPotNodeListener` which stores data about all node devices including their mountpoints (used here as databrokers), `SfcPotRSPDataListener`, `RenderedPathListener`. `RenderedPathListener` is used to listen for RSP changes. `SfcPotRSPDataListener` implements RPC services to enable or disable SFC Proof of Transit on a particular RSP. When the SFC Proof of Transit is invoked, RSP listeners and service implementations are setup to receive SFCPOT configurations. When a user configures via a POST RPC call to enable SFCPOT on a particular RSP, the configuration drives the creation of necessary augmentations to the RSP to effect the SFCPOT configurations.

SFC Proof of Transit details

Several deployments use traffic engineering, policy routing, segment routing or service function chaining (SFC) to steer packets through a specific set of nodes. In certain cases regulatory obligations or a compliance policy require to prove that all packets that are supposed to follow a specific path are indeed being forwarded across the exact set of nodes specified. I.e. if a packet flow is supposed to go through a series of service functions or network nodes, it has to be proven that all packets of the flow actually went through the service chain or collection of nodes specified by the policy. In case the packets of a flow weren't appropriately processed, a proof of transit egress device would be required to identify the policy violation and take corresponding actions (e.g. drop or redirect the packet, send an alert etc.) corresponding to the policy.

The SFCPOT approach is based on meta-data which is added to every packet. The meta data is updated at every hop and is used to verify whether a packet traversed all required nodes. A particular path is either described by a set of secret keys, or a set of shares of a single secret. Nodes on the path retrieve their individual keys or shares of a key (using for e.g. Shamir's Shared Sharing Secret scheme) from a central controller. The complete key set is only known to the verifier- which is typically the ultimate node on a path that requires proof of transit. Each node in the path uses its secret or share of the secret to update the meta-data of the packets as the packets pass through the node. When the verifier receives a packet, it can use its key(s) along with the meta-data to validate whether the packet traversed the service chain correctly.

SFC Proof of Transit entities

In order to implement SFC Proof of Transit for a service function chain, an RSP is a pre-requisite to identify the SFC to enable SFC PoT on. SFC Proof of Transit for a particular RSP is enabled by an RPC request to the controller along with necessary parameters to control some of the aspects of the SFC Proof of Transit process.

The RPC handler identifies the RSP and generates SFC Proof of Transit parameters like secret share, secret etc., and adds the generated SFCPOT configuration parameters to SFC main as well as the various SFC hops. The last node in the SFC is configured as a verifier node to allow SFCPOT Proof of Transit process to be completed.

The SFCPOT configuration generators and related handling are done by `SfcPotAPI`, `SfcPotConfigGenerator`, `SfcPotListener`, `SfcPotPolyAPI`, `SfcPotPolyClassAPI` and `SfcPotPolyClass`.

Administering SFC Proof of Transit

To use the SFC Proof of Transit Karaf, at least the following Karaf features must be installed:

- odl-sfc-model
- odl-sfc-provider
- odl-sfc-netconf
- odl-restconf
- odl-netconf-topology
- odl-netconf-connector-all
- odl-sfc-pot

SFC Proof of Transit Tutorial

Overview

This tutorial is a simple example how to configure Service Function Chain Proof of Transit using SFC POT feature.

Preconditions

To enable a device to handle SFC Proof of Transit, it is expected that the netconf server device advertise capability as under ioam-scv.yang present under src/main/yang folder of sfc-pot feature. It is also expected that netconf notifications be enabled and its support capability advertised as capabilities.

It is also expected that the devices are netconf mounted and available in the topology-netconf store.

Instructions

When SFC Proof of Transit is installed, all netconf nodes in topology-netconf are processed and all capable nodes with accessible mountpoints are cached.

First step is to create the required RSP as usually done.

Once RSP name is avaiable it is used to send a POST RPC to the controller similar to below:

```
POST ./restconf/operations/sfc-ioam-nb-pot:enable-sfc-ioam-pot-rendered-path

{
  "input": {
    "sfc-ioam-pot-rsp-name": "rsp1"
  }
}
```

The following can be used to disable the SFC Proof of Transit on an RSP which removes the augmentations and stores back the RSP without the SFCPOT enabled features and also sending down a delete configuration to the SFCPOT configuration sub-tree in the nodes.

```
POST ./restconf/operations/sfc-ioam-nb-pot:disable-sfc-ioam-pot-rendered-path

{
  "input": {
    "sfc-ioam-pot-rsp-name": "rsp1"
  }
}
```

SNBI User Guide

This section describes how to use the SNBI feature in OpenDaylight and contains configuration, administration, and management section for the feature.

Overview

Key distribution in a scaled network has always been a challenge. Typically, operators must perform some manual key distribution process before secure communication is possible between a set of network devices. The Secure Network Bootstrapping Infrastructure (SNBI) project securely and automatically brings up an integrated set of network devices and controllers, simplifying the process of bootstrapping network devices with the keys required for secure communication. SNBI enables connectivity to the network devices by assigning unique IPv6 addresses and bootstrapping devices with the required keys. Admission control of devices into a specific domain is achieved using whitelist of authorized devices.

SNBI Architecture

At a high level, SNBI architecture consists of the following components:

- SNBI Registrar
- SNBI Forwarding Element (FE)

SNBI Registrar

Registrar is a device in a network that validates device against a whitelist and delivers device domain certificate. Registrar includes the following:

- RESCONF API for Domain Whitelist Configuration
- SNBI Southbound Plugin
- Certificate Authority

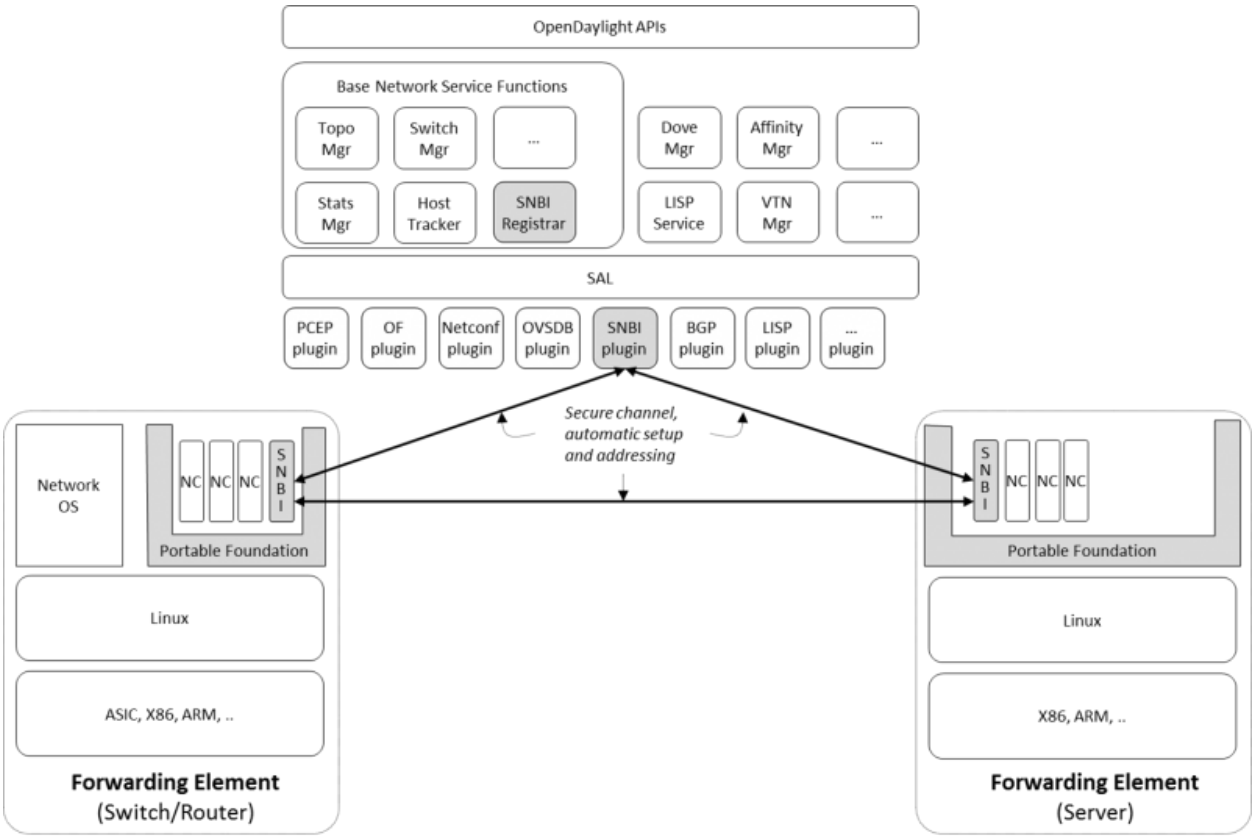

Fig. 104: SNBI Architecture Diagram

RESTCONF API for Domain Whitelist Configuration:.

Below is the YANG model to configure the whitelist of devices for a particular domain.

```
module snbi {
    //The yang version - today only 1 version exists. If omitted defaults to 1.
    yang-version 1;

    //a unique namespace for this SNBI module, to uniquely identify it from other␣
↪modules that may have the same name.
    namespace "http://netconfcentral.org/ns/snbi";

    //a shorter prefix that represents the namespace for references used below
    prefix snbi;

    //Defines the organization which defined / owns this .yang file.
    organization "Netconf Central";

    //defines the primary contact of this yang file.
    contact "snbi-dev";

    //provides a description of this .yang file.
    description "YANG version for SNBI.";

    //defines the dates of revisions for this yang file
    revision "2024-07-02" {
        description "SNBI module";
    }

    typedef UDI {
        type string;
        description "Unique Device Identifier";
    }

    container snbi-domain {
        leaf domain-name {
            type string;
            description "The SNBI domain name";
        }

        list device-list {
            key "list-name";

            leaf list-name {
                type string;
                description "Name of the device list";
            }

            leaf list-type {
                type enumeration {
                    enum "white";
                }
                description "Indicates the type of the list";
            }

            leaf active {
                type boolean;
                description "Indicates whether the list is active or not";
```

(continues on next page)

(continued from previous page)

```
            }

            list devices {
                key "device-identifier";
                leaf device-identifier {
                    type union {
                        type UDI;
                    }
                }
             }
         }
    }
}
```

Southbound Plugin:.

The Southbound Plugin implements the protocol state machine necessary to exchange device identifiers, and deliver certificates.

Certificate Authority:.

A simple certificate authority is implemented using the Bouncy Castle package. The Certificate Authority creates the certificates from the device CSR requests received from the devices. The certificates thus generated are delivered to the devices using the Southbound Plugin.

SNBI Forwarding Element

The forwarding element must be installed or unpacked on a Linux host whose network layer traffic must be secured. The FE performs the following functions:

- Neighour Discovery
- Bootstrap
- Host Configuration

Neighbour Discovery:.

Neighbour Discovery (ND) is the first step in accommodating devices in a secure network. SNBI performs periodic neighbour discovery of SNBI agents by transmitting ND hello packets. The discovered devices are populated in an ND table. Neighbour Discovery is periodic and bidirectional. ND hello packets are transmitted every 10 seconds. A 40 second refresh timer is set for each discovered neighbour. On expiry of the refresh timer, the Neighbour Adjacency is removed from the ND table as the Neighbour Adjacency is no longer valid. It is possible that the same SNBI neighbour is discovered on multiple links, the expiry of a device on one link does not automatically remove the device entry from the ND table.

Bootstrapping:.

Bootstrapping a device involves the following sequential steps:

- Authenticate a device using device identifier (UDI or SUDI)
- Allocate the appropriate device ID and IPv6 address to uniquely identify the device in the network
- Allocate the required keys by installing a Device Domain Certificate
- Accommodate the device in the domain

Host Configuration:.

Involves configuring a host to create a secure overlay network, assigning appropriate ipv6 address, setting up gre tunnels, securing the tunnels traffic via IPsec and enabling connectivity via a routing protocol.

The SNBI Forwarding Element is packaged in a docker container available at this link: https://hub.docker.com/r/snbi/boron/. For more information on docker, refer to this link: https://docs.docker.com/linux/.

Prerequisites for Configuring SNBI

Before proceeding further, ensure that the following system requirements are met:

- 64bit Ubunutu 14.04 LTS
- 4GB RAM
- 4GB of hard disk space, sufficient enough to store certificates
- Java Virtual Machine 1.8 or above
- Apache Maven 3.3.3 or above
- Make sure the time on all the devices or synced either manually or using NTP
- The docker version must be greater than 1.0 on a 14.04 Ubuntu

Configuring SNBI

This section contains the following:

- Setting up SNBI Registrar on the controller
- Configuring Whitelist
- Setting up SNBI FE on Linux Hosts

Setting up SNBI Registrar on the controller

This section contains the following:

- Configuring the Registrar Host
- Installing Karaf Package
- Configuring SNBI Registrar

Configuring the Registrar Host:.

Before enabling the SNBI registrar service, assign an IPv6 address to an interface on the registrar host. This is to bind the registrar service to an IPv6 address (**fd08::aaaa:bbbb:1/128**).

```
sudo ip link add snbi-ra type dummy
sudo ip addr add fd08::aaaa:bbbb:1/128 dev snbi-ra
sudo ifconfig snbi-ra up
```

Installing Karaf Package:.

Download the karaf package from this link: http://www.opendaylight.org/software/downloads, unzip and run the `karaf` executable present in the bin folder. Here is an example of this step:

```
cd distribution-karaf-0.3.0-Boron/bin
./karaf
```

Additional information on useful Karaf commands are available at this link: https://wiki.opendaylight.org/view/CrossProject:Integration_Group:karaf.

Configuring SNBI Registrar:.

Before you perform this step, ensure that you have completed the tasks *above*:

To use RESTCONF APIs, install the RESTCONF feature available in the Karaf package. If required, install mdsal-apidocs module for access to documentation. Refer https://wiki.opendaylight.org/view/OpenDaylight_Controller:MD-SAL:Restconf_API_Explorer for more information on MDSAL API docs.

Use the commands below to install the required features and verify the same.

```
feature:install odl-restconf
feature:install odl-mdsal-apidocs
feature:install odl-snbi-all
feature:list -i
```

After confirming that the features are installed, use the following command to start SNBI registrar:

```
snbi:start <domain-name>
```

Configuring Whitelist

The registrar must be configured with a whitelist of devices that are accommodated in a specific domain. The YANG for configuring the domain and the associated whitelist in the controller is avaialble at this link: https://wiki.opendaylight.org/view/SNBI_Architecture_and_Design#Registrar_YANG_Definition. It is recommended to use Postman to configure the registrar using RESTCONF.

This section contains the following:

- Installing PostMan
- Configuring Whitelist using REST API

Installing PostMan:.

Follow the steps below to install postman on your Google Chrome Browser.

- Install Postman via Google Chrome browser available at this link: https://chrome.google.com/webstore/detail/postman-rest-client/fdmmgilgnpjigdojojpjoooidkmcomcm?hl=en
- In the chrome browser address bar, enter: chrome://apps/
- Click Postman.
- Enter the URL.
- Click Headers.
- Enter Accept: header.
- Click Basic Auth tab to create user credentials, such as user name and password.
- Send.

You can download a sample Postman configuration to get started from this link: https://www.getpostman.com/collections/c929a2a4007ffd0a7b51

Configuring Whitelist using REST API:.

The POST method below configures a domain - “secure-domain” and configures a whitelist set of devices to be accommodated to the domain.

```
{
  "snbi-domain": {
    "domain-name": "secure-domain",
    "device-list": [
      {
        "list-name": "demo list",
        "list-type": "white",
        "active": true,
        "devices": [
          {
            "device-id": "UDI-FirstFE"
          },
          {
            "device-id": "UDI-dev1"
          },
          {
            "device-id": "UDI-dev2"
          }
        ]
      }
     ]
  }
}
```

The associated device ID must be configured on the SNBI FE (see below). You can also use REST APIs using the API docs interface to push the domain and whitelist information. The API docs could be accessed at link:http://localhost:8080/apidoc/explorer. More details on the API docs is available at link:https://wiki.opendaylight.org/view/OpenDaylight_Controller:MD-SAL:Restconf_API_Explorer

Setting up SNBI FE on Linux Hosts

The SNBI Daemon is used to bootstrap the host device with a valid device domain certificate and IP address for connectivity and to create a reachable overlay network by interacting with multiple software modules.

Device UDI:.

The Device UDI or the device Unique Identifier can be derived from a multitude of parameters in the host machine, but most derived parameters are already known or do not remain constant across reloads. Therefore, every SNBI FE must be configured explicitly with a UDI that is present in the device whitelist.

First Forwarding Element:.

The registrar service IP address must be provided to the first host (Forwarding Element) to be bootstrapped. As mentioned in the "Configuring the Registrar Host" section, the registrar service IP address is **fd08::aaaa:bbbb:1**. The First Forwarding Element must be configured with this IPv6 address.

Running the SNBI docker image:.

The SNBI FE in the docker image picks the UDI of the ForwardingElement via an environment variable provided when executing docker instance. If the Forwarding Element is a first forwarding element, the IP address of the registrar service should also be provided.

```
sudo docker run -v /etc/timezone:/etc/timezone:ro --net=host --privileged=true
--rm -t -i -e SNBI_UDI=UDI-FirstFE  -e SNBI_REGISTRAR=fd08::aaaa:bbbb:1 snbi/
↪boron:latest /bin/bash
```

After the docker image is executed, you are placed in the snbi.d command prompt.

A new Forwarding Element is bootstrapped in the same way, except that the registrar IP address is not required while running the docker image.

```
sudo docker run --net=host --privileged=true --rm -t -i -e SNBI_UDI=UDI-dev1 snbi/
↪boron:latest /bin/bash
```

Administering or Managing SNBI

The SNBI daemon provides various show commands to verify the current state of the daemon. The commands are completed automatically when you press Tab in your keyboard. There are help strings "?" to list commands.

```
snbi.d > show snbi
        device                  Host deevice
        neighbors               SNBI Neighbors
        debugs                  Debugs enabled
        certificate             Certificate information
```

SNMP Plugin User Guide

Installing Feature

The SNMP Plugin can be installed using a single karaf feature: **odl-snmp-plugin**

After starting Karaf:

- Install the feature: **feature:install odl-snmp-plugin**
- Expose the northbound API: **feature:install odl-restconf**

Northbound APIs

There are two exposed northbound APIs: snmp-get & snmp-set

SNMP GET

Default URL: http://localhost:8181/restconf/operations/snmp:snmp-get

POST Input

Field Name	Type	Description	Example	Required?
ip-address	Ipv4 Address	The IPv4 Address of the desired network node	10.86.3.13	Yes
oid	String	The Object Identifier of the desired MIB table/object	1.3.6.1.2.1.1.1	Yes
get-type	ENUM (GET, GET-NEXT, GET-BULK, GET-WALK)	The type of get request to send	GET-BULK	Yes
community	String	The community string to use for the SNMP request	private	No. (Default: public)

Example.

```
{
    "input": {
        "ip-address": "10.86.3.13",
        "oid" : "1.3.6.1.2.1.1.1",
        "get-type" : "GET-BULK",
        "community" : "private"
    }
}
```

POST Output

Field Name	Type	Description
results	List of { "value" : String } pairs	The results of the SNMP query

Example.

```
{
    "snmp:results": [
        {
            "value": "Ethernet0/0/0",
            "oid": "1.3.6.1.2.1.2.2.1.2.1"
        },
        {
            "value": "FastEthernet0/0/0",
            "oid": "1.3.6.1.2.1.2.2.1.2.2"
        },
        {
            "value": "GigabitEthernet0/0/0",
            "oid": "1.3.6.1.2.1.2.2.1.2.3"
        }
    ]
}
```

SNMP SET

Default URL: http://localhost:8181/restconf/operations/snmp:snmp-set

POST Input

Field Name	Type	Description	Example	Required?
ip-address	Ipv4 Address	The Ipv4 address of the desired network node	10.86.3.13	Yes
oid	String	The Object Identifier of the desired MIB object	1.3.6.2.1.1.1	Yes
value	String	The value to set on the network device	"Hello World"	Yes
community	String	The community string to use for the SNMP request	private	No. (Default: public)

Example.

```
{
    "input": {
        "ip-address": "10.86.3.13",
        "oid" : "1.3.6.1.2.1.1.1.0",
        "value" : "Sample description",
        "community" : "private"
    }
}
```

POST Output

On a successful SNMP-SET, no output is presented, just a HTTP status of 200.

Errors

If any errors happen in the set request, you will be presented with an error message in the output.

For example, on a failed set request you may see an error like:

```
{
    "errors": {
        "error": [
            {
                "error-type": "application",
                "error-tag": "operation-failed",
                "error-message": "SnmpSET failed with error status: 17, error index:␣
↪1. StatusText: Not writable"
            }
        ]
    }
}
```

which corresponds to Error status 17 in the SNMPv2 RFC: https://tools.ietf.org/html/rfc1905.

SNMP4SDN User Guide

Overview

We propose a southbound plugin that can control the off-the-shelf commodity Ethernet switches for the purpose of building SDN using Ethernet switches. For Ethernet switches, forwarding table, VLAN table, and ACL are where one can install flow configuration on, and this is done via SNMP and CLI in the proposed plugin. In addition, some settings required for Ethernet switches in SDN, e.g., disabling STP and flooding, are proposed.

Configuration

Just follow the steps:

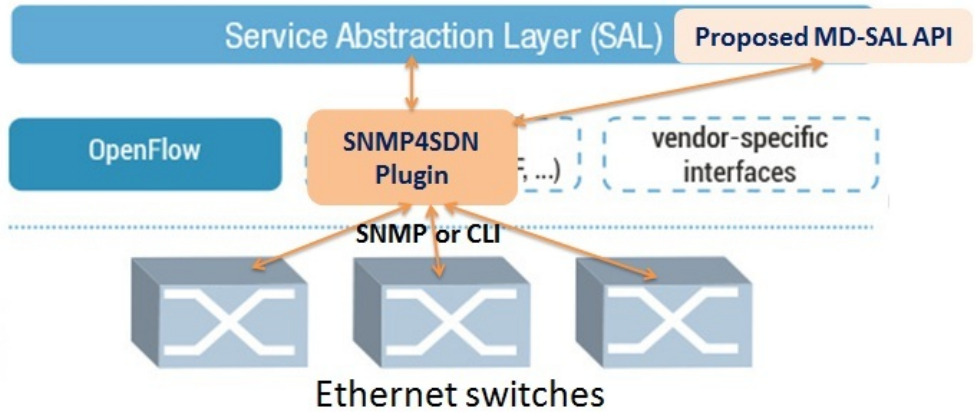

Fig. 105: SNMP4SDN as an OpenDaylight southbound plugin

Prepare the switch list database file

A sample is here, and we suggest to save it as */etc/snmp4sdn_swdb.csv* so that SNMP4SDN Plugin can automatically load this file. Note that the first line is title and should not be removed.

Prepare the vendor-specific configuration file

A sample is here, and we suggest to save it as */etc/snmp4sdn_VendorSpecificSwitchConfig.xml* so that SNMP4SDN Plugin can automatically load this file.

Install SNMP4SDN Plugin

If using SNMP4SDN Plugin provided in OpenDaylight release, just do the following from the Karaf CLI:

```
feature:install odl-snmp4sdn-all
```

Troubleshooting

Installation Troubleshooting

Feature installation failure

When trying to install a feature, if the following failure occurs:

```
Error executing command: Could not start bundle ...
Reason: Missing Constraint: Require-Capability: osgi.ee; filter="(&(osgi.
↪ee=JavaSE)(version=1.7))"
```

A workaround: exit Karaf, and edit the file <karaf_directory>/etc/config.properties, remove the line *${services-${karaf.framework}}* and the ", " in the line above.

Runtime Troubleshooting

Problem starting SNMP Trap Interface

It is possible to get the following exception during controller startup. (The error would not be printed in Karaf console, one may see it in <karaf_directory>/data/log/karaf.log)

```
2014-01-31 15:00:44.688 CET [fileinstall-./plugins] WARN  o.o.snmp4sdn.internal.
↪SNMPListener - Problem starting SNMP Trap Interface: {}
 java.net.BindException: Permission denied
        at java.net.PlainDatagramSocketImpl.bind0(Native Method) ~[na:1.7.0_51]
        at java.net.AbstractPlainDatagramSocketImpl.
↪bind(AbstractPlainDatagramSocketImpl.java:95) ~[na:1.7.0_51]
        at java.net.DatagramSocket.bind(DatagramSocket.java:376) ~[na:1.7.0_51]
        at java.net.DatagramSocket.<init>(DatagramSocket.java:231) ~[na:1.7.0_51]
        at java.net.DatagramSocket.<init>(DatagramSocket.java:284) ~[na:1.7.0_51]
        at java.net.DatagramSocket.<init>(DatagramSocket.java:256) ~[na:1.7.0_51]
        at org.snmpj.SNMPTrapReceiverInterface.<init>(SNMPTrapReceiverInterface.
↪java:126) ~[org.snmpj-1.4.3.jar:na]
        at org.snmpj.SNMPTrapReceiverInterface.<init>(SNMPTrapReceiverInterface.
↪java:99) ~[org.snmpj-1.4.3.jar:na]
        at org.opendaylight.snmp4sdn.internal.SNMPListener.<init>(SNMPListener.
↪java:75) ~[bundlefile:na]
        at org.opendaylight.snmp4sdn.core.internal.Controller.start(Controller.
↪java:174) [bundlefile:na]
...
```

This indicates that the controller is being run as a user which does not have sufficient OS privileges to bind the SNMPTRAP port (162/UDP)

Switch list file missing

The SNMP4SDN Plugin needs a switch list file, which is necessary for topology discovery and should be provided by the administrator (so please prepare one for the first time using SNMP4SDN Plugin, here is the sample). The default file path is /etc/snmp4sdn_swdb.csv. SNMP4SDN Plugin would automatically load this file and start topology discovery. If this file is not ready there, the following message like this will pop up:

```
2016-02-02 04:21:52,476 | INFO| Event Dispatcher | CmethUtil                          |␣
↪466 - org.opendaylight.snmp4sdn - 0.3.0.SNAPSHOT | CmethUtil.readDB() err: {}
java.io.FileNotFoundException: /etc/snmp4sdn_swdb.csv (No such file or directory)
    at java.io.FileInputStream.open0(Native Method)[:1.8.0_65]
    at java.io.FileInputStream.open(FileInputStream.java:195)[:1.8.0_65]
    at java.io.FileInputStream.<init>(FileInputStream.java:138)[:1.8.0_65]
    at java.io.FileInputStream.<init>(FileInputStream.java:93)[:1.8.0_65]
    at java.io.FileReader.<init>(FileReader.java:58)[:1.8.0_65]
    at org.opendaylight.snmp4sdn.internal.util.CmethUtil.readDB(CmethUtil.java:66)
    at org.opendaylight.snmp4sdn.internal.util.CmethUtil.<init>(CmethUtil.java:43)
...
```

Configuration

Just follow the steps:

1. Prepare the switch list database file

A sample is here, and we suggest to save it as */etc/snmp4sdn_swdb.csv* so that SNMP4SDN Plugin can automatically load this file.

Note: The first line is title and should not be removed.

2. Prepare the vendor-specific configuration file

A sample is here, and we suggest to save it as */etc/snmp4sdn_VendorSpecificSwitchConfig.xml* so that SNMP4SDN Plugin can automatically load this file.

3. Install SNMP4SDN Plugin

If using SNMP4SDN Plugin provided in OpenDaylight release, just do the following:

Launch Karaf in Linux console:

```
cd <Boron_controller_directory>/bin
(for example, cd distribution-karaf-x.x.x-Boron/bin)
```

```
./karaf
```

Then in Karaf console, execute:

```
feature:install odl-snmp4sdn-all
```

4. Load switch list

For initialization, we need to feed SNMP4SDN Plugin the switch list. Actually SNMP4SDN Plugin automatically try to load the switch list at /etc/snmp4sdn_swdb.csv if there is. If so, this step could be skipped. In Karaf console, execute:

```
snmp4sdn:ReadDB <switch_list_path>
(For example, snmp4sdn:ReadDB /etc/snmp4sdn_swdb.csv)
(in Windows OS, For example, snmp4sdn:ReadDB D://snmp4sdn_swdb.csv)
```

A sample is here, and we suggest to save it as */etc/snmp4sdn_swdb.csv* so that SNMP4SDN Plugin can automatically load this file.

Note: The first line is title and should not be removed.

5. Show switch list

```
snmp4sdn:PrintDB
```

Tutorial

Topology Service

Execute topology discovery

The SNMP4SDN Plugin automatically executes topology discovery on startup. One may use the following commands to invoke topology discovery manually. Note that you may need to wait for seconds for itto complete.

Note: Currently, one needs to manually execute *snmp4sdn:TopoDiscover* first (just once), then later the automatic topology discovery can be successful. If switches change (switch added or removed), *snmp4sdn:TopoDiscover* is also required. A future version will fix it to eliminate these requirements.

```
snmp4sdn:TopoDiscover
```

If one like to discover all inventory (i.e. switches and their ports) but not edges, just execute "TopoDiscoverSwitches":

```
snmp4sdn:TopoDiscoverSwitches
```

If one like to only discover all edges but not inventory, just execute "TopoDiscoverEdges":

```
snmp4sdn:TopoDiscoverEdges
```

You can also trigger topology discovery via the REST API by using `curl` from the Linux console (or any other REST client):

```
curl --user "admin":"admin" -H "Accept: application/json" -H "Content-type:␣
↪application/json" -X POST http://localhost:8181/restconf/operations/
↪topology:rediscover
```

You can change the periodic topology discovery interval via a REST API:

```
curl --user "admin":"admin" -H "Accept: application/json" -H "Content-type:␣
↪application/json" -X POST http://localhost:8181/restconf/operations/topology:set-
↪discovery-interval -d "{"input":{"interval-second":'<interval_time>'}}"
For example, set the interval as 15 seconds:
curl --user "admin":"admin" -H "Accept: application/json" -H "Content-type:␣
↪application/json" -X POST http://localhost:8181/restconf/operations/topology:set-
↪discovery-interval -d "{"input":{"interval-second":'15'}}"
```

Show the topology

SNMP4SDN Plugin supports to show topology via REST API:

- Get topology

```
curl --user "admin":"admin" -H "Accept: application/json" -H "Content-type:␣
↪application/json" -X POST http://localhost:8181/restconf/operations/
↪topology:get-edge-list
```

- Get switch list

```
curl --user "admin":"admin" -H "Accept: application/json" -H "Content-type:␣
↪application/json" -X POST http://localhost:8181/restconf/operations/
↪topology:get-node-list
```

- Get switches' ports list

```
curl --user "admin":"admin" -H "Accept: application/json" -H "Content-type:␣
↪application/json" -X POST http://localhost:8181/restconf/operations/
↪topology:get-node-connector-list
```

- The three commands above are just for user to get the latest topology discovery result, it does not trigger SNMP4SDN Plugin to do topology discovery.
- To trigger SNMP4SDN Plugin to do topology discover, as described in aforementioned *Execute topology discovery*.

Flow configuration

FDB configuration

SNMP4SDN supports to add entry on FDB table via REST API:

- Get FDB table

```
curl --user "admin":"admin" -H "Accept: application/json" -H "Content-type:␣
↪application/json" -X POST http://<controller_ip_address>:8181/restconf/
↪operations/fdb:get-fdb-table -d "{input:{"node-id":<switch-mac-address-in-
↪number>}}"

For example:
curl --user "admin":"admin" -H "Accept: application/json" -H "Content-type:␣
↪application/json" -X POST http://localhost:8181/restconf/operations/fdb:get-fdb-
↪table -d "{input:{"node-id":158969157063648}}"
```

- Get FDB table entry

```
curl --user "admin":"admin" -H "Accept: application/json" -H "Content-type:␣
↪application/json" -X POST http://<controller_ip_address>:8181/restconf/
↪operations/fdb:get-fdb-entry -d "{input:{"node-id":<switch-mac-address-in-
↪number>, "vlan-id":<vlan-id-in-number>, "dest-mac-addr":<destination-mac-
↪address-in-number>}}"

For example:
curl --user "admin":"admin" -H "Accept: application/json" -H "Content-type:␣
↪application/json" -X POST http://localhost:8181/restconf/operations/fdb:get-fdb-
↪entry -d "{input:{"node-id":158969157063648, "vlan-id":1, "dest-mac-addr
↪":158969157063648}}"
```

- Set FDB table entry

  (Notice invalid value: (1) non unicast mac (2) port not in the VLAN)

```
curl --user "admin":"admin" -H "Accept: application/json" -H "Content-type:␣
↪application/json" -X POST http://<controller_ip_address>:8181/restconf/
↪operations/fdb:set-fdb-entry -d "{input:{"node-id":<switch-mac-address-in-
↪number>, "vlan-id":<vlan-id-in-number>, "dest-mac-addr":<destination-mac-
↪address-in-number>, "port":<port-in-number>, "type":'<type>'}}"

For example:
curl --user "admin":"admin" -H "Accept: application/json" -H "Content-type:␣
↪application/json" -X POST http://localhost:8181/restconf/operations/fdb:set-fdb-
↪entry -d "{input:{"node-id":158969157063648, "vlan-id":1, "dest-mac-addr
↪":187649984473770, "port":23, "type":'MGMT'}}"
```

- Delete FDB table entry

```
curl --user "admin":"admin" -H "Accept: application/json" -H "Content-type:␣
↪application/json" -X POST http://<controller_ip_address>:8181/restconf/
↪operations/fdb:del-fdb-entry -d "{input:{"node-id":<switch-mac-address-in-
↪number>, "vlan-id":<vlan-id-in-number>, "dest-mac-addr":<destination-mac-
↪address-in-number>}}"

For example:
curl --user "admin":"admin" -H "Accept: application/json" -H "Content-type:␣
↪application/json" -X POST http://localhost:8181/restconf/operations/fdb:del-fdb-
↪entry -d "{input:{"node-id":158969157063648, "vlan-id":1, "dest-mac-addr
↪":187649984473770}}"
```

VLAN configuration

SNMP4SDN supports to add entry on VLAN table via REST API:

- Get VLAN table

```
curl --user "admin":"admin" -H "Accept: application/json" -H "Content-type:␣
↪application/json" -X POST http://<controller_ip_address>:8181/restconf/
↪operations/vlan:get-vlan-table -d "{input:{node-id:<switch-mac-address-in-
↪number>}}"

For example:
curl --user "admin":"admin" -H "Accept: application/json" -H "Content-type:␣
↪application/json" -X POST http://localhost:8181/restconf/operations/vlan:get-
↪vlan-table -d "{input:{node-id:158969157063648}}"
```

- Add VLAN

```
curl --user "admin":"admin" -H "Accept: application/json" -H "Content-type:␣
↪application/json" -X POST http://<controller_ip_address>:8181/restconf/
↪operations/vlan:add-vlan -d "{"input":{"node-id":<switch-mac-address-in-number>,
↪ "vlan-id":<vlan-id-in-number>, "vlan-name":'<vlan-name>'}}"

For example:
curl --user "admin":"admin" -H "Accept: application/json" -H "Content-type:␣
↪application/json" -X POST http://localhost:8181/restconf/operations/vlan:add-
↪vlan -d "{"input":{"node-id":158969157063648, "vlan-id":123, "vlan-name":'v123'}
↪}"
```

- Delete VLAN

```
curl --user "admin":"admin" -H "Accept: application/json" -H "Content-type:␣
↪application/json" -X POST http://<controller_ip_address>:8181/restconf/
↪operations/vlan:delete-vlan -d "{"input":{"node-id":<switch-mac-address-in-
↪number>, "vlan-id":<vlan-id-in-number>}}"

For example:
curl --user "admin":"admin" -H "Accept: application/json" -H "Content-type:␣
↪application/json" -X POST http://localhost:8181/restconf/operations/vlan:delete-
↪vlan -d "{"input":{"node-id":158969157063648, "vlan-id":123}}"
```

- Add VLAN and set ports

```
curl --user "admin":"admin" -H "Accept: application/json" -H "Content-type:␣
↪application/json" -X POST http://<controller_ip_address>:8181/restconf/
↪operations/vlan:add-vlan-and-set-ports -d "{"input":{"node-id":<switch-mac-
↪address-in-number>, "vlan-id":<vlan-id-in-number>, "vlan-name":'<vlan-name>',
↪"tagged-port-list":'<tagged-ports-separated-by-comma>', "untagged-port-list":'
↪<untagged-ports-separated-by-comma>'}}"

For example:
curl --user "admin":"admin" -H "Accept: application/json" -H "Content-type:␣
↪application/json" -X POST http://localhost:8181/restconf/operations/vlan:add-
↪vlan-and-set-ports -d "{"input":{"node-id":158969157063648, "vlan-id":123,
↪"vlan-name":'v123', "tagged-port-list":'1,2,3', "untagged-port-list":'4,5,6'}}"
```

- Set VLAN ports

```
curl --user "admin":"admin" -H "Accept: application/json" -H "Content-type:␣
↪application/json" -X POST http://<controller_ip_address>:8181/restconf/
↪operations/vlan:set-vlan-ports -d "{"input":{"node-id":<switch-mac-address-in-
↪number>, "vlan-id":<vlan-id-in-number>, "tagged-port-list":'<tagged-ports-
↪separated-by-comma>', "untagged-port-list":'<untagged-ports-separated-by-comma>
↪'}}"

For example:
curl --user "admin":"admin" -H "Accept: application/json" -H "Content-type:␣
↪application/json" -X POST http://localhost:8181/restconf/operations/vlan:set-
↪vlan-ports -d "{"input":{"node-id":"158969157063648", "vlan-id":"123", "tagged-
↪port-list":'4,5', "untagged-port-list":'2,3'}}"
```

ACL configuration

SNMP4SDN supports to add flow on ACL table via REST API. However, it is so far only implemented for the D-Link DGS-3120 switch.

ACL configuration via CLI is vendor-specific, and SNMP4SDN will support configuration with vendor-specific CLI in future release.

To do ACL configuration using the REST APIs, use commands like the following:

- Clear ACL table

```
curl --user "admin":"admin" -H "Accept: application/json" -H "Content-type:␣
↪application/json" -X POST http://<controller_ip_address>:8181/restconf/
↪operations/acl:clear-acl-table -d "{"input":{"nodeId":<switch-mac-address-in-
↪number>}}"
```

(continues on next page)

(continued from previous page)

```
For example:
curl --user "admin":"admin" -H "Accept: application/json" -H "Content-type:␣
↪application/json" -X POST http://localhost:8181/restconf/operations/acl:clear-
↪acl-table -d "{"input":{"nodeId":158969157063648}}"
```

- Create ACL profile (IP layer)

```
curl --user "admin":"admin" -H "Accept: application/json" -H "Content-type:␣
↪application/json" -X POST http://<controller_ip_address>:8181/restconf/
↪operations/acl:create-acl-profile -d "{input:{"nodeId":<switch-mac-address-in-
↪number>,"profile-id":<profile_id_in_number>,"profile-name":'<profile_name>',
↪"acl-layer":'IP',"vlan-mask":<vlan_mask_in_number>,"src-ip-mask":'<src_ip_mask>
↪',"dst-ip-mask":"<destination_ip_mask>"}}"

For example:
curl --user "admin":"admin" -H "Accept: application/json" -H "Content-type:␣
↪application/json" -X POST http://localhost:8181/restconf/operations/acl:create-
↪acl-profile -d "{input:{"nodeId":158969157063648,"profile-id":1,"profile-name":
↪'profile_1',"acl-layer":'IP',"vlan-mask":1,"src-ip-mask":'255.255.0.0',"dst-ip-
↪mask":'255.255.255.255'}}"
```

- Create ACL profile (MAC layer)

```
curl --user "admin":"admin" -H "Accept: application/json" -H "Content-type:␣
↪application/json" -X POST http://<controller_ip_address>:8181/restconf/
↪operations/acl:create-acl-profile -d "{input:{"nodeId":<switch-mac-address-in-
↪number>,"profile-id":<profile_id_in_number>,"profile-name":'<profile_name>',
↪"acl-layer":'ETHERNET',"vlan-mask":<vlan_mask_in_number>}}"

For example:
curl --user "admin":"admin" -H "Accept: application/json" -H "Content-type:␣
↪application/json" -X POST http://localhost:8181/restconf/operations/acl:create-
↪acl-profile -d "{input:{"nodeId":158969157063648,"profile-id":2,"profile-name":
↪'profile_2',"acl-layer":'ETHERNET',"vlan-mask":4095}}"
```

- Delete ACL profile

```
curl --user "admin":"admin" -H "Accept: application/json" -H "Content-type:␣
↪application/json" -X POST http://localhost:8181/restconf/operations/acl:del-acl-
↪profile -d "{input:{"nodeId":<switch-mac-address-in-number>,"profile-id":
↪<profile_id_in_number>}}"

For example:
curl --user "admin":"admin" -H "Accept: application/json" -H "Content-type:␣
↪application/json" -X POST http://localhost:8181/restconf/operations/acl:del-acl-
↪profile -d "{input:{"nodeId":158969157063648,"profile-id":1}}"
```

```
curl --user "admin":"admin" -H "Accept: application/json" -H "Content-type:␣
↪application/json" -X POST http://<controller_ip_address>:8181/restconf/
↪operations/acl:del-acl-profile -d "{input:{"nodeId":<switch-mac-address-in-
↪number>,"profile-name":"<profile_name>"}}"

For example:
curl --user "admin":"admin" -H "Accept: application/json" -H "Content-type:␣
↪application/json" -X POST http://localhost:8181/restconf/operations/acl:del-acl-
↪profile -d "{input:{"nodeId":158969157063648,"profile-name":'profile_2'}}"
```

- Set ACL rule

```
curl --user "admin":"admin" -H "Accept: application/json" -H "Content-type:␣
↪application/json" -X POST http://<controller_ip_address>:8181/restconf/
↪operations/acl:set-acl-rule -d "{input:{"nodeId":<switch-mac-address-in-number>,
↪"profile-id":<profile_id_in_number>,"profile-name":'<profile_name>',"rule-id":
↪<rule_id_in_number>,"port-list":[<port_number>,<port_number>,...],"acl-layer":'
↪<acl_layer>',"vlan-id":<vlan_id_in_number>,"src-ip":"<src_ip_address>","dst-ip":
↪'<dst_ip_address>',"acl-action":'<acl_action>'}}"
(<acl_layer>: IP or ETHERNET)
(<acl_action>: PERMIT as permit, DENY as deny)

For example:
curl --user "admin":"admin" -H "Accept: application/json" -H "Content-type:␣
↪application/json" -X POST http://localhost:8181/restconf/operations/acl:set-acl-
↪rule -d "{input:{"nodeId":158969157063648,"profile-id":1,"profile-name":
↪'profile_1',"rule-id":1,"port-list":[1,2,3],"acl-layer":'IP',"vlan-id":2,"src-ip
↪":'1.1.1.1',"dst-ip":'2.2.2.2',"acl-action":'PERMIT'}}"
```

- Delete ACL rule

```
curl --user "admin":"admin" -H "Accept: application/json" -H "Content-type:␣
↪application/json" -X POST http://<controller_ip_address>:8181/restconf/
↪operations/acl:del-acl-rule -d "{input:{"nodeId":<switch-mac-address-in-number>,
↪"profile-id":<profile_id_in_number>,"profile-name":'<profile_name>',"rule-id":
↪<rule_id_in_number>}}"

For example:
curl --user "admin":"admin" -H "Accept: application/json" -H "Content-type:␣
↪application/json" -X POST http://localhost:8181/restconf/operations/acl:del-acl-
↪rule -d "{input:{"nodeId":158969157063648,"profile-id":1,"profile-name":
↪'profile_1',"rule-id":1}}"
```

Special configuration

SNMP4SDN supports setting the following special configurations via REST API:

- Set STP port state

```
curl --user "admin":"admin" -H "Accept: application/json" -H "Content-type:␣
↪application/json" -X POST http://<controller_ip_address>:8181/restconf/
↪operations/config:set-stp-port-state -d "{input:{"node-id":<switch-mac-address-
↪in-number>, "port":<port_number>, enable:<true_or_false>}}"
(true: enable, false: disable)

For example:
curl --user "admin":"admin" -H "Accept: application/json" -H "Content-type:␣
↪application/json" -X POST http://localhost:8181/restconf/operations/config:set-
↪stp-port-state -d "{input:{"node-id":158969157063648, "port":2, enable:false}}"
```

- Get STP port state

```
curl --user "admin":"admin" -H "Accept: application/json" -H "Content-type:␣
↪application/json" -X POST http://<controller_ip_address>:8181/restconf/
↪operations/config:get-stp-port-state -d "{input:{"node-id":<switch-mac-address-
↪in-number>, "port":<port_number>}}"
```

(continues on next page)

(continued from previous page)

```
For example:
curl --user "admin":"admin" -H "Accept: application/json" -H "Content-type:␣
↪application/json" -X POST http://localhost:8181/restconf/operations/config:get-
↪stp-port-state -d "{input:{"node-id":158969157063648, "port":2}}"
```

- Get STP port root

```
curl --user "admin":"admin" -H "Accept: application/json" -H "Content-type:␣
↪application/json" -X POST http://<controller_ip_address>:8181/restconf/
↪operations/config:get-stp-port-root -d "{input:{"node-id":<switch-mac-address-
↪in-number>, "port":<port_number>}}"

For example:
curl --user "admin":"admin" -H "Accept: application/json" -H "Content-type:␣
↪application/json" -X POST http://localhost:8181/restconf/operations/config:get-
↪stp-port-root -d "{input:{"node-id":158969157063648, "port":2}}"
```

- Enable STP

```
curl --user "admin":"admin" -H "Accept: application/json" -H "Content-type:␣
↪application/json" -X POST http://<controller_ip_address>:8181/restconf/
↪operations/config:enable-stp -d "{input:{"node-id":<switch-mac-address-in-
↪number>}}"

For example:
curl --user "admin":"admin" -H "Accept: application/json" -H "Content-type:␣
↪application/json" -X POST http://localhost:8181/restconf/operations/
↪config:enable-stp -d "{input:{"node-id":158969157063648}}"
```

- Disable STP

```
curl --user "admin":"admin" -H "Accept: application/json" -H "Content-type:␣
↪application/json" -X POST http://<controller_ip_address>:8181/restconf/
↪operations/config:disable-stp -d "{input:{"node-id":<switch-mac-address-in-
↪number>}}"

For example:
curl --user "admin":"admin" -H "Accept: application/json" -H "Content-type:␣
↪application/json" -X POST http://localhost:8181/restconf/operations/
↪config:disable-stp -d "{input:{"node-id":158969157063648}}"
```

- Get ARP table

```
curl --user "admin":"admin" -H "Accept: application/json" -H "Content-type:␣
↪application/json" -X POST http://<controller_ip_address>:8181/restconf/
↪operations/config:get-arp-table -d "{input:{"node-id":<switch-mac-address-in-
↪number>}}"

For example:
curl --user "admin":"admin" -H "Accept: application/json" -H "Content-type:␣
↪application/json" -X POST http://localhost:8181/restconf/operations/config:get-
↪arp-table -d "{input:{"node-id":158969157063648}}"
```

- Set ARP entry

  (Notice to give IP address with subnet prefix)

```
curl --user "admin":"admin" -H "Accept: application/json" -H "Content-type:␣
↪application/json" -X POST http://<controller_ip_address>:8181/restconf/
↪operations/config:set-arp-entry -d "{input:{"node-id":<switch-mac-address-in-
↪number>, "ip-address":'<ip_address>', "mac-address":<mac_address_in_number>}}"

For example:
curl --user "admin":"admin" -H "Accept: application/json" -H "Content-type:␣
↪application/json" -X POST http://localhost:8181/restconf/operations/config:set-
↪arp-entry -d "{input:{"node-id":158969157063648, "ip-address":'10.217.9.9',
↪"mac-address":1}}"
```

- Get ARP entry

```
curl --user "admin":"admin" -H "Accept: application/json" -H "Content-type:␣
↪application/json" -X POST http://<controller_ip_address>:8181/restconf/
↪operations/config:get-arp-entry -d "{input:{"node-id":<switch-mac-address-in-
↪number>, "ip-address":'<ip_address>'}}"

For example:
curl --user "admin":"admin" -H "Accept: application/json" -H "Content-type:␣
↪application/json" -X POST http://localhost:8181/restconf/operations/config:get-
↪arp-entry -d "{input:{"node-id":158969157063648, "ip-address":'10.217.9.9'}}"
```

- Delete ARP entry

```
curl --user "admin":"admin" -H "Accept: application/json" -H "Content-type:␣
↪application/json" -X POST http://<controller_ip_address>:8181/restconf/
↪operations/config:delete-arp-entry -d "{input:{"node-id":<switch-mac-address-in-
↪number>, "ip-address":'<ip_address>'}}"

For example:
curl --user "admin":"admin" -H "Accept: application/json" -H "Content-type:␣
↪application/json" -X POST http://localhost:8181/restconf/operations/
↪config:delete-arp-entry -d "{input:{"node-id":158969157063648, "ip-address":'10.
↪217.9.9'}}"
```

Using Postman to invoke REST API

Besides using the curl tool to invoke REST API, like the examples aforementioned, one can also use GUI tool like Postman for better data display.

- Install Postman: Install Postman in the Chrome browser
- In the chrome browser bar enter

```
chrome://apps/
```

- Click on Postman.

Example: Get VLAN table using Postman

As the screenshot shown below, one needs to fill in required fields.

```
URL:
http://<controller_ip_address>:8181/restconf/operations/vlan:get-vlan-table

Accept header:
application/json

Content-type:
application/json

Body:
{input:{"node-id":<node_id>}}
for example:
{input:{"node-id":158969157063648}}
```

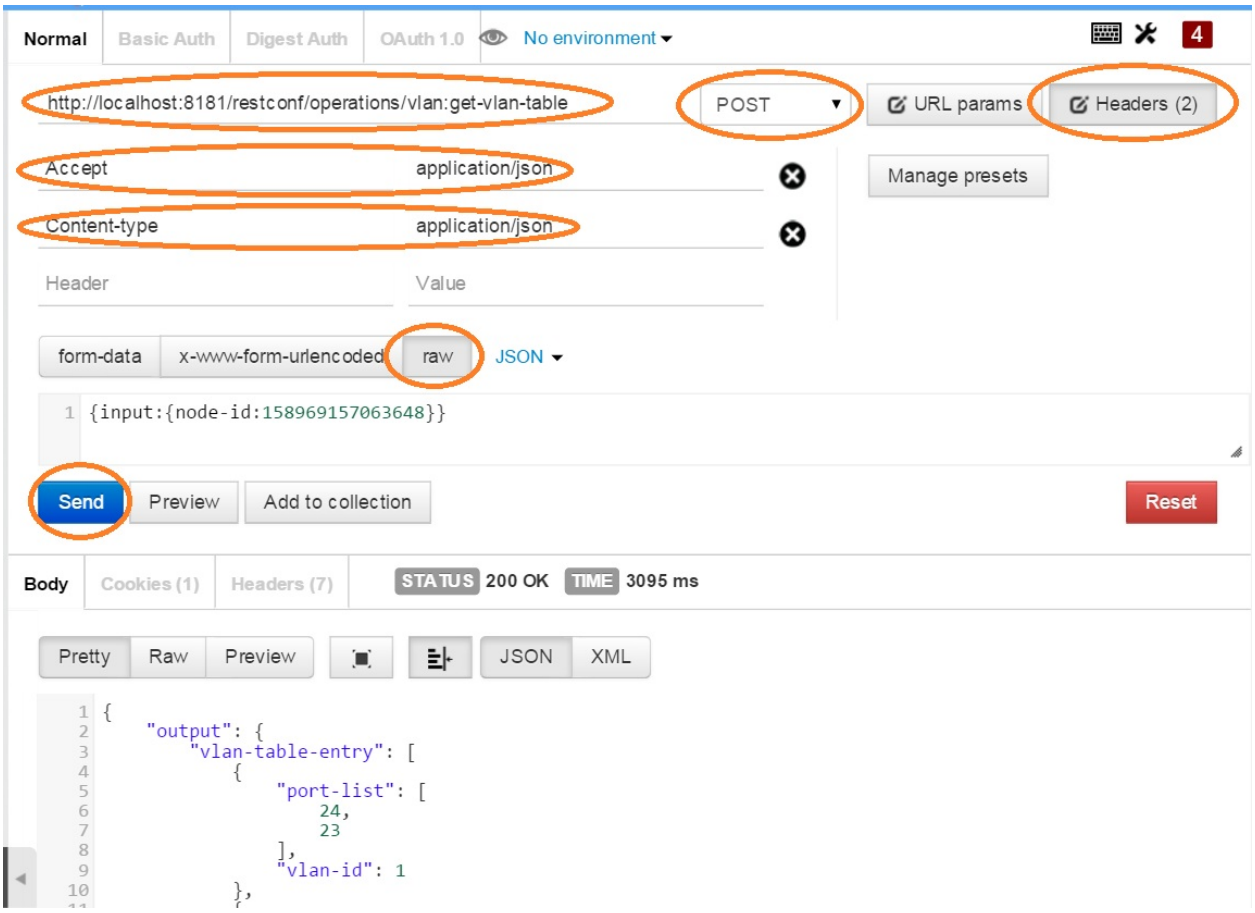

Fig. 106: Example: Get VLAN table using Postman

Multi-vendor support

So far the supported vendor-specific configurations:

- Add VLAN and set ports
- (More functions are TBD)

The SNMP4SDN Plugin would examine whether the configuration is described in the vendor-specific configuration

file. If yes, the configuration description would be adopted, otherwise just use the default configuration. For example, adding VLAN and setting the ports is supported via SNMP standard MIB. However we found some special cases, for example, certain Accton switch requires to add VLAN first and then allows to set the ports. So one may describe this in the vendor-specific configuration file.

A vendor-specific configuration file sample is here, and we suggest to save it as */etc/snmp4sdn_VendorSpecificSwitchConfig.xml* so that SNMP4SDN Plugin can automatically load it.

Help

- SNMP4SDN Wiki
- SNMP4SDN Mailing Lists: (user, developer)
- Latest troubleshooting in Wiki

SXP User Guide

Overview

SXP (Source-Group Tag eXchange Protocol) project is an effort to enhance OpenDaylight platform with IP-SGT (IP Address to Source Group Tag) bindings that can be learned from connected SXP-aware network nodes. The current implementation supports SXP protocol version 4 according to the Smith, Kandula - SXP IETF draft and grouping of peers and creating filters based on ACL/Prefix-list syntax for filtering outbound and inbound IP-SGT bindings. All protocol legacy versions 1-3 are supported as well. Additionally, version 4 adds bidirectional connection type as an extension of a unidirectional one.

SXP Architecture

The SXP Server manages all connected clients in separate threads and a common SXP protocol agreement is used between connected peers. Each SXP network peer is modelled with its pertaining class, e.g., SXP Server represents the SXP Speaker, SXP Listener the Client. The server program creates the ServerSocket object on a specified port and waits until a client starts up and requests connect on the IP address and port of the server. The client program opens a Socket that is connected to the server running on the specified host IP address and port.

The SXP Listener maintains connection with its speaker peer. From an opened channel pipeline, all incoming SXP messages are processed by various handlers. Message must be decoded, parsed and validated.

The SXP Speaker is a counterpart to the SXP Listener. It maintains a connection with its listener peer and sends composed messages.

The SXP Binding Handler extracts the IP-SGT binding from a message and pulls it into the SXP-Database. If an error is detected during the IP-SGT extraction, an appropriate error code and sub-code is selected and an error message is sent back to the connected peer. All transitive messages are routed directly to the output queue of SXP Binding Dispatcher.

The SXP Binding Dispatcher represents a selector that will decides how many data from the SXP-database will be sent and when. It is responsible for message content composition based on maximum message length.

The SXP Binding Filters handles filtering of outgoing and incoming IP-SGT bindings according to BGP filtering using ACL and Prefix List syntax for specifying filter or based on Peer-sequence length.

The SXP Domains feature provides isolation of SXP peers and bindings learned between them, also exchange of Bindings is possible across SXP-Domains by ACL, Prefix List or Peer-Sequence filters

Configuring SXP

The OpenDaylight Karaf distribution comes pre-configured with baseline SXP configuration. Configuration of SXP Nodes is also possible via NETCONF.

- **22-sxp-controller-one-node.xml** (defines the basic parameters)

Administering or Managing SXP

By RPC (response is XML document containing requested data or operation status):

- Get Connections POST http://127.0.0.1:8181/restconf/operations/sxp-controller:get-connections

```
<input xmlns:xsi="urn:opendaylight:sxp:controller">
 <domain-name>global</domain-name>
 <requested-node>0.0.0.100</requested-node>
</input>
```

- Add Connection POST http://127.0.0.1:8181/restconf/operations/sxp-controller:add-connection

```
<input xmlns:xsi="urn:opendaylight:sxp:controller">
 <requested-node>0.0.0.100</requested-node>
 <domain-name>global</domain-name>
 <connections>
  <connection>
   <peer-address>172.20.161.50</peer-address>
   <tcp-port>64999</tcp-port>
   <!-- Password setup: default | none leave empty -->
   <password>default</password>
   <!-- Mode: speaker/listener/both -->
   <mode>speaker</mode>
   <version>version4</version>
   <description>Connection to ASR1K</description>
   <!-- Timers setup: 0 to disable specific timer usability, the default value will␣
↪be used -->
   <connection-timers>
    <!-- Speaker -->
    <hold-time-min-acceptable>45</hold-time-min-acceptable>
    <keep-alive-time>30</keep-alive-time>
   </connection-timers>
  </connection>
  <connection>
   <peer-address>172.20.161.178</peer-address>
   <tcp-port>64999</tcp-port>
   <!-- Password setup: default | none leave empty-->
   <password>default</password>
   <!-- Mode: speaker/listener/both -->
   <mode>listener</mode>
   <version>version4</version>
   <description>Connection to ISR</description>
   <!-- Timers setup: 0 to disable specific timer usability, the default value will␣
↪be used -->
   <connection-timers>
    <!-- Listener -->
    <reconciliation-time>120</reconciliation-time>
    <hold-time>90</hold-time>
    <hold-time-min>90</hold-time-min>
```

(continues on next page)

(continued from previous page)

```
    <hold-time-max>180</hold-time-max>
   </connection-timers>
  </connection>
 </connections>
</input>
```

- Delete Connection POST http://127.0.0.1:8181/restconf/operations/sxp-controller:delete-connection

```
<input xmlns:xsi="urn:opendaylight:sxp:controller">
 <requested-node>0.0.0.100</requested-node>
 <domain-name>global</domain-name>
 <peer-address>172.20.161.50</peer-address>
</input>
```

- Add Binding Entry POST http://127.0.0.1:8181/restconf/operations/sxp-controller:add-entry

```
<input xmlns:xsi="urn:opendaylight:sxp:controller">
 <requested-node>0.0.0.100</requested-node>
 <domain-name>global</domain-name>
 <ip-prefix>192.168.2.1/32</ip-prefix>
 <sgt>20</sgt >
</input>
```

- Update Binding Entry POST http://127.0.0.1:8181/restconf/operations/sxp-controller:update-entry

```
<input xmlns:xsi="urn:opendaylight:sxp:controller">
 <requested-node>0.0.0.100</requested-node>
 <domain-name>global</domain-name>
 <original-binding>
  <ip-prefix>192.168.2.1/32</ip-prefix>
  <sgt>20</sgt>
 </original-binding>
 <new-binding>
  <ip-prefix>192.168.3.1/32</ip-prefix>
  <sgt>30</sgt>
 </new-binding>
</input>
```

- Delete Binding Entry POST http://127.0.0.1:8181/restconf/operations/sxp-controller:delete-entry

```
<input xmlns:xsi="urn:opendaylight:sxp:controller">
 <requested-node>0.0.0.100</requested-node>
 <domain-name>global</domain-name>
 <ip-prefix>192.168.3.1/32</ip-prefix>
 <sgt>30</sgt >
</input>
```

- Get Node Bindings

  This RPC gets particular device bindings. An SXP-aware node is identified with a unique Node-ID. If a user requests bindings for a Speaker 20.0.0.2, the RPC will search for an appropriate path, which contains 20.0.0.2 Node-ID, within locally learnt SXP data in the SXP database and replies with associated bindings. POST http://127.0.0.1:8181/restconf/operations/sxp-controller:get-node-bindings

```
<input xmlns:xsi="urn:opendaylight:sxp:controller">
 <requested-node>20.0.0.2</requested-node>
```

(continues on next page)

(continued from previous page)

```
 <bindings-range>all</bindings-range>
 <domain-name>global</domain-name>
</input>
```

- Get Binding SGTs POST http://127.0.0.1:8181/restconf/operations/sxp-controller:get-binding-sgts

```
<input xmlns:xsi="urn:opendaylight:sxp:controller">
 <requested-node>0.0.0.100</requested-node>
 <domain-name>global</domain-name>
 <ip-prefix>192.168.12.2/32</ip-prefix>
</input>
```

- Add PeerGroup with or without filters to node. POST http://127.0.0.1:8181/restconf/operations/sxp-controller:add-peer-group

```
<input xmlns="urn:opendaylight:sxp:controller">
 <requested-node>127.0.0.1</requested-node>
 <sxp-peer-group>
  <name>TEST</name>
  <sxp-peers>
  </sxp-peers>
  <sxp-filter>
   <filter-type>outbound</filter-type>
   <acl-entry>
    <entry-type>deny</entry-type>
    <entry-seq>1</entry-seq>
    <sgt-start>1</sgt-start>
    <sgt-end>100</sgt-end>
   </acl-entry>
   <acl-entry>
    <entry-type>permit</entry-type>
    <entry-seq>45</entry-seq>
    <matches>1</matches>
    <matches>3</matches>
    <matches>5</matches>
   </acl-entry>
  </sxp-filter>
 </sxp-peer-group>
</input>
```

- Delete PeerGroup with peer-group-name from node request-node. POST http://127.0.0.1:8181/restconf/operations/sxp-controller:delete-peer-group

```
<input xmlns="urn:opendaylight:sxp:controller">
 <requested-node>127.0.0.1</requested-node>
 <peer-group-name>TEST</peer-group-name>
</input>
```

- Get PeerGroup with peer-group-name from node request-node. POST http://127.0.0.1:8181/restconf/operations/sxp-controller:get-peer-group

```
<input xmlns="urn:opendaylight:sxp:controller">
 <requested-node>127.0.0.1</requested-node>
 <peer-group-name>TEST</peer-group-name>
</input>
```

- Add Filter to peer group on node request-node. POST http://127.0.0.1:8181/restconf/operations/sxp-controller:

add-filter

```
<input xmlns="urn:opendaylight:sxp:controller">
 <requested-node>127.0.0.1</requested-node>
 <peer-group-name>TEST</peer-group-name>
 <sxp-filter>
  <filter-type>outbound</filter-type>
  <acl-entry>
   <entry-type>deny</entry-type>
   <entry-seq>1</entry-seq>
   <sgt-start>1</sgt-start>
   <sgt-end>100</sgt-end>
  </acl-entry>
  <acl-entry>
   <entry-type>permit</entry-type>
   <entry-seq>45</entry-seq>
   <matches>1</matches>
   <matches>3</matches>
   <matches>5</matches>
  </acl-entry>
 </sxp-filter>
</input>
```

- Delete Filter from peer group on node request-node. POST http://127.0.0.1:8181/restconf/operations/sxp-controller:delete-filter

```
<input xmlns="urn:opendaylight:sxp:controller">
 <requested-node>127.0.0.1</requested-node>
 <peer-group-name>TEST</peer-group-name>
 <filter-type>outbound</filter-type>
</input>
```

- Update Filter of the same type in peer group on node request-node. POST http://127.0.0.1:8181/restconf/operations/sxp-controller:update-filter

```
<input xmlns="urn:opendaylight:sxp:controller">
 <requested-node>127.0.0.1</requested-node>
 <peer-group-name>TEST</peer-group-name>
 <sxp-filter>
  <filter-type>outbound</filter-type>
  <acl-entry>
   <entry-type>deny</entry-type>
   <entry-seq>1</entry-seq>
   <sgt-start>1</sgt-start>
   <sgt-end>100</sgt-end>
  </acl-entry>
  <acl-entry>
   <entry-type>permit</entry-type>
   <entry-seq>45</entry-seq>
   <matches>1</matches>
   <matches>3</matches>
   <matches>5</matches>
  </acl-entry>
 </sxp-filter>
</input>
```

- Add new SXP aware Node POST http://127.0.0.1:8181/restconf/operations/sxp-controller:add-node

```
<input xmlns="urn:opendaylight:sxp:controller">
    <node-id>1.1.1.1</node-id>
    <source-ip>0.0.0.0</source-ip>
    <timers>
        <retry-open-time>5</retry-open-time>
        <hold-time-min-acceptable>120</hold-time-min-acceptable>
        <delete-hold-down-time>120</delete-hold-down-time>
        <hold-time-min>90</hold-time-min>
        <reconciliation-time>120</reconciliation-time>
        <hold-time>90</hold-time>
        <hold-time-max>180</hold-time-max>
        <keep-alive-time>30</keep-alive-time>
    </timers>
    <mapping-expanded>150</mapping-expanded>
    <security>
        <password>password</password>
    </security>
    <tcp-port>64999</tcp-port>
    <version>version4</version>
    <description>ODL SXP Controller</description>
    <master-database></master-database>
</input>
```

- Delete SXP aware node POST http://127.0.0.1:8181/restconf/operations/sxp-controller:delete-node

```
<input xmlns="urn:opendaylight:sxp:controller">
 <node-id>1.1.1.1</node-id>
</input>
```

- Add SXP Domain on node request-node. POST http://127.0.0.1:8181/restconf/operations/sxp-controller:add-domain

```
<input xmlns="urn:opendaylight:sxp:controller">
  <node-id>1.1.1.1</node-id>
  <domain-name>global</domain-name>
</input>
```

- Delete SXP Domain on node request-node. POST http://127.0.0.1:8181/restconf/operations/sxp-controller:delete-domain

```
<input xmlns="urn:opendaylight:sxp:controller">
 <node-id>1.1.1.1</node-id>
 <domain-name>global</domain-name>
</input>
```

Use cases for SXP

Cisco has a wide installed base of network devices supporting SXP. By including SXP in OpenDaylight, the binding of policy groups to IP addresses can be made available for possible further processing to a wide range of devices, and applications running on OpenDaylight. The range of applications that would be enabled is extensive. Here are just a few of them:

OpenDaylight based applications can take advantage of the IP-SGT binding information. For example, access control can be defined by an operator in terms of policy groups, while OpenDaylight can configure access control lists on network elements using IP addresses, e.g., existing technology.

Interoperability between different vendors. Vendors have different policy systems. Knowing the IP-SGT binding for Cisco makes it possible to maintain policy groups between Cisco and other vendors.

OpenDaylight can aggregate the binding information from many devices and communicate it to a network element. For example, a firewall can use the IP-SGT binding information to know how to handle IPs based on the group-based ACLs it has set. But to do this with SXP alone, the firewall has to maintain a large number of network connections to get the binding information. This incurs heavy overhead costs to maintain all of the SXP peering and protocol information. OpenDaylight can aggregate the IP-group information so that the firewall need only connect to OpenDaylight. By moving the information flow outside of the network elements to a centralized position, we reduce the overhead of the CPU consumption on the enforcement element. This is a huge savings - it allows the enforcement point to only have to make one connection rather than thousands, so it can concentrate on its primary job of forwarding and enforcing.

OpenDaylight can relay the binding information from one network element to others. Changes in group membership can be propagated more readily through a centralized model. For example, in a security application a particular host (e.g., user or IP Address) may be found to be acting suspiciously or violating established security policies. The defined response is to put the host into a different source group for remediation actions such as a lower quality of service, restricted access to critical servers, or special routing conditions to ensure deeper security enforcement (e.g., redirecting the host's traffic through an IPS with very restrictive policies). Updated group membership for this host needs to be communicated to multiple network elements as soon as possible; a very efficient and effective method of propagation can be performed using OpenDaylight as a centralized point for relaying the information.

OpenDaylight can create filters for exporting and receiving IP-SGT bindings used on specific peer groups, thus can provide more complex maintaining of policy groups.

Although the IP-SGT binding is only one specific piece of information, and although SXP is implemented widely in a single vendor's equipment, bringing the ability of OpenDaylight to process and distribute the bindings, is a very specific immediate useful implementation of policy groups. It would go a long way to develop both the usefulness of OpenDaylight and of policy groups.

TSDR User Guide

This document describes how to use HSQLDB, HBase, and Cassandra data stores to capture time series data using Time Series Data Repository (TSDR) features in OpenDaylight. This document contains configuration, administration, management, usage, and troubleshooting sections for the features.

Overview

The Time Series Data Repository (TSDR) project in OpenDaylight (ODL) creates a framework for collecting, storing, querying, and maintaining time series data. TSDR provides the framework for plugging in proper data collectors to collect various time series data and store the data into TSDR Data Stores. With a common data model and generic TSDR data persistence APIs, the user can choose various data stores to be plugged into the TSDR persistence framework. Currently, three types of data stores are supported: HSQLDB relational database, HBase NoSQL database, and Cassandra NoSQL database.

With the capabilities of data collection, storage, query, aggregation, and purging provided by TSDR, network administrators can leverage various data driven appliations built on top of TSDR for security risk detection, performance analysis, operational configuration optimization, traffic engineering, and network analytics with automated intelligence.

TSDR Architecture

TSDR has the following major components:

- Data Collection Service

- Data Storage Service
- TSDR Persistence Layer with data stores as plugins
- TSDR Data Stores
- Data Query Service
- Grafana integration for time series data visualization
- Data Aggregation Service
- Data Purging Service

The Data Collection Service handles the collection of time series data into TSDR and hands it over to the Data Storage Service. The Data Storage Service stores the data into TSDR through the TSDR Persistence Layer. The TSDR Persistence Layer provides generic Service APIs allowing various data stores to be plugged in. The Data Aggregation Service aggregates time series fine-grained raw data into course-grained roll-up data to control the size of the data. The Data Purging Service periodically purges both fine-grained raw data and course-granined aggregated data according to user-defined schedules.

We have implemented The Data Collection Service, Data Storage Service, TSDR Persistence Layer, TSDR HSQLDB Data Store, TSDR HBase Data Store, and TSDR Cassandra Datastore. Among these services and components, time series data is communicated using a common TSDR data model, which is designed and implemented for the abstraction of time series data commonalities. With these functions, TSDR is able to collect the data from the data sources and store them into one of the TSDR data stores: HSQLDB Data Store, HBase Data Store or Cassandra Data Store. Besides a simple query command from Karaf console to retrieve data from the TSDR data stores, we also provided a Data Query Service for the user to use REST API to query the data from the data stores. Moreover, the user can use Grafana, which is a time series visualization tool to view the data stored in TSDR in various charting formats.

Configuring TSDR Data Stores

To Configure HSQLDB Data Store

The HSQLDB based storage files get stored automatically in <karaf install folder>/tsdr/ directory. If you want to change the default storage location, the configuration file to change can be found in <karaf install folder>/etc directory. The filename is org.ops4j.datasource-metric.cfg. Change the last portion of the url=jdbc:hsqldb:./tsdr/metric to point to different directory.

To Configure HBase Data Store

After installing HBase Server on the same machine as OpenDaylight, if the user accepts the default configuration of the HBase Data Store, the user can directly proceed with the installation of HBase Data Store from Karaf console.

Optionally, the user can configure TSDR HBase Data Store following HBase Data Store Configuration Procedure.

- HBase Data Store Configuration Steps
  - Open the file etc/tsdr-persistence-hbase.peroperties under karaf distribution directory.
  - Edit the following parameters:
    - HBase server name
    - HBase server port
    - HBase client connection pool size
    - HBase client write buffer size

After the configuration of HBase Data Store is complete, proceed with the installation of HBase Data Store from Karaf console.

- HBase Data Store Installation Steps
  - Start Karaf Console
  - Run the following commands from Karaf Console: feature:install odl-tsdr-hbase

To Configure Cassandra Data Store

Currently, there's no configuration needed for Cassandra Data Store. The user can use Cassandra data store directly after installing the feature from Karaf console.

Additionally separate commands have been implemented to install various data collectors.

Administering or Managing TSDR Data Stores

To Administer HSQLDB Data Store

Once the TSDR default datastore feature (odl-tsdr-hsqldb-all) is enabled, the TSDR captured OpenFlow statistics metrics can be accessed from Karaf Console by executing the command

```
tsdr:list <metric-category> <starttimestamp> <endtimestamp>
```

wherein

- <metric-category> = any one of the following categories FlowGroupStats, FlowMeterStats, FlowStats, FlowTableStats, PortStats, QueueStats
- <starttimestamp> = to filter the list of metrics starting this timestamp
- <endtimestamp> = to filter the list of metrics ending this timestamp
- <starttimestamp> and <endtimestamp> are optional.
- Maximum 1000 records will be displayed.

To Administer HBase Data Store

- Using Karaf Command to retrieve data from HBase Data Store

The user first need to install hbase data store from karaf console:

feature:install odl-tsdr-hbase

The user can retrieve the data from HBase data store using the following commands from Karaf console:

```
tsdr:list
tsdr:list <CategoryName> <StartTime> <EndTime>
```

Typing tab will get the context prompt of the arguments when typeing the command in Karaf console.

To Administer Cassandra Data Store

The user first needs to install Cassandra data store from Karaf console:

```
feature:install odl-tsdr-cassandra
```

Then the user can retrieve the data from Cassandra data store using the following commands from Karaf console:

```
tsdr:list
tsdr:list <CategoryName> <StartTime> <EndTime>
```

Typing tab will get the context prompt of the arguments when typeing the command in Karaf console.

Installing TSDR Data Collectors

When the user uses HSQLDB data store and installed "odl-tsdr-hsqldb-all" feature from Karaf console, besides the HSQLDB data store, OpenFlow data collector is also installed with this command. However, if the user needs to use other collectors, such as NetFlow Collector, Syslog Collector, SNMP Collector, and Controller Metrics Collector, the user needs to install them with separate commands. If the user uses HBase or Cassandra data store, no collectors will be installed when the data store is installed. Instead, the user needs to install each collector separately using feature install command from Karaf console.

The following is the list of supported TSDR data collectors with the associated feature install commands:

- OpenFlow Data Collector

  ```
  feature:install odl-tsdr-openflow-statistics-collector
  ```

- SNMP Data Collector

  ```
  feature:install odl-tsdr-snmp-data-collector
  ```

- NetFlow Data Collector

  ```
  feature:install odl-tsdr-netflow-statistics-collector
  ```

- sFlow Data Collector feature:install odl-tsdr-sflow-statistics-colletor
- Syslog Data Collector

  ```
  feature:install odl-tsdr-syslog-collector
  ```

- Controller Metrics Collector

  ```
  feature:install odl-tsdr-controller-metrics-collector
  ```

In order to use controller metrics collector, the user needs to install Sigar library.

The following is the instructions for installing Sigar library on Ubuntu:

- Install back end library by "sudo apt-get install libhyperic-sigar-java"
- Execute "export LD_LIBRARY_PATH=/usr/lib/jni/:/usr/lib:/usr/local/lib" to set the path of the JNI (you can add this to the ".bashrc" in your home directory)
- Download the file "sigar-1.6.4.jar". It might be also in your ".m2" directory under "~/.m2/resources/org/fusesource/sigar/1.6.4"

- Create the directory "org/fusesource/sigar/1.6.4" under the "system" directory in your controller home directory and place the "sigar-1.6.4.jar" there

Configuring TSDR Data Collectors

- SNMP Data Collector Device Credential Configuration

After installing SNMP Data Collector, a configuration file under etc/ directory of ODL distribution is generated: etc/tsdr.snmp.cfg is created.

The following is a sample tsdr.snmp.cfg file:

credentials=[192.168.0.2,public],[192.168.0.3,public]

The above credentials indicate that TSDR SNMP Collector is going to connect to two devices. The IPAddress and Read community string of these two devices are (192.168.0.2, public), and (192.168.0.3) respectively.

The user can make changes to this configuration file any time during runtime. The configuration will be picked up by TSDR in the next cycle of data collection.

Polling interval configuration for SNMP Collector and OpenFlow Stats Collector

The default polling interval of SNMP Collector and OpenFlow Stats Collector is 30 seconds and 15 seconds respectively. The user can change the polling interval through restconf APIs at any time. The new polling interval will be picked up by TSDR in the next collection cycle.

- Retrieve Polling Interval API for SNMP Collector
  - URL: http://localhost:8181/restconf/config/tsdr-snmp-data-collector:TSDRSnmpDataCollectorConfig
  - Verb: GET
- Update Polling Interval API for SNMP Collector
  - URL: http://localhost:8181/restconf/operations/tsdr-snmp-data-collector:setPollingInterval
  - Verb: POST
  - Content Type: application/json
  - Input Payload:

```
{
   "input": {
       "interval": "15000"
   }
}
```

- Retrieve Polling Interval API for OpenFlowStats Collector
  - URL: http://localhost:8181/restconf/config/tsdr-openflow-statistics-collector:TSDROSCConfig
  - Verb: GET
- Update Polling Interval API for OpenFlowStats Collector
  - URL: http://localhost:8181/restconf/operations/tsdr-openflow-statistics-collector:setPollingInterval
  - Verb: POST
  - Content Type: application/json
  - Input Payload:

```
{
   "input": {
       "interval": "15000"
   }
}
```

Querying TSDR from REST APIs

TSDR provides two REST APIs for querying data stored in TSDR data stores.

- Query of TSDR Metrics
  - URL: http://localhost:8181/tsdr/metrics/query
  - Verb: GET
  - Parameters:
    - tsdrkey=[NID=][DC=][MN=][RK=]

```
The TSDRKey format indicates the NodeID(NID), DataCategory(DC),␣
↪MetricName(MN), and RecordKey(RK) of the monitored objects.
For example, the following is a valid tsdrkey:
[NID=openflow:1][DC=FLOWSTATS][MN=PacketCount][RK=Node:openflow:1,Table:0,
↪Flow:3]
The following is also a valid tsdrkey:
tsdrkey=[NID=][DC=FLOWSTATS][MN=][RK=]
In the case when the sections in the tsdrkey is empty, the query will␣
↪return all the records in the TSDR data store that matches the filled␣
↪tsdrkey. In the above example, the query will return all the data in␣
↪FLOWSTATS data category.
The query will return only the first 1000 records that match the query␣
↪criteria.
```

    - from=<time_in_seconds>
    - until=<time_in_seconds>

The following is an example curl command for querying metric data from TSDR data store:

curl -G -v -H "Accept: application/json" -H "Content-Type: application/json" "http://localhost:8181/tsdr/metrics/query" –data-urlencode "tsdrkey=[NID=][DC=FLOWSTATS][MN=][RK=]" –data-urlencode "from=0" –data-urlencode "until=240000000000"|more

- Query of TSDR Log type of data
  - URL:http://localhost:8181/tsdr/logs/query
  - Verb: GET
  - Parameters:
    - tsdrkey=tsdrkey=[NID=][DC=][RK=]

```
The TSDRKey format indicates the NodeID(NID), DataCategory(DC), and␣
↪RecordKey(RK) of the monitored objects.
For example, the following is a valid tsdrkey:
[NID=openflow:1][DC=NETFLOW][RK]
The query will return only the first 1000 records that match the query␣
↪criteria.
```

* from=<time_in_seconds>
* until=<time_in_seconds>

The following is an example curl command for querying log type of data from TSDR data store:

curl -G -v -H "Accept: application/json" -H "Content-Type: application/json" "http://localhost:8181/tsdr/logs/query" –data-urlencode "tsdrkey=[NID=][DC=NETFLOW][RK=]" –data-urlencode "from=0" –data-urlencode "until=240000000000"lmore

Grafana integration with TSDR

TSDR provides northbound integration with Grafana time series data visualization tool. All the metric type of data stored in TSDR data store can be visualized using Grafana.

For the detailed instruction about how to install and configure Grafana to work with TSDR, please refer to the following link:

https://wiki.opendaylight.org/view/Grafana_Integration_with_TSDR_Step-by-Step

Purging Service configuration

After the data stores are installed from Karaf console, the purging service will be installed as well. A configuration file called tsdr.data.purge.cfg will be generated under etc/ directory of ODL distribution.

The following is the sample default content of the tsdr.data.purge.cfg file:

host=127.0.0.1 data_purge_enabled=true data_purge_time=23:59:59 data_purge_interval_in_minutes=1440 retention_time_in_hours=168

The host indicates the IPAddress of the data store. In the case when the data store is together with ODL controller, 127.0.0.1 should be the right value for the host IP. The other attributes are self-explained. The user can change those attributes at any time. The configuration change will be picked up right away by TSDR Purging service at runtime.

How to use TSDR to collect, store, and view OpenFlow Interface Statistics

Overview

This tutorial describes an example of using TSDR to collect, store, and view one type of time series data in OpenDaylight environment.

Prerequisites

You would need to have the following as prerequisits:

- One or multiple OpenFlow enabled switches. Alternatively, you can use mininet to simulate such a switch.
- Successfully installed OpenDaylight Controller.
- Successfully installed HBase Data Store following TSDR HBase Data Store Installation Guide.
- Connect the OpenFlow enabled switch(es) to OpenDaylight Controller.

Target Environment

HBase data store is only supported in Linux operation system.

Instructions

- Start OpenDaylight.
- Connect OpenFlow enabled switch(es) to the controller.
  - If using mininet, run the following commands from mininet command line:
    * mn –topo single,3 –controller *remote,ip=172.17.252.210,port=6653* –switch ovsk,protocols=OpenFlow13
- Install tsdr hbase feature from Karaf:
  - feature:install odl-tsdr-hbase
- Install OpenFlow Statistics Collector from Karaf:
  - feature:install odl-tsdr-openflow-statistics-collector
- run the following command from Karaf console:
  - tsdr:list PORTSTATS

You should be able to see the interface statistics of the switch(es) from the HBase Data Store. If there are too many rows, you can use "tsdr:list InterfaceStats|more" to view it page by page.

By tabbing after "tsdr:list", you will see all the supported data categories. For example, "tsdr:list FlowStats" will output the Flow statistics data collected from the switch(es).

Troubleshooting

Karaf logs

All TSDR features and components write logging information including information messages, warnings, errors and debug messages into karaf.log.

HBase and Cassandra logs

For HBase and Cassandra data stores, the database level logs are written into HBase log and Cassandra logs.

- HBase log
  - HBase log is under <HBase-installation-directory>/logs/.
- Cassandra log
  - Cassandra log is under {cassandra.logdir}/system.log. The default {cassandra.logdir} is /var/log/cassandra/.

Security

TSDR gets the data from a variety of sources, which can be secured in different ways.

- OpenFlow Security
  - The OpenFlow data can be configured with Transport Layer Security (TLS) since the OpenFlow Plugin that TSDR depends on provides this security support.
- SNMP Security
  - The SNMP version3 has security support. However, since ODL SNMP Plugin that TSDR depends on does not support version 3, we (TSDR) will not have security support at this moment.
- NetFlow Security
  - NetFlow, which cannot be configured with security so we recommend making sure it flows only over a secured management network.
- Syslog Security
  - Syslog, which cannot be configured with security so we recommend making sure it flows only over a secured management network.

Support multiple data stores simultaneously at runtime

TSDR supports running multiple data stores simultaneously at runtim. For example, it is possible to configure TSDR to push log type of data into Cassandra data store, while pushing metrics type of data into HBase.

When you install one TSDR data store from karaf console, such as using feature:install odl-tsdr-hsqldb, a properties file will be generated under <Karaf-distribution-directory>/etc/. For example, when you install hsqldb, a file called tsdr-persistence-hsqldb.properties is generated under that directory.

By default, all the types of data are supported in the data store. For example, the default content of tsdr-persistence-hsqldb.properties is as follows:

```
metric-persistency=true
log-persistency=true
binary-persistency=true
```

When the user would like to use different data stores to support different types of data, he/she could enable or disable a particular type of data persistence in the data stores by configuring the properties file accordingly.

For example, if the user would like to store the log type of data in HBase, and store the metric and binary type of data in Cassandra, he/she needs to install both hbase and cassandra data stores from Karaf console. Then the user needs to modify the properties file under <Karaf-distribution-directory>/etc as follows:

- tsdr-persistence-hbase.properties

```
metric-persistency=false
log-persistency=true
binary-persistency=true
```

- tsdr-persistence-cassandra.properties

```
metric-psersistency=true
log-persistency=false
binary-persistency=false
```

TTP CLI Tools User Guide

Overview

Table Type Patterns are a specification developed by the Open Networking Foundation to enable the description and negotiation of subsets of the OpenFlow protocol. This is particularly useful for hardware switches that support OpenFlow as it enables the to describe what features they do (and thus also what features they do not) support. More details can be found in the full specification listed on the OpenFlow specifications page.

TTP CLI Tools Architecture

The TTP CLI Tools use the TTP Model and the YANG Tools/RESTCONF codecs to translate between the Data Transfer Objects (DTOs) and JSON/XML.

User Network Interface Manager Plug-in (Unimgr)

Overview

The User Network Interface Manager (Unimgr) is an experimental/proof of concept (PoC) project formed to initiate the development of data models and APIs facilitating the use by software applications and/or service orchestrators of OpenDaylight to configure and provision connectivity services, in particular Carrier Ethernet services as defined by Metro Ethernet Forum (MEF), in physical or virtual network elements.

MEF as defined the LSO Reference Architecture for the management and control of domains and entities that enable cooperative LSO capabilities across one or more service provider networks. The architecture also identifies the Management Interface Reference Points (LSO Reference Points), the logical points of interaction between specific functional management components. These LSO Reference Points are further defined by interface profiles and instantiated by APIs.

The LSO High Level Management Reference Architecture is shown below. Note that this is a functional architecture that does not describe how the management components are implemented (e.g., single vs. multiple instances), but rather identifies management components that provide logical functionality as well as the points of interaction among them.

Unimgr provides support for both the Legato as well as the Presto interfaces. These interface, and the APIs associated with them, are defined by YANG models developed within MEF in collaboration with ONF and IETF. For the Boron release, these are as follows:

Legato YANG modules: https://git.opendaylight.org/gerrit/gitweb?p=unimgr.git;a=tree;f=legato-api/src/main/yang;hb=refs/heads/stable/boron

Presto YANG modules: https://git.opendaylight.org/gerrit/gitweb?p=unimgr.git;a=tree;f=presto-api/src/main/yang;hb=refs/heads/stable/boron

An application/user can interact with Unimgr at either the service orchestration layer (Legato) or the network resource provisioning layer (Presto).

Unimgr Components

Unimgr is comprised of the following OpenDaylight Karaf features:

odl-unimgr-api	OpenDaylight :: UniMgr :: api
odl-unimgr	OpenDaylight :: UniMgr
odl-unimgr-console	OpenDaylight :: UniMgr :: CLI
odl-unimgr-rest	OpenDaylight :: UniMgr :: REST
odl-unimgr-ui	OpenDaylight :: UniMgr :: UI

Installing Unimgr

After launching OpenDaylight, install the feature for Unimgr. From the karaf command prompt execute the following command:

```
$ feature:install odl-unimgr-ui
```

Explore and exercise the Unimgr REST API

To see the Unimgr API, browse to this URL: http://localhost:8181/apidoc/explorer/index.html

Replace localhost with the IP address or hostname where OpenDaylight is running if you are not running OpenDaylight locally on your machine.

See also the Unimgr Developer Guide for a full listing of the API.

Unified Secure Channel

This document describes how to use the Unified Secure Channel (USC) feature in OpenDaylight. This document contains configuration, administration, and management sections for the feature.

Overview

In enterprise networks, more and more controller and network management systems are being deployed remotely, such as in the cloud. Additionally, enterprise networks are becoming more heterogeneous - branch, IoT, wireless (including cloud access control). Enterprise customers want a converged network controller and management system solution. This feature is intended for device and network administrators looking to use unified secure channels for their systems.

USC Channel Architecture

- USC Agent
  - The USC Agent provides proxy and agent functionality on top of all standard protocols supported by the device. It initiates call-home with the controller, maintains live connections with with the controller, acts as a demuxer/muxer for packets with the USC header, and authenticates the controller.
- USC Plugin
  - The USC Plugin is responsible for communication between the controller and the USC agent . It responds to call-home with the controller, maintains live connections with the devices, acts as a muxer/demuxer for packets with the USC header, and provides support for TLS/DTLS.
- USC Manager

  - The USC Manager handles configurations, high availability, security, monitoring, and clustering support for USC.
- USC UI
  - The USC UI is responsible for displaying a graphical user interface representing the state of USC in the OpenDaylight DLUX UI.

Installing USC Channel

To install USC, download OpenDaylight and use the Karaf console to install the following feature:

odl-usc-channel-ui

Configuring USC Channel

This section gives details about the configuration settings for various components in USC.

The USC configuration files for the Karaf distribution are located in distribution/karaf/target/assembly/etc/usc

- certificates
  - The certificates folder contains the client key, pem, and rootca files as is necessary for security.
- akka.conf
  - This file contains configuration related to clustering. Potential configuration properties can be found on the akka website at http://doc.akka.io
- usc.properties
  - This file contains configuration related to USC. Use this file to set the location of certificates, define the source of additional akka configurations, and assign default settings to the USC behavior.

Administering or Managing USC Channel

After installing the odl-usc-channel-ui feature from the Karaf console, users can administer and manage USC channels from the the UI or APIDOCS explorer.

Go to http://\protect\T1\textdollar\protect\T1\textbraceleftipaddress\protect\T1\textbraceright:8181/index.html, sign in, and click on the USC side menu tab. From there, users can view the state of USC channels.

Go to http://\protect\T1\textdollar\protect\T1\textbraceleftipaddress\protect\T1\textbraceright:8181/apidoc/explorer/index.html, sign in, and expand the usc-channel panel. From there, users can execute various API calls to test their USC deployment such as add-channel, delete-channel, and view-channel.

Tutorials

Below are tutorials for USC Channel

Viewing USC Channel

The purpose of this tutorial is to view USC Channel

Overview

This tutorial walks users through the process of viewing the USC Channel environment topology including established channels connecting the controllers and devices in the USC topology.

Prerequisites

For this tutorial, we assume that a device running a USC agent is already installed.

Instructions

- Run the OpenDaylight distribution and install odl-usc-channel-ui from the Karaf console.
- Go to http://\protect\T1\textdollar\protect\T1\textbraceleftipaddress\protect\T1\textbraceright:8181/apidoc/explorer/index.html
- Execute add-channel with the following json data:
  - {"input":{"channel":{"hostname":"127.0.0.1","port":1068,"remote":false}}}
- Go to http://\protect\T1\textdollar\protect\T1\textbraceleftipaddress\protect\T1\textbraceright:8181/index.html
- Click on the USC side menu tab.
- The UI should display a table including the added channel from step 3.

Usecplugin-AAA User Guide

The Usecplugin User Guide contains information about configuration, administration, management, using and troubleshooting the feature.

Overview

AAA plugin provides authorization, authentication and accounting services to OpenDaylight. A user logs in to OpenDaylight through the username and password provided by AAA plugin. Usecplugin-AAA collects and stores information about both successful and failed login attempts to OpenDaylight.

Usecplugin-AAA Architecture

AAA plugin creates log messages about successful and failed login attempts to OpenDaylight. Usecplugin-AAA continuously reads this log file and checks for either successful and failed attempt information. Whenever Usecpluin-AAA identifies a new attempt entry in the log file it is stored in YANG Data Store and its own log file.

Administering or Managing Usecplugin-AAA

- Install feature `odl-usecplugin-aaa`
- Enable odl-aaa log using command `log:set DEBUG org.opendaylight.aaa.shiro.filters`
- Login to the RESTCONF documentation.
- Check operational datastore for login attempts.

- POST URI :: http://localhost:8181/restconf/operations/usecpluginaaa:attemptFromIP
  - Sample Input :: {"usecpluginaaa:input":{"ScrIP":"10.0.0.1"}}
- POST URI :: http://localhost:8181/restconf/operations/usecpluginaaa:attemptOnDateTime
  - Sample Input :: {"usecpluginaaa:input":{"dateTime":"2016-07-27 14:11:18"}}

Usecplugin-OpenFlow User Guide

The Usecplugin-OpenFlow User Guide contains information about configuration, administration, management, using and troubleshooting the feature.

Overview

Usecplugin-OpenFlow collects information about potential OpenFlow Packet_In attacks to OpenDaylight. A threshold (water mark) can be set for the Packet_In rate which when breached will trigger Packet_In message information collection.

Usecplugin Architecture

Usecplugin listens on OpenFlow southbound interface for Packet_In messages. When the rate of Packet_In breaches the high water mark the application parses the message for header information which is subsequently stored in YANG Data Store and a log file. Usecplugin has PacketHandler class that implements the PacketProcessing interface to override the OnPacketReceived notification by which the application is notified of Packet_In messages.

Configuring Usecplugin-OpenFlow

Install the Usecplugin-OpenFlow feautre in OpenDaylight with the `feature:install odl-usecplugin-openflow` at the Karaf CLI.

A user can set the low water mark and high water mark for Packet_In rates as well as number of samples for checking the time interval to calculate Packet_In rate.

URI http://localhost:8181/apidoc/explorer/index.html#!/usecplugin(2015-01-05)

High Water Mark Configuration

PUT URI http://localhost:8181/restconf/config/usecplugin:sample-data-hwm/

Sample Input
```
{"usecplugin:sample-data-hwm": { "samples":"3000",
    "highWaterMark":"3000"}}
```

Low Water Mark Configuration

PUT URI http://localhost:8181/restconf/config/usecplugin:sample-data-lwm/

Sample Input
```
{"usecplugin:sample-data-lwm": { "samples-lwm":"2000",
    "lowWaterMark-lwm":"2000"}}
```

Administering or Managing Usecplugin-OpenFlow

Use RPC POST APIs in the following format for getting the attack related information.

attackID

URI http://localhost:8181/restconf/operations/usecplugin:attackID

Sample Input `{"usecplugin:input": { "NodeID":"openflow:1"}}`

attacksFromIP

URI http://localhost:8181/restconf/operations/usecplugin:attacksFromIP

Sample Input `{"usecplugin:input": { "SrcIP":"10.0.0.1"}}`

attacksToIP

URI http://localhost:8181/restconf/operations/usecplugin:attacksToIP

Sample Input `{"usecplugin:input": { "DstIP":"10.0.0.2"}}`

Virtual Tenant Network (VTN)

VTN Overview

OpenDaylight Virtual Tenant Network (VTN) is an application that provides multi-tenant virtual network on an SDN controller.

Conventionally, huge investment in the network systems and operating expenses are needed because the network is configured as a silo for each department and system. So, various network appliances must be installed for each tenant and those boxes cannot be shared with others. It is a heavy work to design, implement and operate the entire complex network.

The uniqueness of VTN is a logical abstraction plane. This enables the complete separation of logical plane from physical plane. Users can design and deploy any desired network without knowing the physical network topology or bandwidth restrictions.

VTN allows the users to define the network with a look and feel of conventional L2/L3 network. Once the network is designed on VTN, it will automatically be mapped into underlying physical network, and then configured on the individual switch leveraging SDN control protocol. The definition of logical plane makes it possible not only to hide the complexity of the underlying network but also to better manage network resources. It achieves reducing reconfiguration time of network services and minimizing network configuration errors.

It is implemented as two major components

- *VTN Manager*
- *VTN Coordinator*

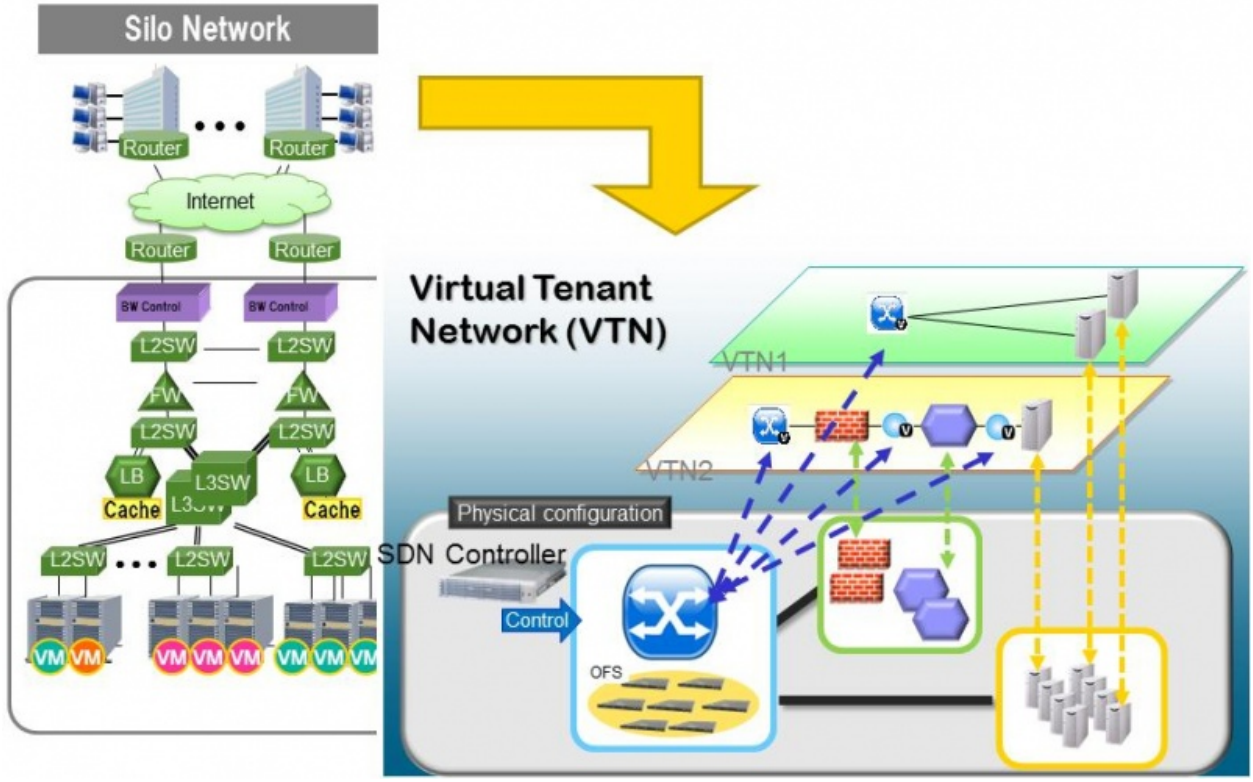

Fig. 107: VTN Overview

VTN Manager

An OpenDaylight Plugin that interacts with other modules to implement the components of the VTN model. It also provides a REST interface to configure VTN components in OpenDaylight. VTN Manager is implemented as one plugin to the OpenDaylight. This provides a REST interface to create/update/delete VTN components. The user command in VTN Coordinator is translated as REST API to VTN Manager by the OpenDaylight Driver component. In addition to the above mentioned role, it also provides an implementation to the OpenStack L2 Network Functions API.

Features Overview

- **odl-vtn-manager** provides VTN Manager's JAVA API.
- **odl-vtn-manager-rest** provides VTN Manager's REST API.
- **odl-vtn-manager-neutron** provides the integration with Neutron interface.

REST API

VTN Manager provides REST API for virtual network functions.

Here is an example of how to create a virtual tenant network.

```
curl --user "admin":"admin" -H "Accept: application/json" -H \
"Content-type: application/json" -X POST \
```

(continues on next page)

(continued from previous page)

```
http://localhost:8181/restconf/operations/vtn:update-vtn \
-d '{"input":{"tenant-name":"vtn1"}}'
```

You can check the list of all tenants by executing the following command.

```
curl --user "admin":"admin" -H "Accept: application/json" -H \
"Content-type: application/json" -X GET \
http://localhost:8181/restconf/operational/vtn:vtns
```

REST Conf documentation for VTN Manager, please refer to: https://nexus.opendaylight.org/content/sites/site/org.opendaylight.vtn/boron/manager.model/apidocs/index.html

VTN Coordinator

The VTN Coordinator is an external application that provides a REST interface for an user to use OpenDaylight VTN Virtualization. It interacts with VTN Manager plugin to implement the user configuration. It is also capable of multiple OpenDaylight orchestration. It realizes Virtual Tenant Network (VTN) provisioning in OpenDaylight instances. In the OpenDaylight architecture VTN Coordinator is part of the network application, orchestration and services layer. VTN Coordinator will use the REST interface exposed by the VTN Manger to realize the virtual network using OpenDaylight. It uses OpenDaylight APIs (REST) to construct the virtual network in OpenDaylight instances. It provides REST APIs for northbound VTN applications and supports virtual networks spanning across multiple OpenDaylight by coordinating across OpenDaylight.

For VTN Coordinator REST API, please refer to: https://wiki.opendaylight.org/view/OpenDaylight_Virtual_Tenant_Network_%28VTN%29:VTN_Coordinator:RestApi

Network Virtualization Function

The user first defines a VTN. Then, the user maps the VTN to a physical network, which enables communication to take place according to the VTN definition. With the VTN definition, L2 and L3 transfer functions and flow-based traffic control functions (filtering and redirect) are possible.

Virtual Network Construction

The following table shows the elements which make up the VTN. In the VTN, a virtual network is constructed using virtual nodes (vBridge, vRouter) and virtual interfaces and links. It is possible to configure a network which has L2 and L3 transfer function, by connecting the virtual intrefaces made on virtual nodes via virtual links.

vBridge	The logical representation of L2 switch function.
vRouter	The logical representation of router function.
vTep	The logical representation of Tunnel End Point - TEP.
vTunnel	The logical representation of Tunnel.
vBypass	The logical representation of connectivity between controlled networks.
Virtual interface	The representation of end point on the virtual node.
Virtual Linkv(vLink)	The logical representation of L1 connectivity between virtual interfaces.

The following figure shows an example of a constructed virtual network. VRT is defined as the vRouter, BR1 and BR2 are defined as vBridges. interfaces of the vRouter and vBridges are connected using vLinks.

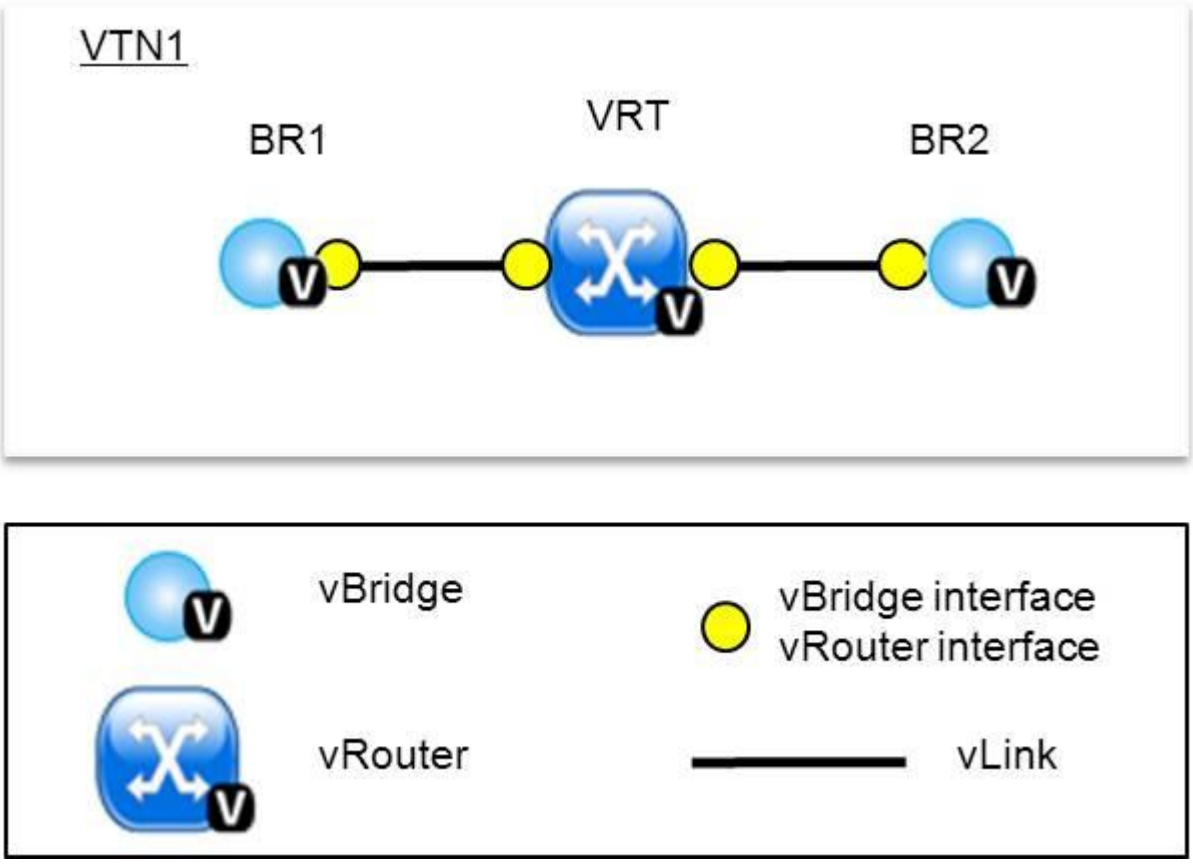

Fig. 108: VTN Construction

Mapping of Physical Network Resources

Map physical network resources to the constructed virtual network. Mapping identifies which virtual network each packet transmitted or received by an OpenFlow switch belongs to, as well as which interface in the OpenFlow switch transmits or receives that packet. There are two mapping methods. When a packet is received from the OFS, port mapping is first searched for the corresponding mapping definition, then VLAN mapping is searched, and the packet is mapped to the relevant vBridge according to the first matching mapping.

Port mapping	Maps physical network resources to an interface of vBridge using Switch ID, Port ID and VLAN ID of the incoming L2 frame. Untagged frame mapping is also supported.
VLAN mapping	Maps physical network resources to a vBridge using VLAN ID of the incoming L2 frame.Maps physical resources of a particular switch to a vBridge using switch ID and VLAN ID of the incoming L2 frame.
MAC mapping	Maps physical resources to an interface of vBridge using MAC address of the incoming L2 frame(The initial contribution does not include this method).

VTN can learn the terminal information from a terminal that is connected to a switch which is mapped to VTN. Further, it is possible to refer that terminal information on the VTN.

- Learning terminal information VTN learns the information of a terminal that belongs to VTN. It will store the MAC address and VLAN ID of the terminal in relation to the port of the switch.
- Aging of terminal information Terminal information, learned by the VTN, will be maintained until the packets from terminal keep flowing in VTN. If the terminal gets disconnected from the VTN, then the aging timer will start clicking and the terminal information will be maintained till timeout.

The following figure shows an example of mapping. An interface of BR1 is mapped to port GBE0/1 of OFS1 using port mapping. Packets received from GBE0/1 of OFS1 are regarded as those from the corresponding interface of BR1. BR2 is mapped to VLAN 200 using VLAN mapping. Packets with VLAN tag 200 received from any ports of any OFSs are regarded as those from an interface of BR2.

vBridge Functions

The vBridge provides the bridge function that transfers a packet to the intended virtual port according to the destination MAC address. The vBridge looks up the MAC address table and transmits the packet to the corresponding virtual interface when the destination MAC address has been learned. When the destination MAC address has not been learned, it transmits the packet to all virtual interfaces other than the receiving port (flooding). MAC addresses are learned as follows.

- MAC address learning The vBridge learns the MAC address of the connected host. The source MAC address of each received frame is mapped to the receiving virtual interface, and this MAC address is stored in the MAC address table created on a per-vBridge basis.
- MAC address aging The MAC address stored in the MAC address table is retained as long as the host returns the ARP reply. After the host is disconnected, the address is retained until the aging timer times out. To have the vBridge learn MAC addresses statically, you can register MAC addresses manually.

vRouter Functions

The vRouter transfers IPv4 packets between vBridges. The vRouter supports routing, ARP learning, and ARP aging functions. The following outlines the functions.

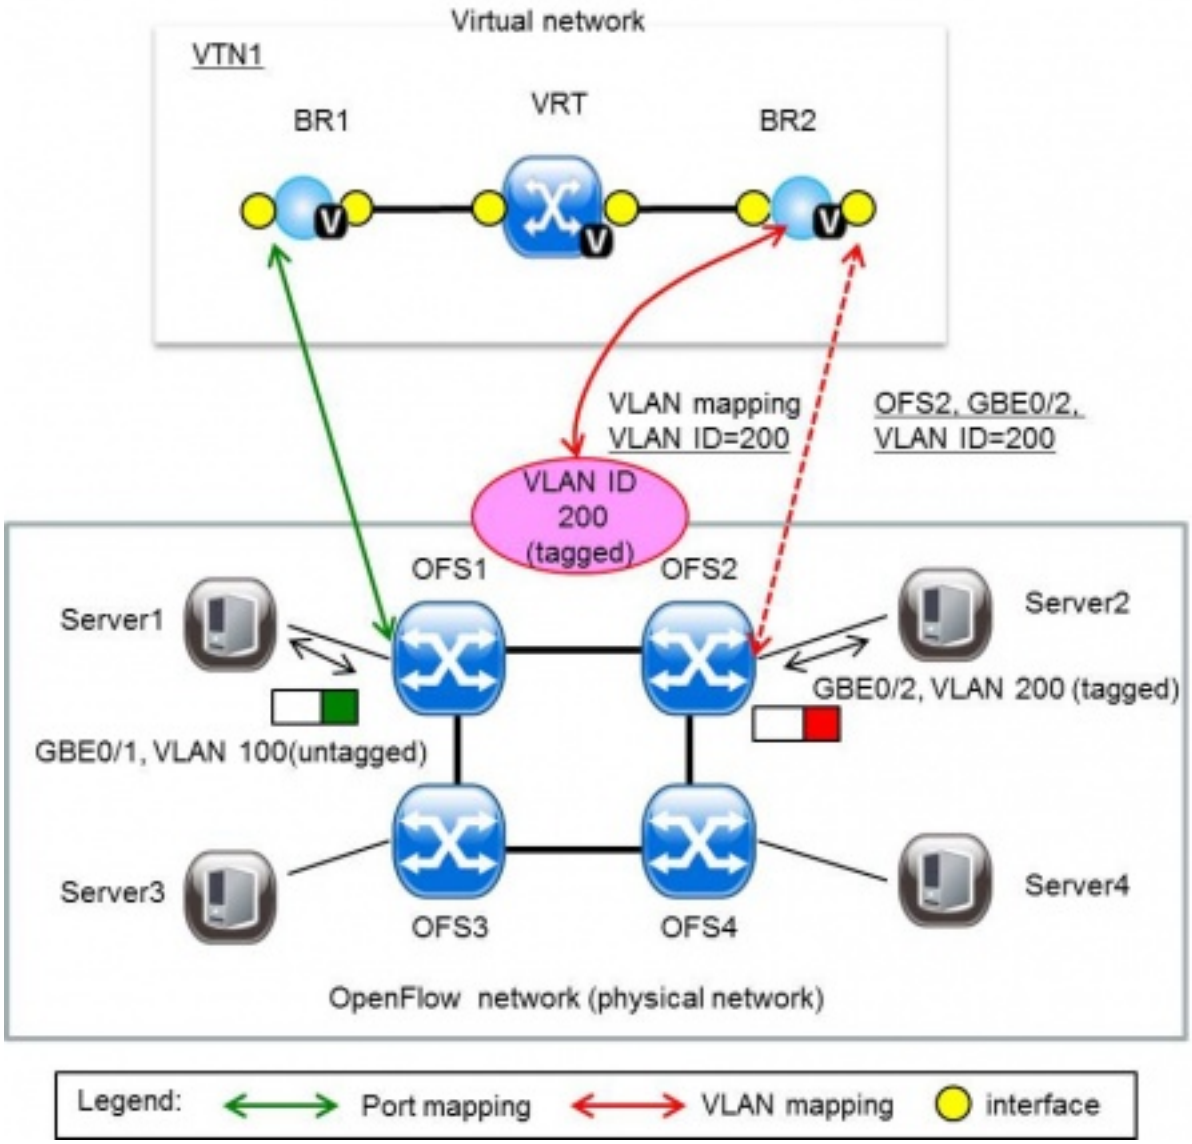

Fig. 109: VTN Mapping

- Routing function When an IP address is registered with a virtual interface of the vRouter, the default routing information for that interface is registered. It is also possible to statically register routing information for a virtual interface.
- ARP learning function The vRouter associates a destination IP address, MAC address and a virtual interface, based on an ARP request to its host or a reply packet for an ARP request, and maintains this information in an ARP table prepared for each routing domain. The registered ARP entry is retained until the aging timer, described later, times out. The vRouter transmits an ARP request on an individual aging timer basis and deletes the associated entry from the ARP table if no reply is returned. For static ARP learning, you can register ARP entry information manually.
- DHCP relay agent function The vRouter also provides the DHCP relay agent function.

Flow Filter Functions

Flow Filter function is similar to ACL. It is possible to allow or prohibit communication with only certain kind of packets that meet a particular condition. Also, it can perform a processing called Redirection - WayPoint routing, which is different from the existing ACL. Flow Filter can be applied to any interface of a vNode within VTN, and it is possible to the control the packets that pass interface. The match conditions that could be specified in Flow Filter are as follows. It is also possible to specify a combination of multiple conditions.

- Source MAC address
- Destination MAC address
- MAC ether type
- VLAN Priority
- Source IP address
- Destination IP address
- DSCP
- IP Protocol
- TCP/UDP source port
- TCP/UDP destination port
- ICMP type
- ICMP code

The types of Action that can be applied on packets that match the Flow Filter conditions are given in the following table. It is possible to make only those packets, which match a particular condition, to pass through a particular server by specifying Redirection in Action. E.g., path of flow can be changed for each packet sent from a particular terminal, depending upon the destination IP address. VLAN priority control and DSCP marking are also supported.

Action	Function
Pass	Pass particular packets matching the specified conditions.
Drop	Discards particular packets matching the specified conditions.
Redirection	Redirects the packet to a desired virtual interface. Both Transparent Redirection (not changing MAC address) and Router Redirection (changing MAC address) are supported.

The following figure shows an example of how the flow filter function works.

If there is any matching condition specified by flow filter when a packet being transferred within a virtual network goes through a virtual interface, the function evaluates the matching condition to see whether the packet matches it.

If the packet matches the condition, the function applies the matching action specified by flow filter. In the example shown in the figure, the function evaluates the matching condition at BR1 and discards the packet if it matches the condition.

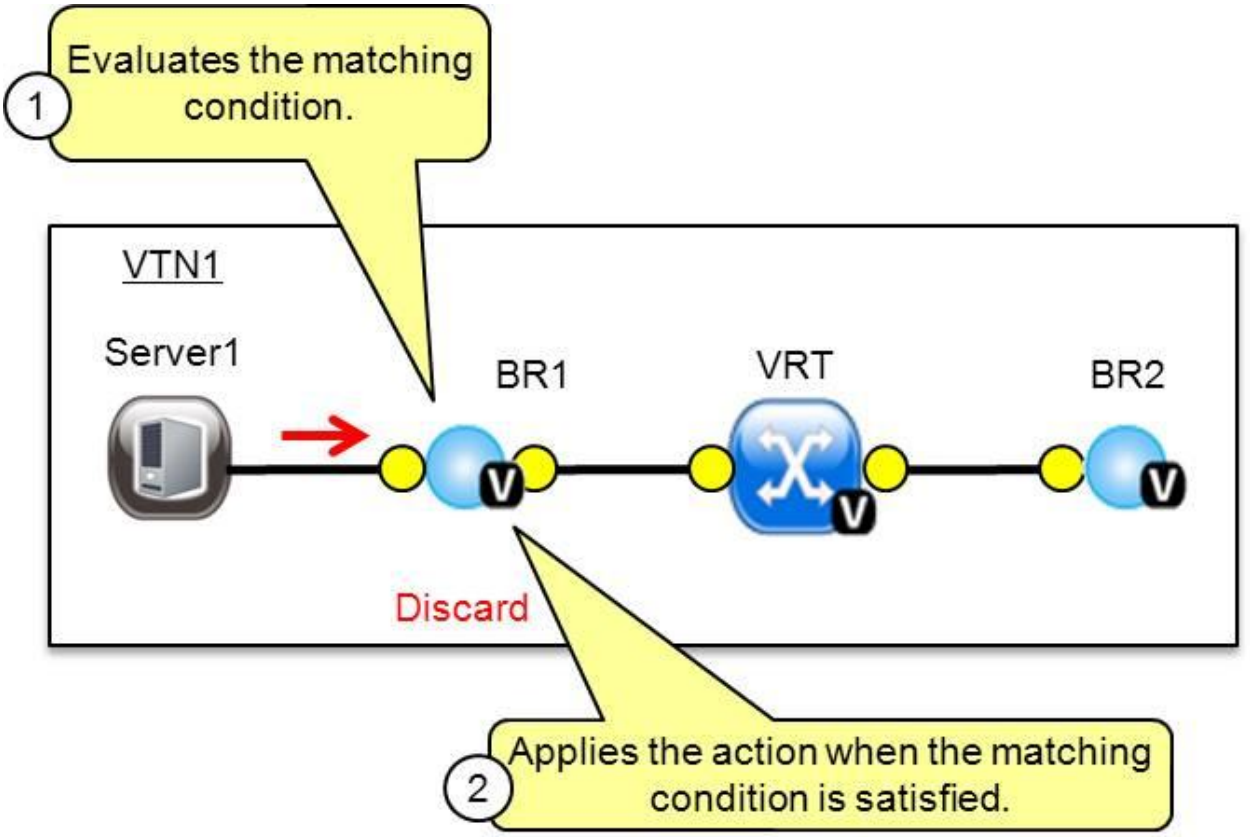

Fig. 110: VTN FlowFilter

Multiple SDN Controller Coordination

With the network abstractions, VTN enables to configure virtual network across multiple SDN controllers. This provides highly scalable network system.

VTN can be created on each SDN controller. If users would like to manage those multiple VTNs with one policy, those VTNs can be integrated to a single VTN.

As a use case, this feature is deployed to multi data center environment. Even if those data centers are geographically separated and controlled with different controllers, a single policy virtual network can be realized with VTN.

Also, one can easily add a new SDN Controller to an existing VTN or delete a particular SDN Controller from VTN.

In addition to this, one can define a VTN which covers both OpenFlow network and Overlay network at the same time.

Flow Filter, which is set on the VTN, will be automatically applied on the newly added SDN Controller.

Coordination between OpenFlow Network and L2/L3 Network

It is possible to configure VTN on an environment where there is mix of L2/L3 switches as well. L2/L3 switch will be shown on VTN as vBypass. Flow Filter or policing cannot be configured for a vBypass. However, it is possible to

treat it as a virtual node inside VTN.

Virtual Tenant Network (VTN) API

VTN provides Web APIs. They are implemented by REST architecture and provide the access to resources within VTN that are identified by URI. User can perform the operations like GET/PUT/POST/DELETE against the virtual network resources (e.g. vBridge or vRouter) by sending a message to VTN through HTTPS communication in XML or JSON format.

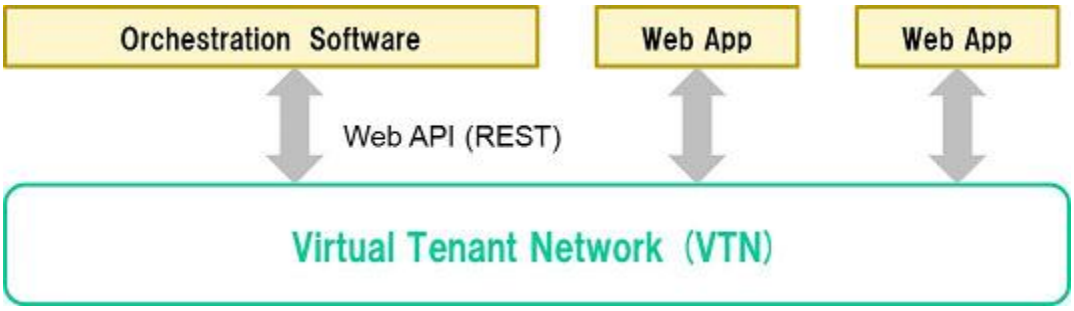

Fig. 111: VTN API

Function Outline

VTN provides following operations for various network resources.

Resources	GET	POST	PUT	DELETE
VTN	Yes	Yes	Yes	Yes
vBridge	Yes	Yes	Yes	Yes
vRouter	Yes	Yes	Yes	Yes
vTep	Yes	Yes	Yes	Yes
vTunnel	Yes	Yes	Yes	Yes
vBypass	Yes	Yes	Yes	Yes
vLink	Yes	Yes	Yes	Yes
Interface	Yes	Yes	Yes	Yes
Port map	Yes	No	Yes	Yes
Vlan map	Yes	Yes	Yes	Yes
Flowfilter (ACL/redirect)	Yes	Yes	Yes	Yes
Controller information	Yes	Yes	Yes	Yes
Physical topology information	Yes	No	No	No
Alarm information	Yes	No	No	No

Example usage

The following is an example of the usage to construct a virtual network.

- Create VTN

```
 curl --user admin:adminpass -X POST -H 'content-type: application/json'  \
-d '{"vtn":{"vtn_name":"VTN1"}}' http://172.1.0.1:8083/vtn-webapi/vtns.json
```

- Create Controller Information

```
 curl --user admin:adminpass -X POST -H 'content-type: application/json'  \
-d '{"controller": {"controller_id":"CONTROLLER1","ipaddr":"172.1.0.1","type":"odc",
↪"username":"admin", \
"password":"admin","version":"1.0"}}' http://172.1.0.1:8083/vtn-webapi/controllers.
↪json
```

- Create vBridge under VTN

```
curl --user admin:adminpass -X POST -H 'content-type: application/json' \
-d '{"vbridge":{"vbr_name":"VBR1","controller_id": "CONTROLLER1","domain_id":
↪"(DEFAULT)"}}' \
http://172.1.0.1:8083/vtn-webapi/vtns/VTN1/vbridges.json
```

- Create the interface under vBridge

```
curl --user admin:adminpass -X POST -H 'content-type: application/json' \
-d '{"interface":{"if_name":"IF1"}}' http://172.1.0.1:8083/vtn-webapi/vtns/VTN1/
↪vbridges/VBR1/interfaces.json
```

VTN OpenStack Configuration

This guide describes how to set up OpenStack for integration with OpenDaylight Controller.

While OpenDaylight Controller provides several ways to integrate with OpenStack, this guide focus on the way which uses VTN features available on OpenDaylight. In the integration, VTN Manager work as network service provider for OpenStack.

VTN Manager features, enable OpenStack to work in pure OpenFlow environment in which all switches in data plane are OpenFlow switch.

Requirements

- OpenDaylight Controller. (VTN features must be installed)
- OpenStack Control Node.
- OpenStack Compute Node.
- OpenFlow Switch like mininet(Not Mandatory).

The VTN features support multiple OpenStack nodes. You can deploy multiple OpenStack Compute Nodes. In management plane, OpenDaylight Controller, OpenStack nodes and OpenFlow switches should communicate with each other. In data plane, Open vSwitches running in OpenStack nodes should communicate with each other through a physical or logical OpenFlow switches. The core OpenFlow switches are not mandatory. Therefore, you can directly connect to the Open vSwitch's.

Sample Configuration

Below steps depicts the configuration of single OpenStack Control node and OpenStack Compute node setup. Our test setup is as follows

Server Preparation

- Install Ubuntu 14.04 LTS in two servers (OpenStack Control node and Compute node respectively)

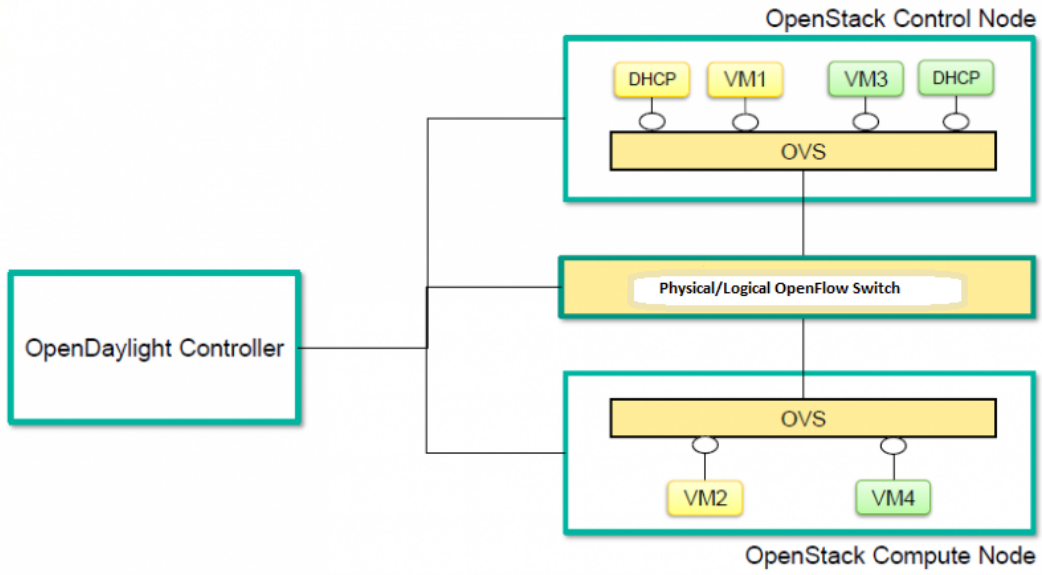

Fig. 112: Openstack Overview

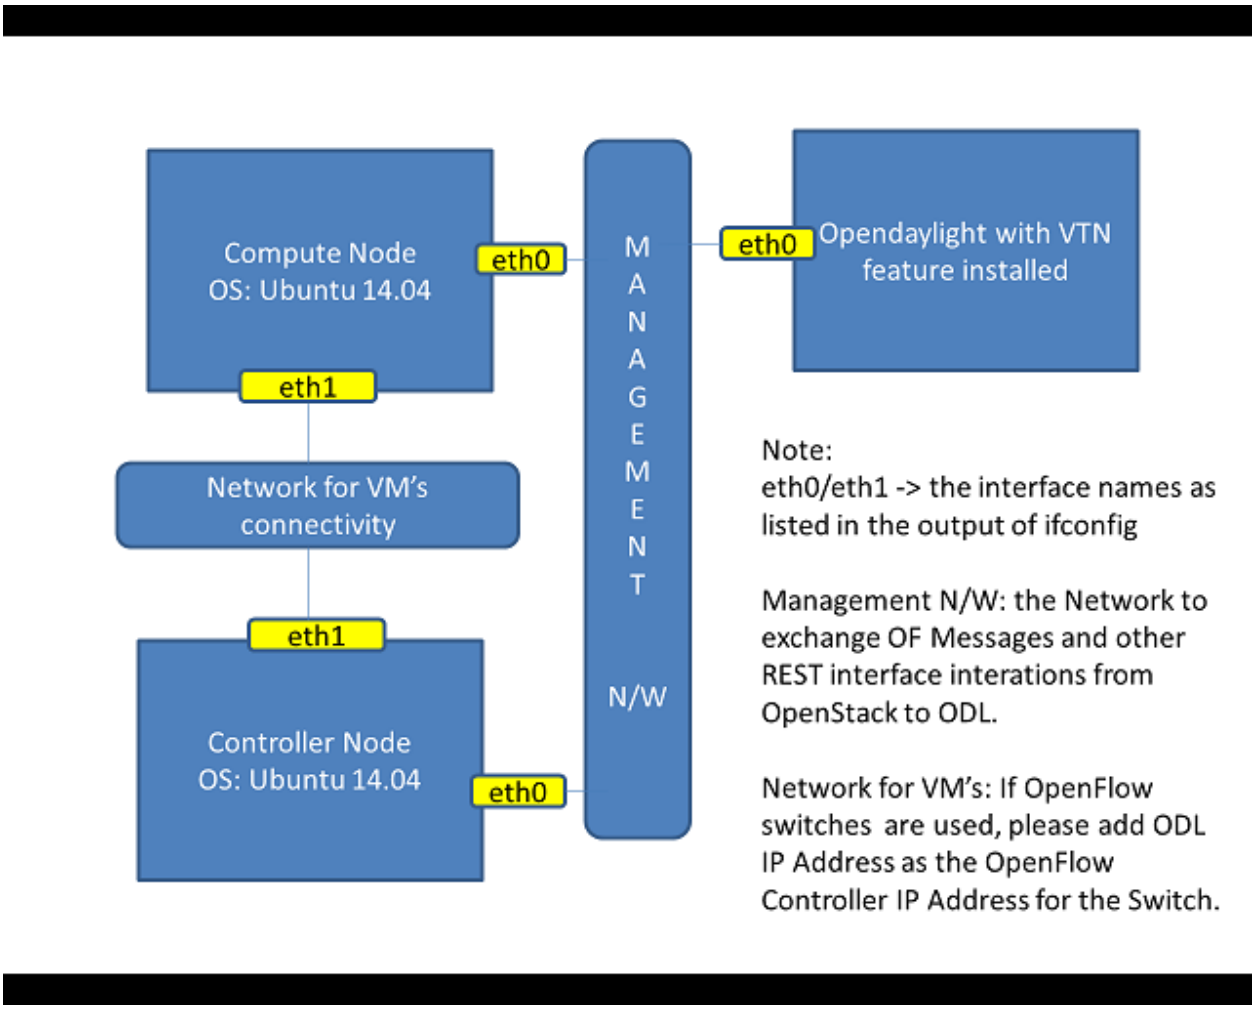

Fig. 113: LAB Setup

- While installing, Ubuntu mandates creation of a User, we created the user "stack"(We will use the same user for running devstack)
- Proceed with the below mentioned User Settings and Network Settings in both the Control and Compute nodes.

User Settings for devstack - Login to both servers - Disable Ubuntu Firewall

```
sudo ufw disable
```

- Install the below packages (optional, provides ifconfig and route coammnds, handy for debugging!!)

```
sudo apt-get install net-tools
```

- Edit sudo vim /etc/sudoers and add an entry as follows

```
stack ALL=(ALL) NOPASSWD: ALL
```

Network Settings - Checked the output of ifconfig -a, two interfaces were listed eth0 and eth1 as indicated in the image above. - We had connected eth0 interface to the Network where OpenDaylight is reachable. - eth1 interface in both servers were connected to a different network to act as data plane for the VM's created using the OpenStack. - Manually edited the file : sudo vim /etc/network/interfaces and made entries as follows

```
 stack@ubuntu-devstack:~/devstack$ cat /etc/network/interfaces
 # This file describes the network interfaces available on your system
 # and how to activate them. For more information, see interfaces(5).
 # The loop-back network interface
 auto lo
 iface lo inet loopback
 # The primary network interface
 auto eth0
 iface eth0 inet static
      address <IP_ADDRESS_TO_REACH_ODL>
      netmask <NET_MASK>
      broadcast <BROADCAST_IP_ADDRESS>
      gateway <GATEWAY_IP_ADDRESS>
auto eth1
iface eth1 inet static
     address <IP_ADDRESS_UNIQ>
     netmask <NETMASK>
```

Note: Please ensure that the eth0 interface is the default route and it is able to reach the ODL_IP_ADDRESS NOTE: The entries for eth1 are not mandatory, If not set, we may have to manually do "ifup eth1" after the stacking is complete to activate the interface

---

Finalize the user and network settings - Please reboot both nodes after the user and network settings to have the network settings applied to the network - Login again and check the output of ifconfig to ensure that both interfaces are listed

OpenDaylight Settings and Execution

VTN Configuration for OpenStack Integration:

- VTN uses the configuration parameters from "90-vtn-neutron.xml" file for the OpenStack integration.
- These values will be set for the OpenvSwitch, in all the participating OpenStack nodes.

- A configuration file "90-vtn-neutron.xml" will be generated automatically by following the below steps,
- Download the latest Boron karaf distribution from the below link,

```
http://www.opendaylight.org/software/downloads
```

- cd "distribution-karaf-0.5.0-Boron" and run karaf by using the following command "./bin/karaf".
- Install the below feature to generate "90-vtn-neutron.xml"

```
feature:install odl-vtn-manager-neutron
```

- Logout from the karaf console and Check "90-vtn-neutron.xml" file from the following path "distribution-karaf-0.5.0-Boron/etc/opendaylight/karaf/".
- The contents of "90-vtn-neutron.xml" should be as follows:

bridgename=br-int portname=eth1 protocols=OpenFlow13 failmode=secure

- The values of the configuration parameters must be changed based on the user environment.
- Especially, "portname" should be carefully configured, because if the value is wrong, OpenDaylight fails to forward packets.
- Other parameters works fine as is for general use cases.
  - bridgename
    - The name of the bridge in Open vSwitch, that will be created by OpenDaylight Controller.
    - It must be "br-int".
  - portname
    - The name of the port that will be created in the vbridge in Open vSwitch.
    - This must be the same name of the interface of OpenStack Nodes which is used for interconnecting OpenStack Nodes in data plane.(in our case:eth1)
    - By default, if 90-vtn-neutron.xml is not created, VTN uses ens33 as portname.
  - protocols
    - OpenFlow protocol through which OpenFlow Switch and Controller communicate.
    - The values can be OpenFlow13 or OpenFlow10.
  - failmode
    - The value can be "standalone" or "secure".
    - Please use "secure" for general use cases.

Start ODL Controller

- Please refer to the Installation Pages to run ODL with VTN Feature enabled.
- After running ODL Controller, please ensure ODL Controller listens to the ports:6633,6653, 6640 and 8080
- Please allow the ports in firewall for the devstack to be able to communicate with ODL Controller.

Note:

- 6633/6653 - OpenFlow Ports

- 6640 - OVS Manager Port
- 8080 - Port for REST API

Devstack Setup

Get Devstack (All nodes)

- Install git application using
  - sudo apt-get install git
- Get devstack
  - git clone https://git.openstack.org/openstack-dev/devstack;
- Switch to stable/Juno Version branch
  - cd devstack

```
git checkout stable/juno
```

Note: If you want to use stable/kilo Version branch, Please execute the below command in devstack folder

```
git checkout stable/kilo
```

Note: If you want to use stable/liberty Version branch, Please execute the below command in devstack folder

```
git checkout stable/liberty
```

Stack Control Node

- local.conf:
- cd devstack in the controller node
- Copy the contents of local.conf for juno (devstack control node) from https://wiki.opendaylight.org/view/OpenDaylight_Virtual_Tenant_Network_(VTN):Scripts:devstack and save it as "local.conf" in the "devstack".
- Copy the contents of local.conf for kilo and liberty (devstack control node) from https://wiki.opendaylight.org/view/OpenDaylight_Virtual_Tenant_Network_(VTN):Scripts:devstack_post_juno_versions and save it as "local.conf" in the "devstack".
- Please modify the IP Address values as required.
- Stack the node

```
./stack.sh
```

Verify Control Node stacking

- stack.sh prints out Horizon is now available at http://<CONTROL_NODE_IP_ADDRESS>:8080/
- Execute the command *sudo ovs-vsctl show* in the control node terminal and verify if the bridge *br-int* is created.
- Typical output of the ovs-vsctl show is indicated below:

```
e232bbd5-096b-48a3-a28d-ce4a492d4b4f
   Manager "tcp:192.168.64.73:6640"
       is_connected: true
   Bridge br-int
       Controller "tcp:192.168.64.73:6633"
           is_connected: true
       fail_mode: secure
       Port "eth1"
          Interface "eth1"
   ovs_version: "2.0.2"
```

Stack Compute Node

- local.conf:
- cd devstack in the controller node
- Copy the contents of local.conf for juno (devstack compute node) from https://wiki.opendaylight.org/view/OpenDaylight_Virtual_Tenant_Network_(VTN):Scripts:devstack and save it as "local.conf" in the "devstack".
- Copy the contents of local.conf file for kilo and liberty (devstack compute node) from https://wiki.opendaylight.org/view/OpenDaylight_Virtual_Tenant_Network_(VTN):Scripts:devstack_post_juno_versions and save it as "local.conf" in the "devstack".
- Please modify the IP Address values as required.
- Stack the node

```
./stack.sh
```

Verify Compute Node Stacking

- stack.sh prints out This is your host ip: <COMPUTE_NODE_IP_ADDRESS>
- Execute the command *sudo ovs-vsctl show* in the control node terminal and verify if the bridge *br-int* is created.
- The output of the ovs-vsctl show will be similar to the one seen in control node.

Additional Verifications

- Please visit the OpenDaylight DLUX GUI after stacking all the nodes, http://<ODL_IP_ADDRESS>:8181/index.html. The switches, topology and the ports that are currently read can be validated.

```
http://<controller-ip>:8181/index.html
```

Tip: If the interconnected between the Open vSwitch is not seen, Please bring up the interface for the dataplane manually using the below comamnd

```
ifup <interface_name>
```

- Please Accept Promiscuous mode in the networks involving the interconnect.

Create VM from Devstack Horizon GUI

- Login to http://<CONTROL_NODE_IP>:8080/ to check the horizon GUI.

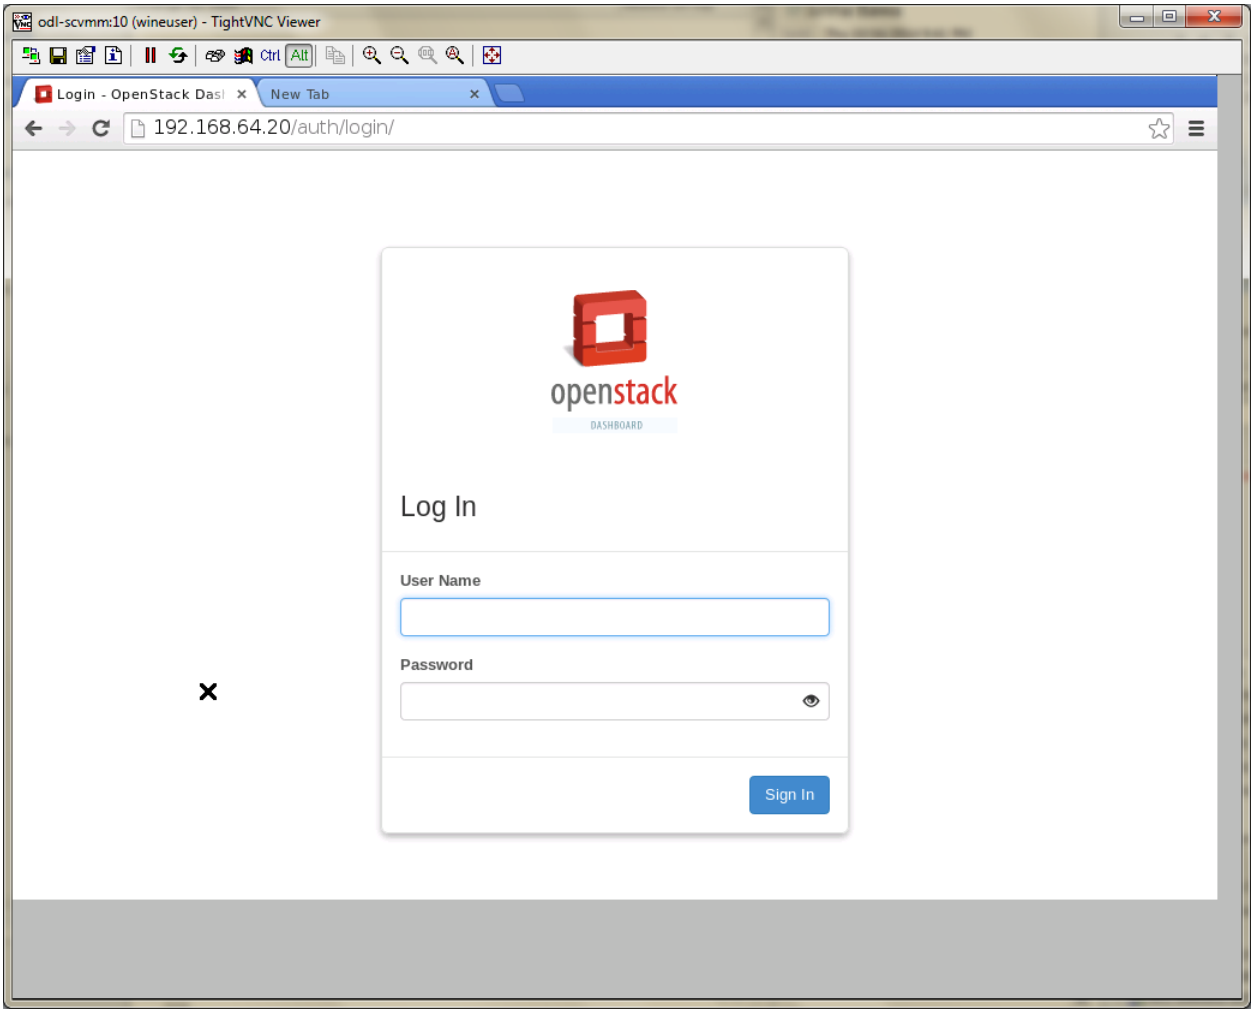

Fig. 114: Horizon GUI

Enter the value for User Name as admin and enter the value for Password as labstack.

- We should first ensure both the hypervisors(control node and compute node) are mapped under hypervisors by clicking on Hpervisors tab.
- Create a new Network from Horizon GUI.
- Click on Networks Tab.

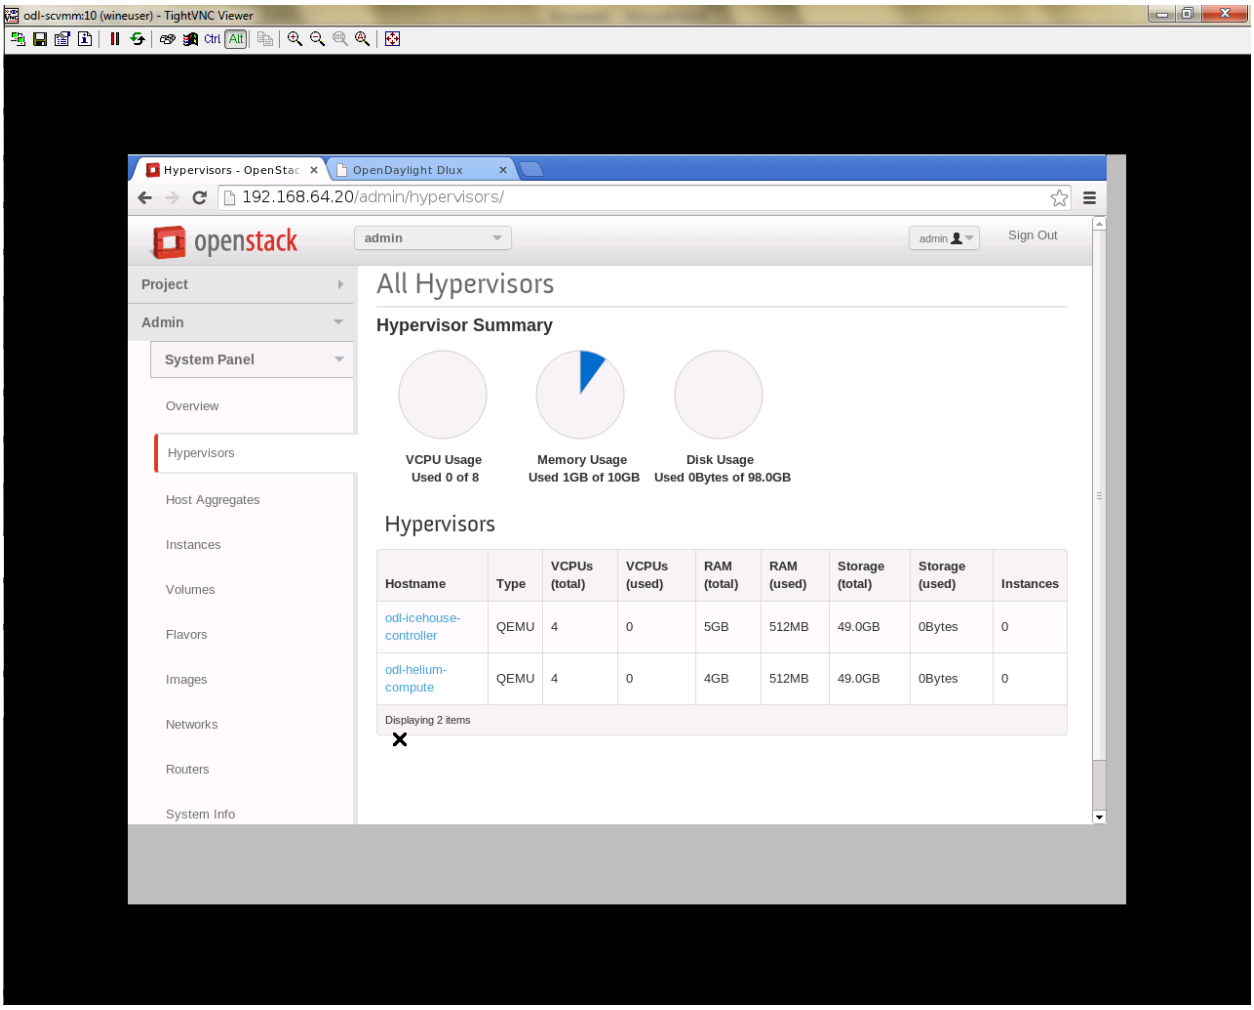

Fig. 115: Hypervisors

- click on the Create Network button.

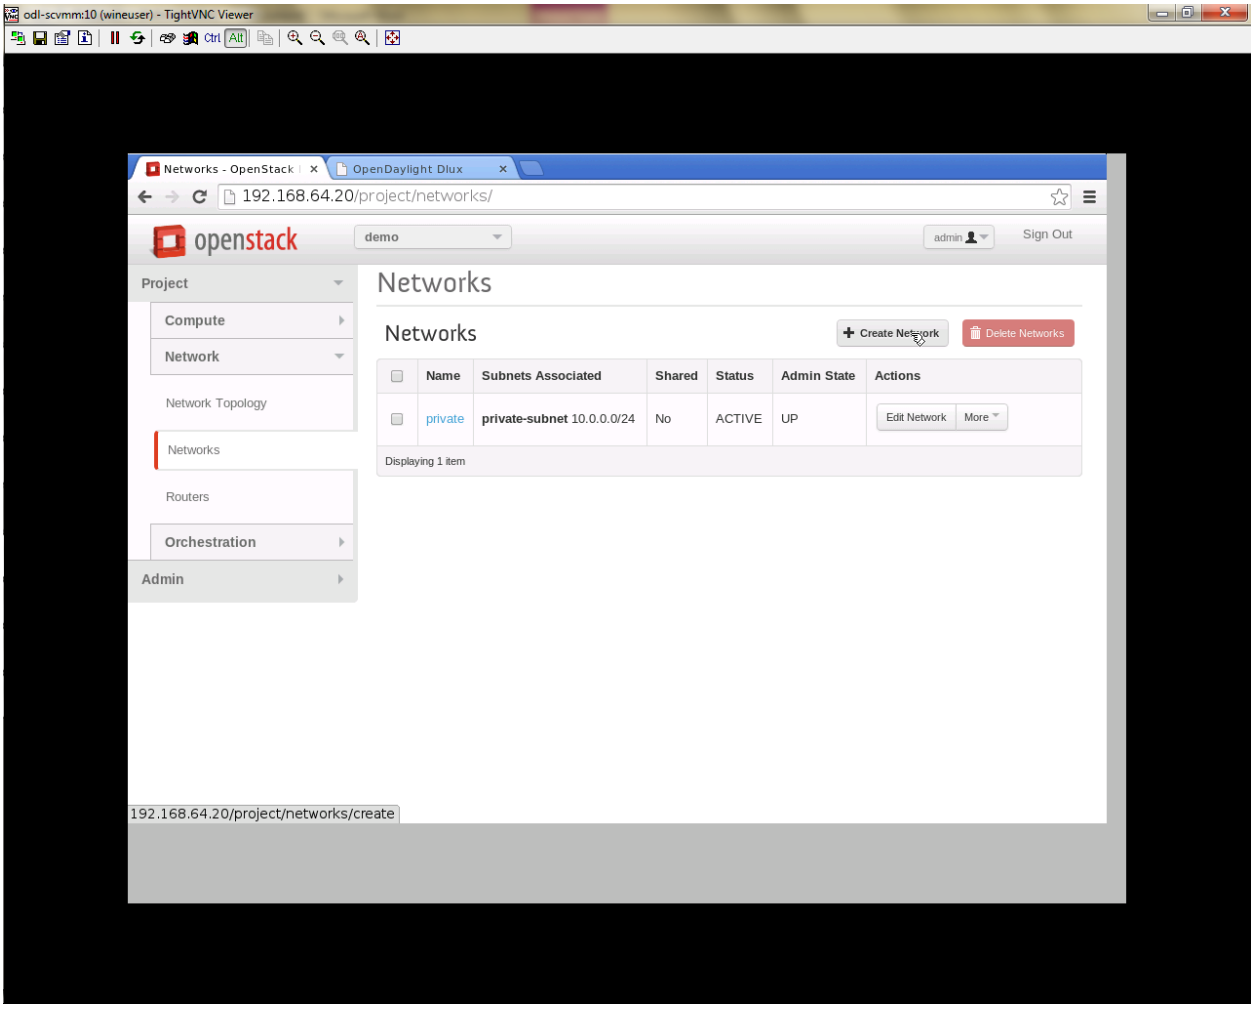

Fig. 116: Create Network

- A popup screen will appear.
- Enter network name and click Next button.
- Create a sub network by giving Network Address and click Next button .
- Specify the additional details for subnetwork (please refer the image for your reference).
- Click Create button
- Create VM Instance
- Navigate to Instances tab in the GUI.
- Click on Launch Instances button.
- Click on Details tab to enter the VM details.For this demo we are creating Ten VM's(instances).
- In the Networking tab, we must select the network,for this we need to drag and drop the Available networks to Selected Networks (i.e.,) Drag vtn1 we created from Available networks to Selected Networks and click Launch to create the instances.
- Ten VM's will be created.

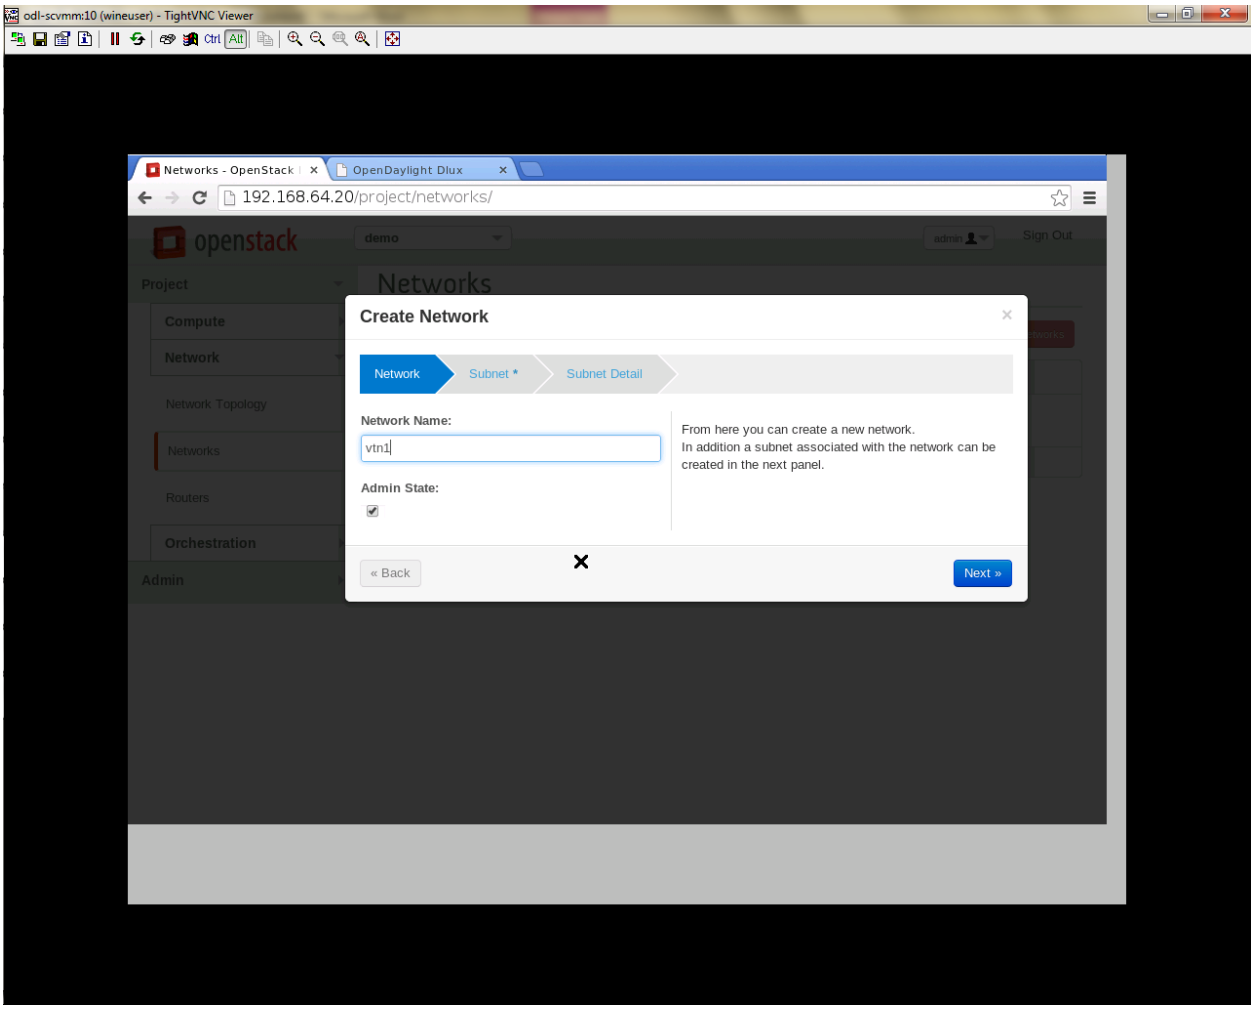

Fig. 117: Step 1

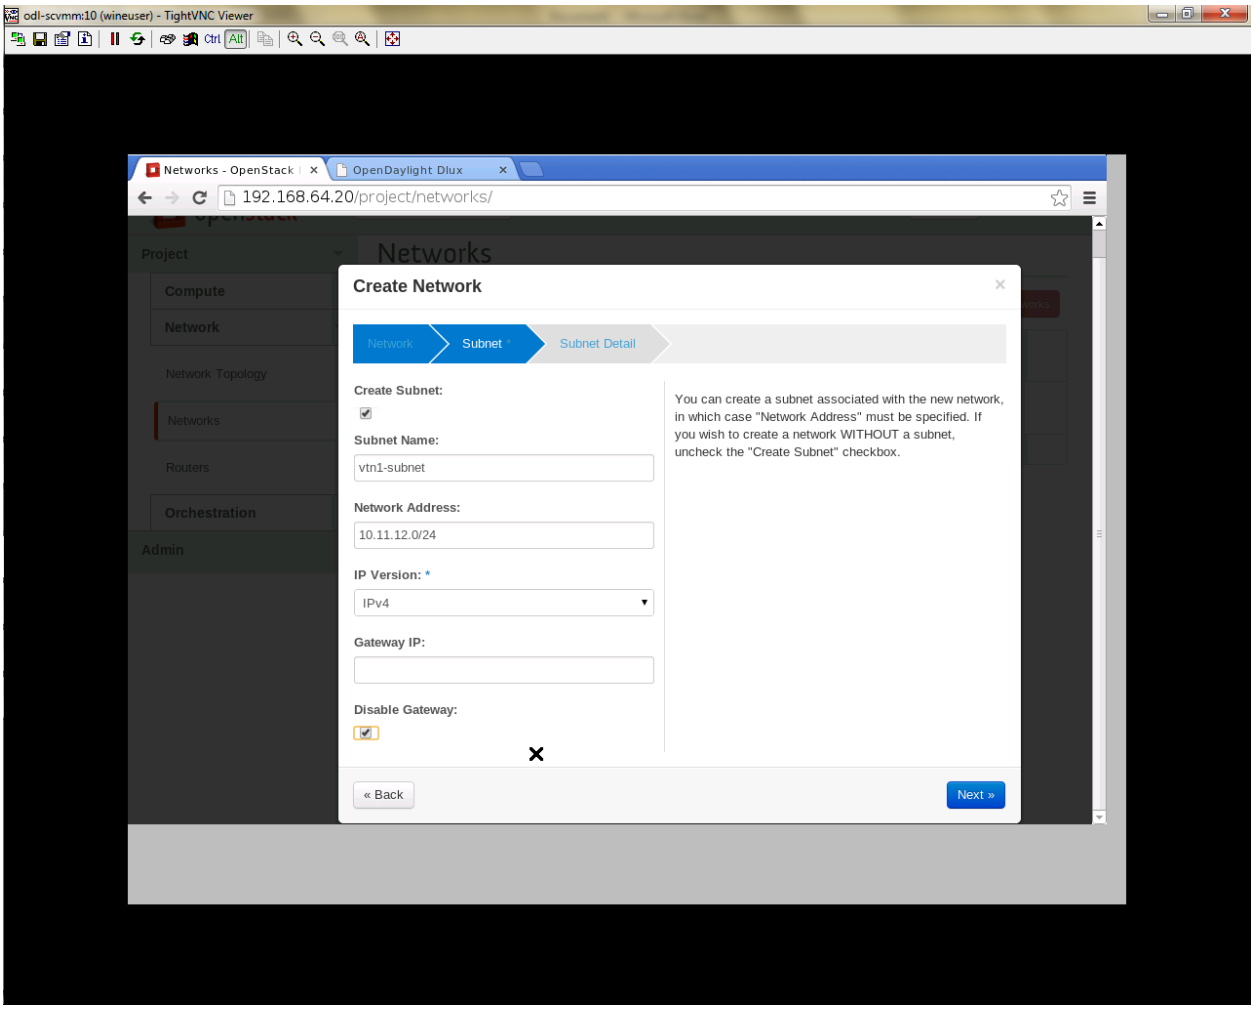

Fig. 118: Step 2

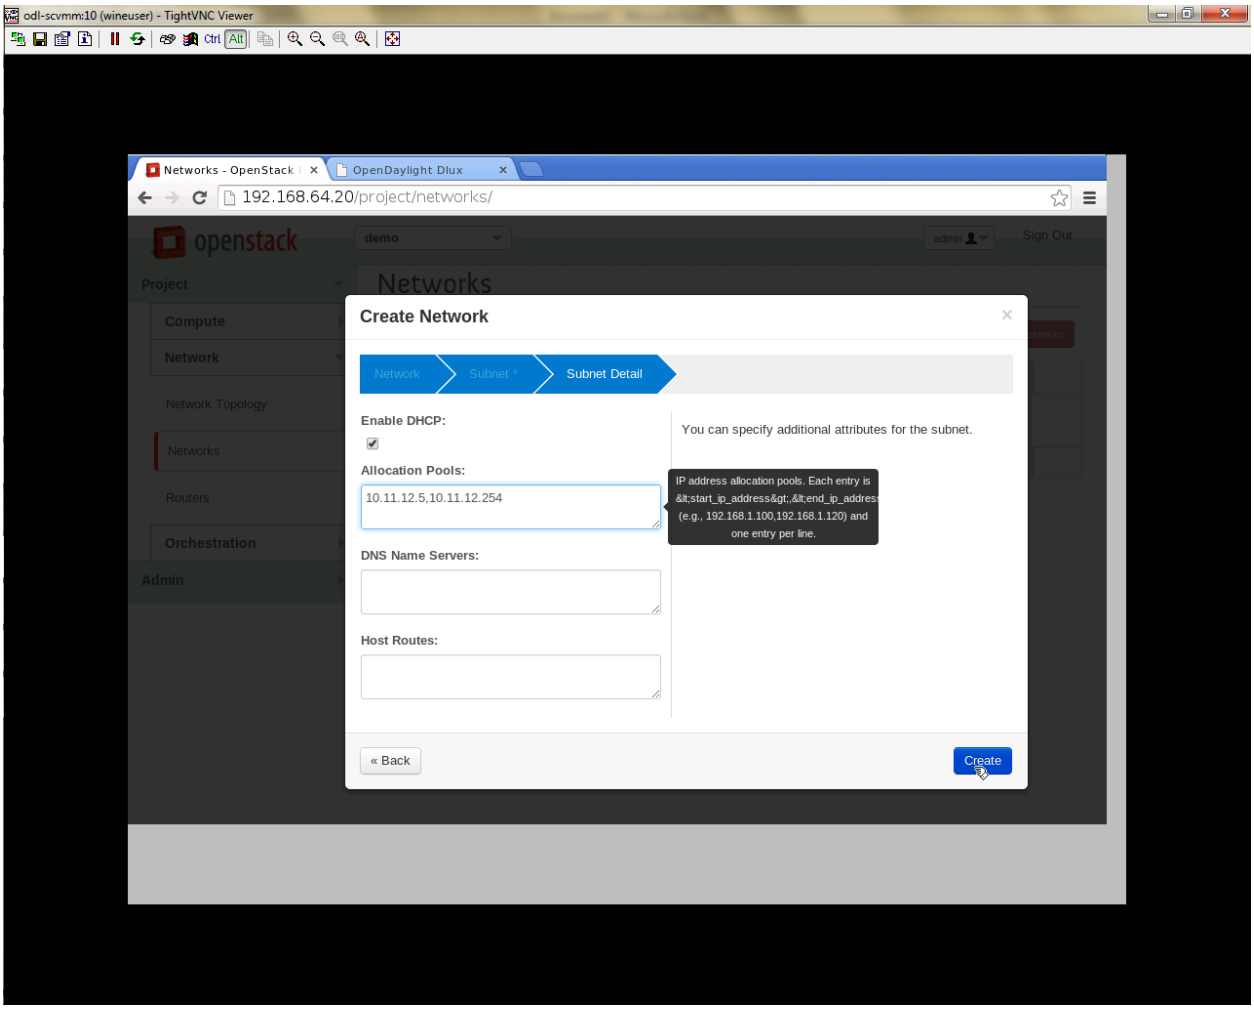

Fig. 119: Step 3

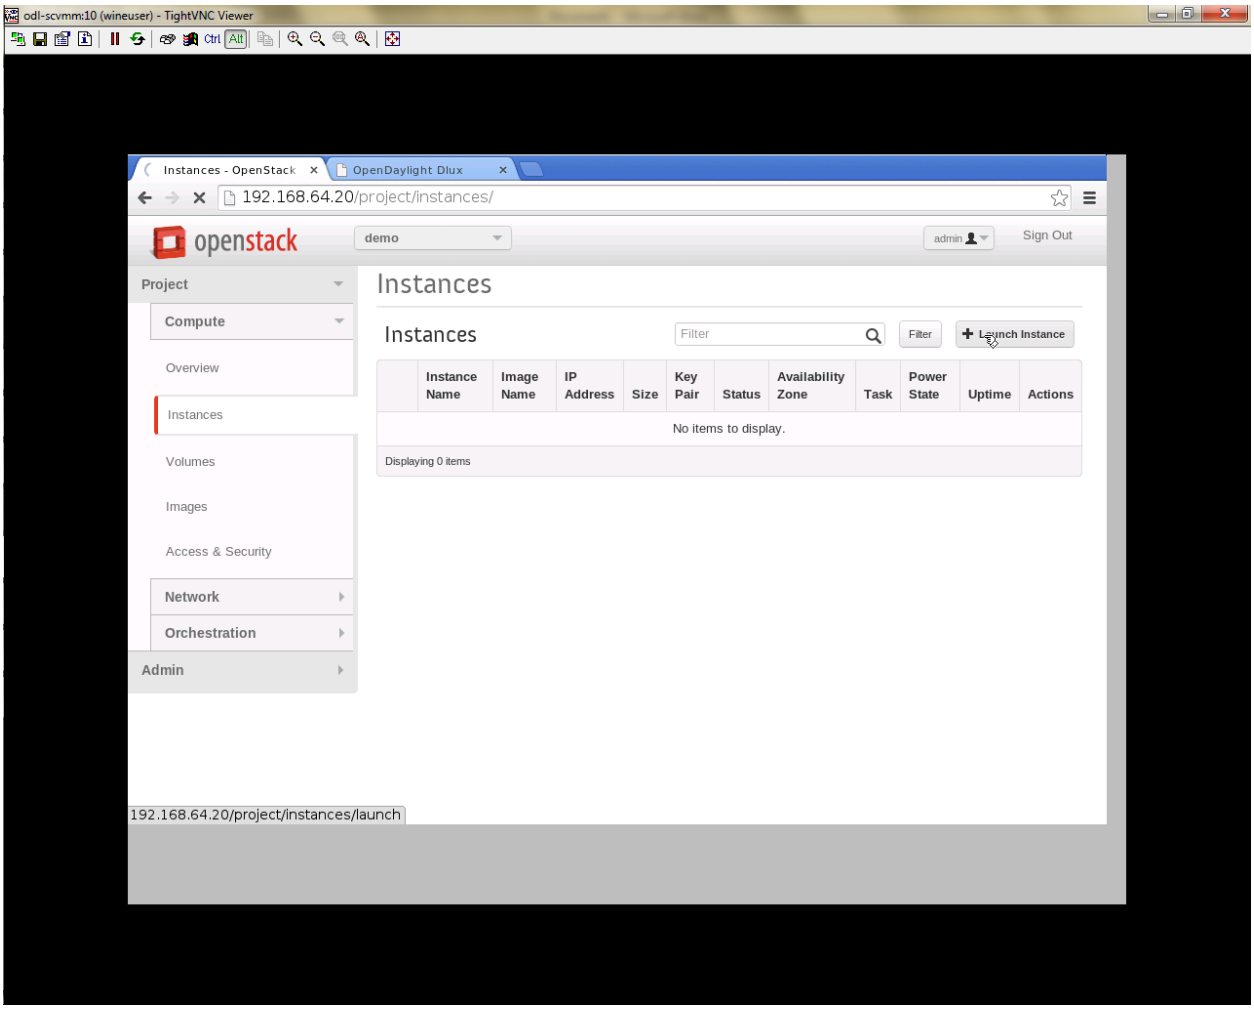

Fig. 120: Instance Creation

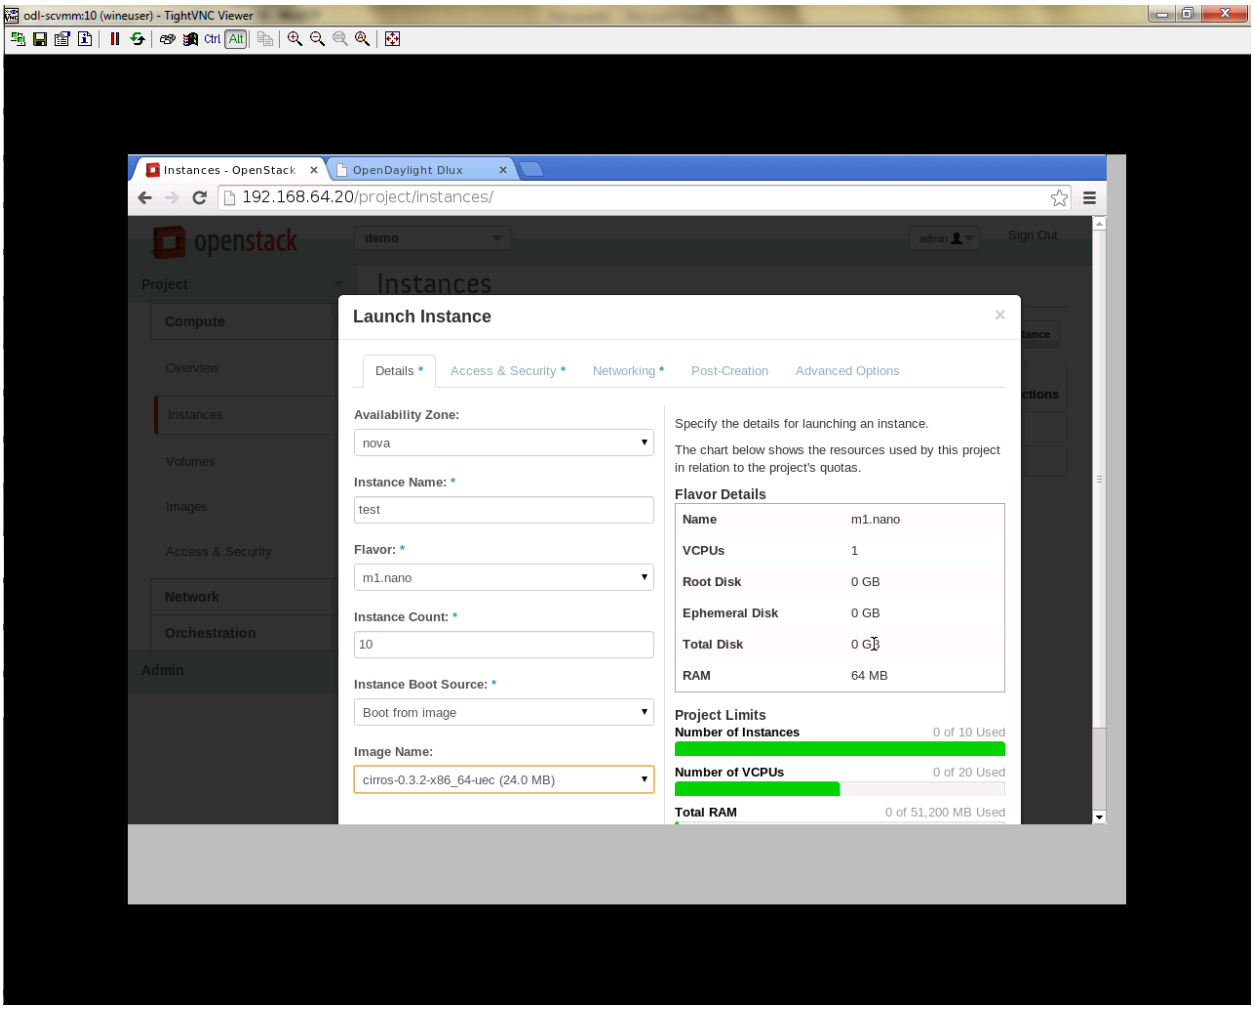

Fig. 121: Launch Instance

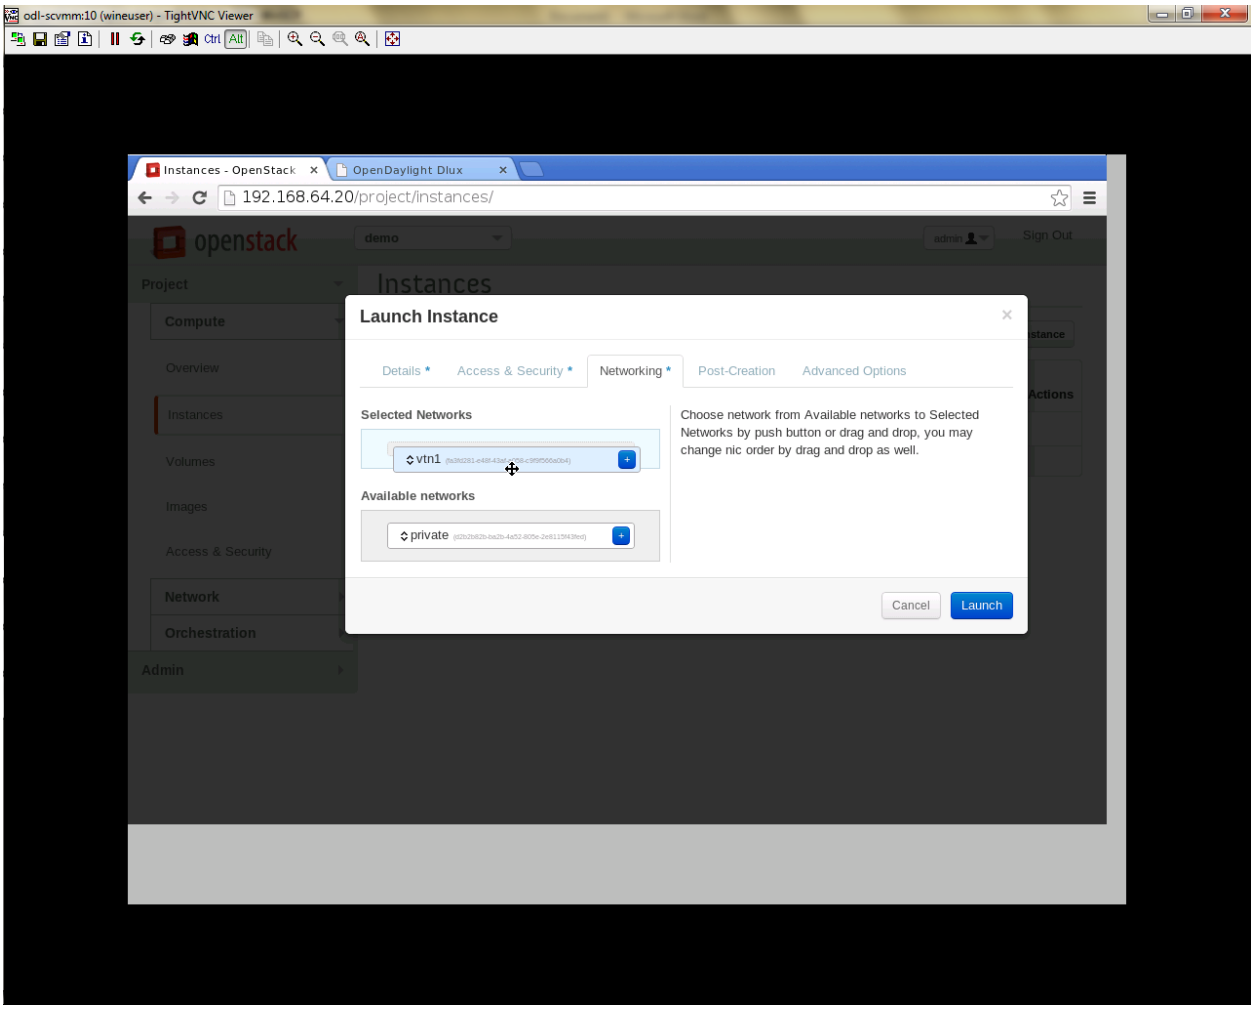

Fig. 122: Launch Network

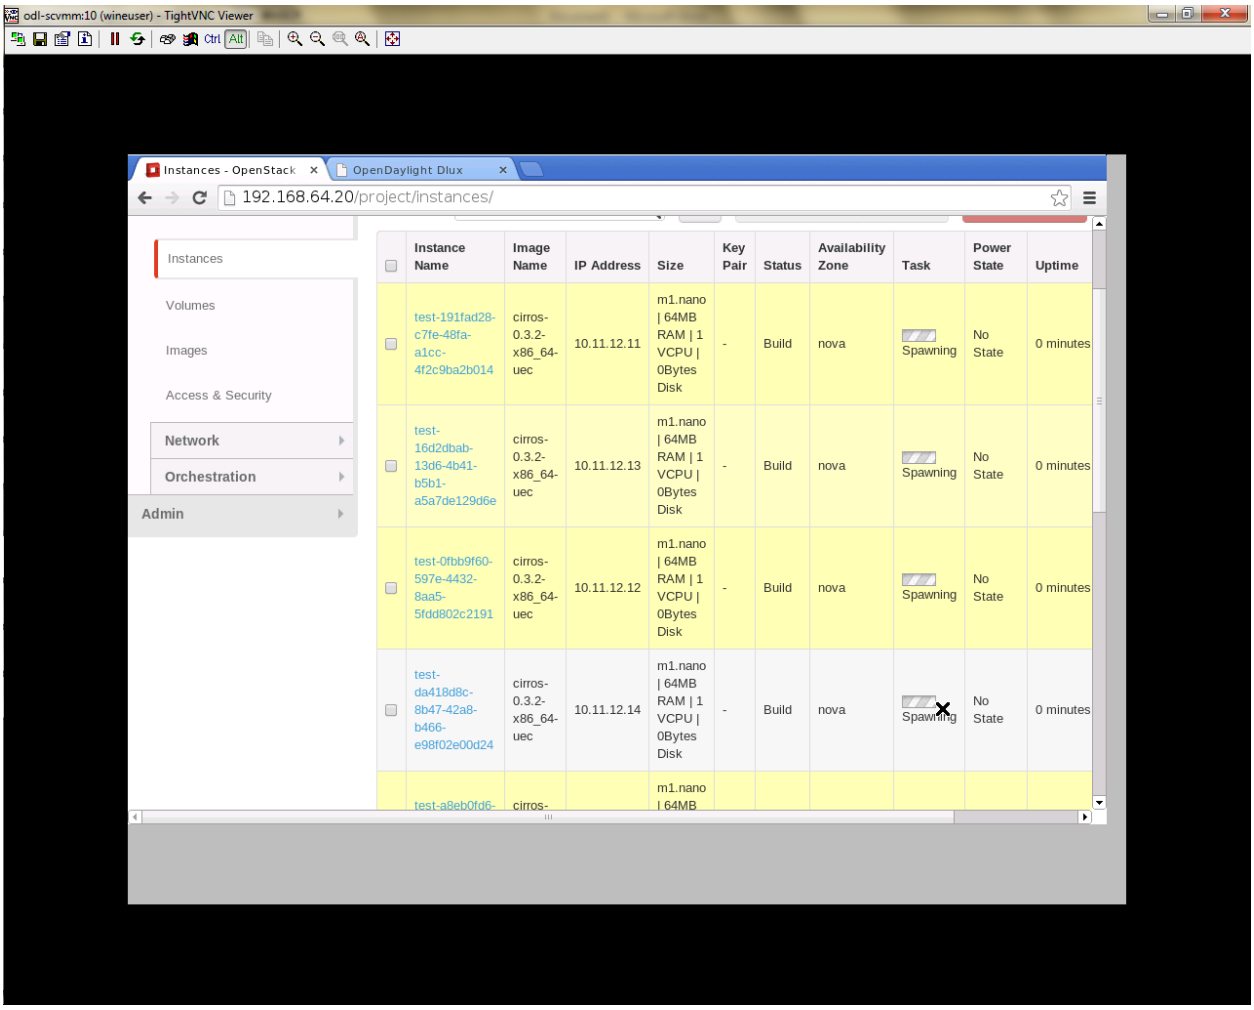

Fig. 123: Load All Instances

- Click on any VM displayed in the Instances tab and click the Console tab.

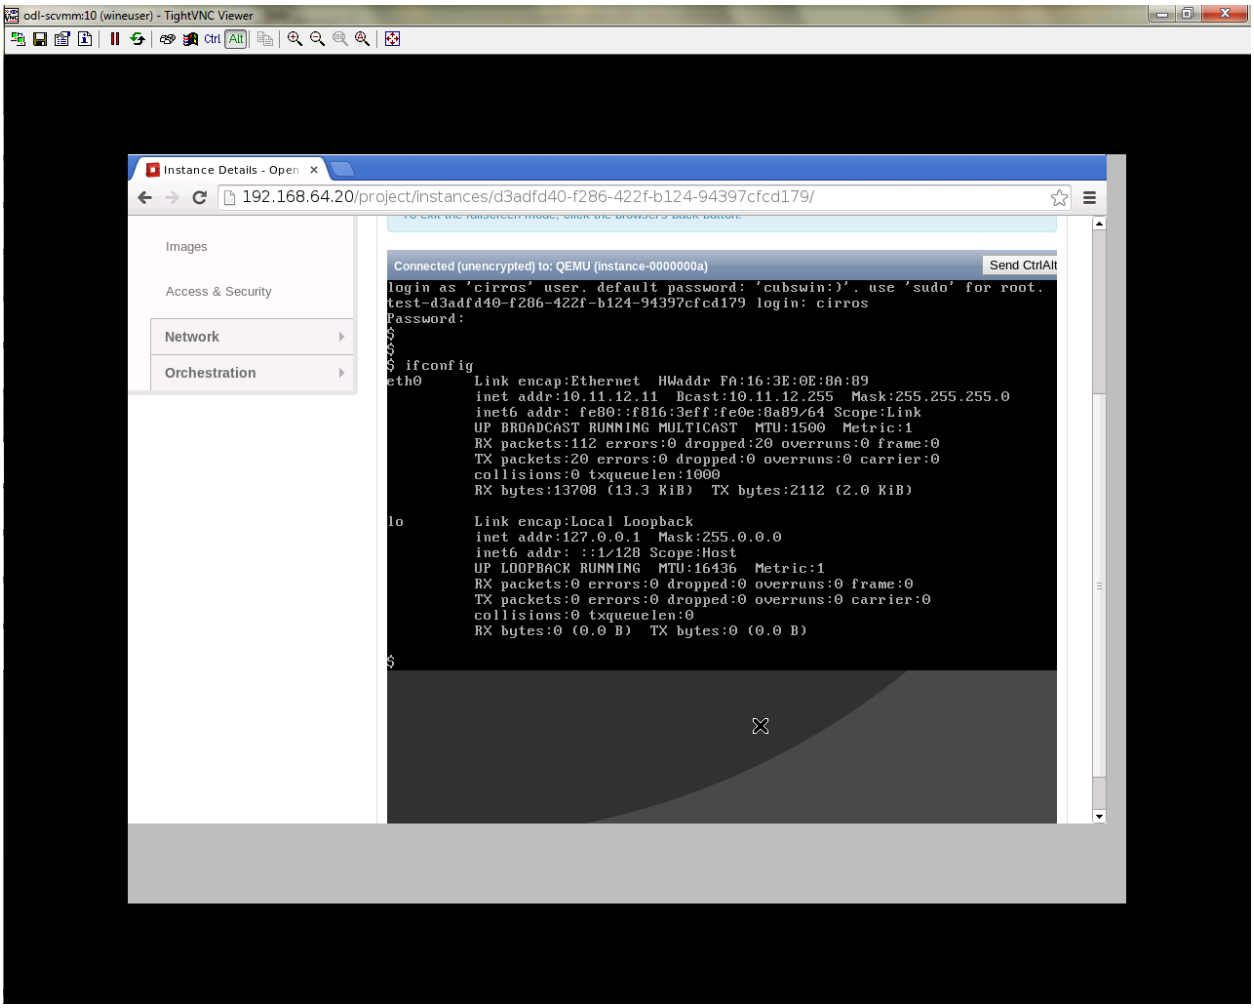

Fig. 124: Instance Console

- Login to the VM console and verify with a ping command.

Verification of Control and Compute Node after VM creation

- Every time a new VM is created, more interfaces are added to the br-int bridge in Open vSwitch.
- Use *sudo ovs-vsctl show* to list the number of interfaces added.
- Please visit the DLUX GUI to list the new nodes in every switch.

Getting started with DLUX

Ensure that you have created a topology and enabled MD-SAL feature in the Karaf distribution before you use DLUX for network management.

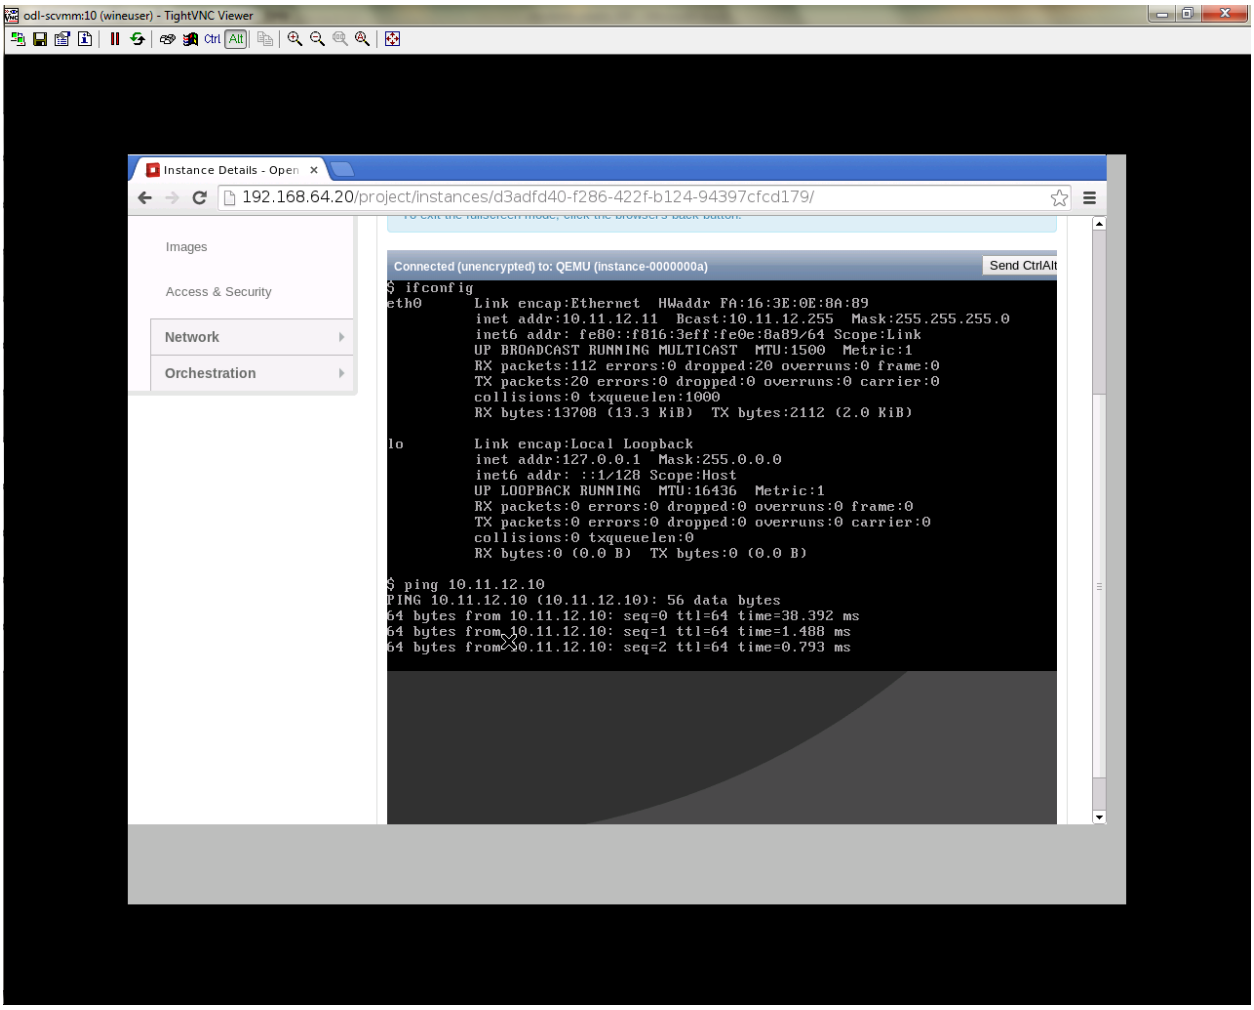
odl-scvmm:10 (wineuser) - TightVNC Viewer
Instance Details - Open
192.168.64.20/project/instances/d3adfd40-f286-422f-b124-94397cfcd179/
Images
Access & Security
Network
Orchestration
Connected (unencrypted) to: QEMU (instance-0000000a)
Send CtrlAlt
$ ifconfig
eth0 Link encap:Ethernet HWaddr FA:16:3E:0E:8A:89
inet addr:10.11.12.11 Bcast:10.11.12.255 Mask:255.255.255.0
inet6 addr: fe80::f816:3eff:fe0e:8a89/64 Scope:Link
UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
RX packets:112 errors:0 dropped:20 overruns:0 frame:0
TX packets:20 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:1000
RX bytes:13708 (13.3 KiB) TX bytes:2112 (2.0 KiB)
lo Link encap:Local Loopback
inet addr:127.0.0.1 Mask:255.0.0.0
inet6 addr: ::1/128 Scope:Host
UP LOOPBACK RUNNING MTU:16436 Metric:1
RX packets:0 errors:0 dropped:0 overruns:0 frame:0
TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:0
RX bytes:0 (0.0 B) TX bytes:0 (0.0 B)
$ ping 10.11.12.10
PING 10.11.12.10 (10.11.12.10): 56 data bytes
64 bytes from 10.11.12.10: seq=0 ttl=64 time=38.392 ms
64 bytes from 10.11.12.10: seq=1 ttl=64 time=1.488 ms
64 bytes from 10.11.12.10: seq=2 ttl=64 time=0.793 ms

Fig. 125: Ping

Logging In

To log in to DLUX, after installing the application: * Open a browser and enter the login URL. If you have installed DLUX as a stand-alone, then the login URL is http://localhost:9000/DLUX/index.html. However if you have deployed DLUX with Karaf, then the login URL is http://<your IP>:8181/dlux/index.html. * Login to the application with user ID and password credentials as admin. NOTE:admin is the only user type available for DLUX in this release.

Working with DLUX

To get a complete DLUX feature list, install restconf, odl l2 switch, and switch while you start the DLUX distribution.

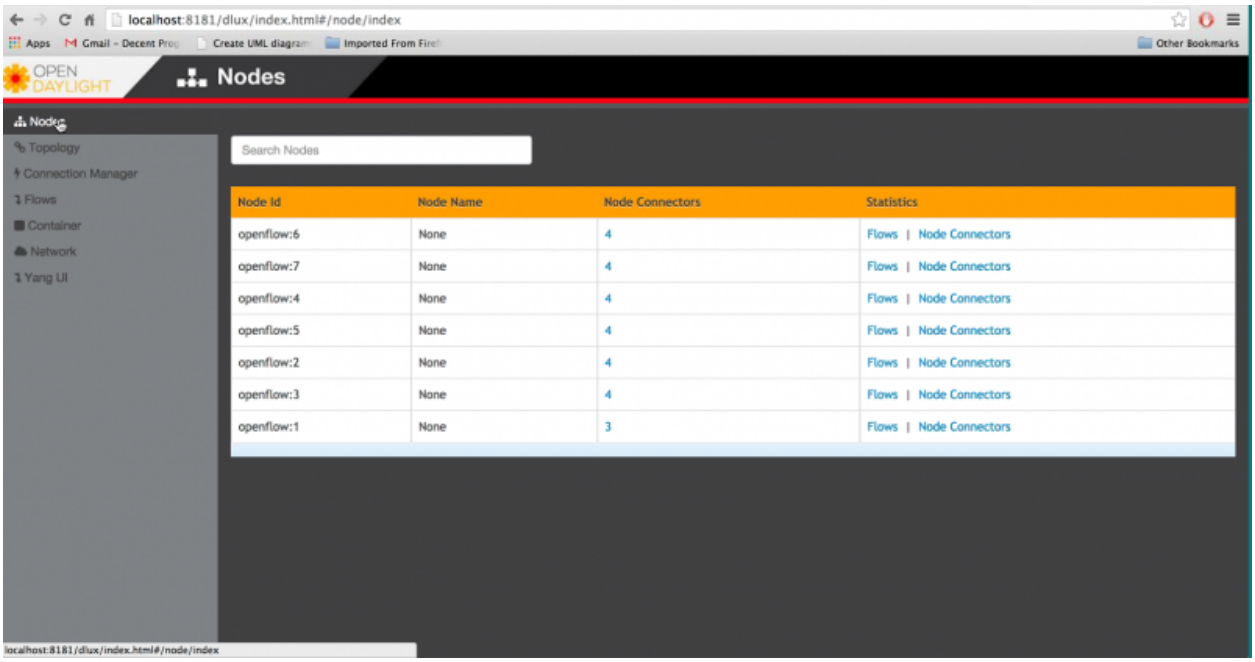

Fig. 126: DLUX_GUI

Note: DLUX enables only those modules, whose APIs are responding. If you enable just the MD-SAL in beginning and then start dlux, only MD-SAL related tabs will be visible. While using the GUI if you enable AD-SAL karaf features, those tabs will appear automatically.

Viewing Network Statistics

The Nodes module on the left pane enables you to view the network statistics and port information for the switches in the network. * To use the Nodes module: ** Select Nodeson the left pane.

```
The right pane displays atable that lists all the nodes, node connectors and the␣
↪statistics.
```

- Enter a node ID in the Search Nodes tab to search by node connectors.
- Click on the Node Connector number to view details such as port ID, port name, number of ports per switch, MAC Address, and so on.

- Click Flows in the Statistics column to view Flow Table Statistics for the particular node like table ID, packet match, active flows and so on.
- Click Node Connectors to view Node Connector Statistics for the particular node ID.

Viewing Network Topology

To view network topology: * Select Topology on the left pane. You will view the graphical representation on the right pane.

```
In the diagram
blue boxes represent the switches,black represents the hosts available, and lines␣
↪represents how switches are connected.
```

Note: DLUX UI does not provide ability to add topology information. The Topology should be created using an open flow plugin. Controller stores this information in the database and displays on the DLUX page, when the you connect to the controller using OpenFlow.

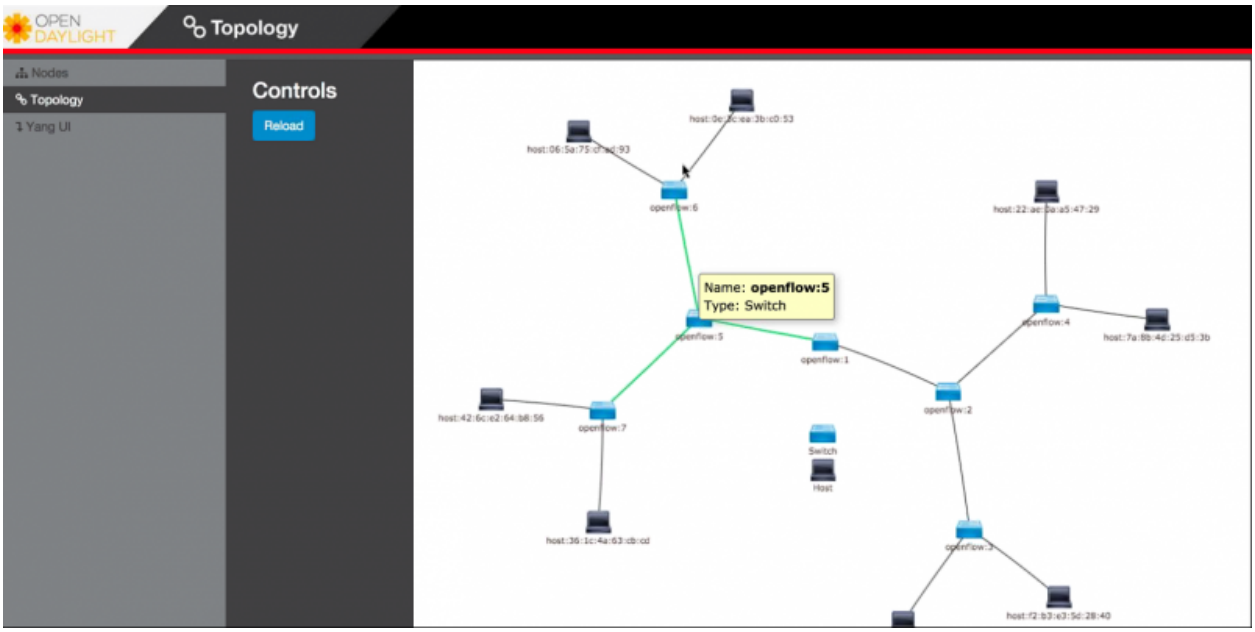

Fig. 127: Topology

OpenStack PackStack Installation Steps

- Please go through the below wiki page for OpenStack PackStack installation steps.
  - https://wiki.opendaylight.org/view/Release/Lithium/VTN/User_Guide/Openstack_Packstack_Support

References

- http://devstack.org/guides/multinode-lab.html

- https://wiki.opendaylight.org/view/File:Vtn_demo_hackfest_2014_march.pdf

VTN Manager Usage Examples

How to provision virtual L2 Network

Overview

This page explains how to provision virtual L2 network using VTN Manager. This page targets Boron release, so the procedure described here does not work in other releases.

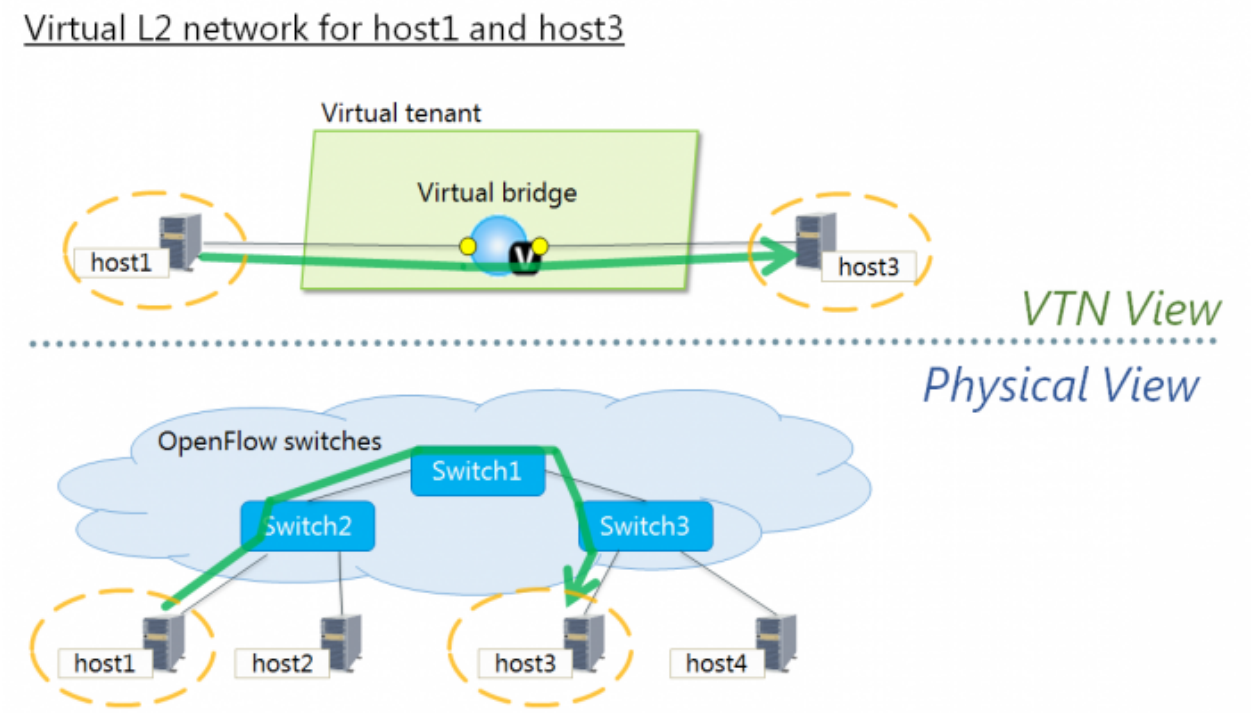

Fig. 128: Virtual L2 network for host1 and host3

Requirements

Mininet

- To provision OpenFlow switches, this page uses Mininet. Mininet details and set-up can be referred at the following page: https://wiki.opendaylight.org/view/OpenDaylight_Controller:Installation#Using_Mininet
- Start Mininet and create three switches(s1, s2, and s3) and four hosts(h1, h2, h3, and h4) in it.

```
mininet@mininet-vm:~$ sudo mn --controller=remote,ip=192.168.0.100 --topo tree,2
```

Note: Replace "192.168.0.100" with the IP address of OpenDaylight controller based on your environment.

- you can check the topology that you have created by executing "net" command in the Mininet console.

```
mininet> net
h1 h1-eth0:s2-eth1
h2 h2-eth0:s2-eth2
h3 h3-eth0:s3-eth1
h4 h4-eth0:s3-eth2
s1 lo:  s1-eth1:s2-eth3 s1-eth2:s3-eth3
s2 lo:  s2-eth1:h1-eth0 s2-eth2:h2-eth0 s2-eth3:s1-eth1
s3 lo:  s3-eth1:h3-eth0 s3-eth2:h4-eth0 s3-eth3:s1-eth2
```

- In this guide, you will provision the virtual L2 network to establish communication between h1 and h3.

Configuration

To provision the virtual L2 network for the two hosts (h1 and h3), execute REST API provided by VTN Manager as follows. It uses curl command to call the REST API.

- Create a virtual tenant named vtn1 by executing the update-vtn RPC.

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn:update-vtn -d '{"input":{"tenant-name":"vtn1
↪"}}'
```

- Create a virtual bridge named vbr1 in the tenant vtn1 by executing the update-vbridge RPC.

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn-vbridge:update-vbridge -d '{"input":{"tenant-
↪name":"vtn1", "bridge-name":"vbr1"}}'
```

- Create two interfaces into the virtual bridge by executing the update-vinterface RPC.

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn-vinterface:update-vinterface -d '{"input":{
↪"tenant-name":"vtn1", "bridge-name":"vbr1", "interface-name":"if1"}}'
```

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn-vinterface:update-vinterface -d '{"input":{
↪"tenant-name":"vtn1", "bridge-name":"vbr1", "interface-name":"if2"}}'
```

- Configure two mappings on the created interfaces by executing the set-port-map RPC.
  - The interface if1 of the virtual bridge will be mapped to the port "s2-eth1" of the switch "openflow:2" of the Mininet.
    - The h1 is connected to the port "s2-eth1".
  - The interface if2 of the virtual bridge will be mapped to the port "s3-eth1" of the switch "openflow:3" of the Mininet.
    - The h3 is connected to the port "s3-eth1".

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn-port-map:set-port-map -d '{"input":{"tenant-
↪name":"vtn1", "bridge-name":"vbr1", "interface-name":"if1", "node":"openflow:2",
↪"port-name":"s2-eth1"}}'
```

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn-port-map:set-port-map -d '{"input":{"tenant-
↪name":"vtn1", "bridge-name":"vbr1", "interface-name":"if2", "node":"openflow:3",
↪"port-name":"s3-eth1"}}'
```

Verification

- Please execute ping from h1 to h3 to verify if the virtual L2 network for h1 and h3 is provisioned successfully.

```
mininet> h1 ping h3
PING 10.0.0.3 (10.0.0.3) 56(84) bytes of data.
64 bytes from 10.0.0.3: icmp_seq=1 ttl=64 time=243 ms
64 bytes from 10.0.0.3: icmp_seq=2 ttl=64 time=0.341 ms
64 bytes from 10.0.0.3: icmp_seq=3 ttl=64 time=0.078 ms
64 bytes from 10.0.0.3: icmp_seq=4 ttl=64 time=0.079 ms
```

- You can also verify the configuration by executing the following REST API. It shows all configuration in VTN Manager.

```
curl --user "admin":"admin" -H "Content-type: application/json" -X GET http://
↪localhost:8181/restconf/operational/vtn:vtns/
```

- The result of the command should be like this.

```
{
  "vtns": {
    "vtn": [
    {
      "name": "vtn1",
        "vtenant-config": {
          "idle-timeout": 300,
          "hard-timeout": 0
        },
        "vbridge": [
        {
          "name": "vbr1",
          "bridge-status": {
            "state": "UP",
            "path-faults": 0
          },
          "vbridge-config": {
            "age-interval": 600
          },
          "vinterface": [
          {
            "name": "if2",
            "vinterface-status": {
              "entity-state": "UP",
              "state": "UP",
              "mapped-port": "openflow:3:3"
            },
            "vinterface-config": {
              "enabled": true
            },
            "port-map-config": {
```

(continues on next page)

(continued from previous page)

```
                "vlan-id": 0,
                "port-name": "s3-eth1",
                "node": "openflow:3"
              }
            },
            {
              "name": "if1",
              "vinterface-status": {
                "entity-state": "UP",
                "state": "UP",
                "mapped-port": "openflow:2:1"
              },
              "vinterface-config": {
                "enabled": true
              },
              "port-map-config": {
                "vlan-id": 0,
                "port-name": "s2-eth1",
                "node": "openflow:2"
              }
            }
            ]
          }
        ]
      }
      ]
    }
}
```

Cleaning Up

- You can delete the virtual tenant vtn1 by executing the remove-vtn RPC.

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn:remove-vtn -d '{"input":{"tenant-name":"vtn1
↪"}}'
```

How To Test Vlan-Map In Mininet Environment

Overview

This page explains how to test Vlan-map in a multi host scenario using mininet. This page targets Boron release, so the procedure described here does not work in other releases.

Requirements

Save the mininet script given below as vlan_vtn_test.py and run the mininet script in the mininet environment where Mininet is installed.

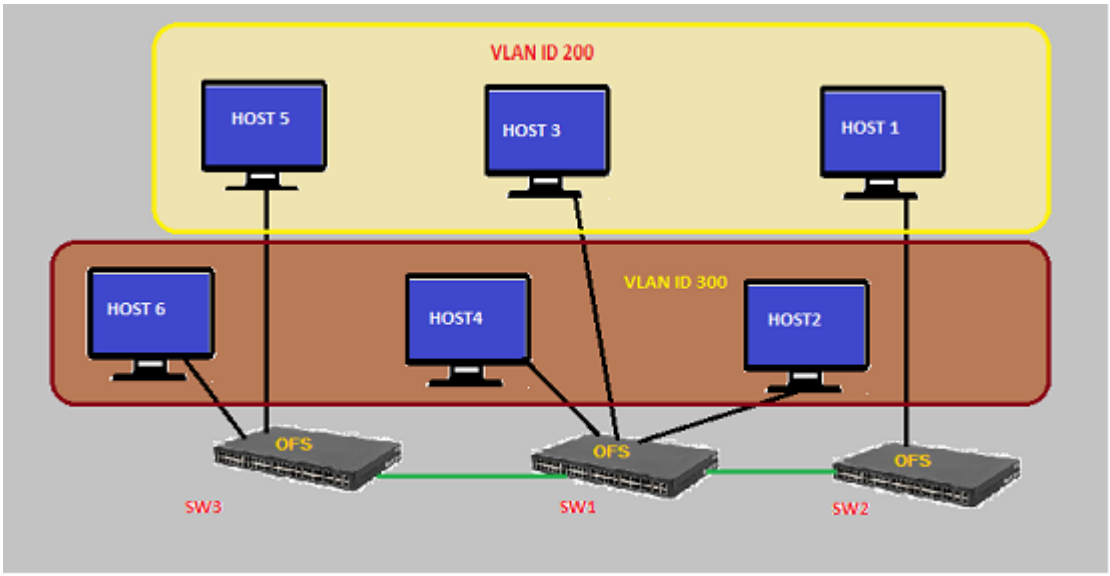

Fig. 129: Example that demonstrates vlanmap testing in Mininet Environment

Mininet Script

https://wiki.opendaylight.org/view/OpenDaylight_Virtual_Tenant_Network_(VTN):Scripts:Mininet#Network_with_hosts_in_different_vlan

- Run the mininet script

```
sudo mn --controller=remote,ip=192.168.64.13 --custom vlan_vtn_test.py --topo mytopo
```

Note: Replace "192.168.64.13" with the IP address of OpenDaylight controller based on your environment.

- You can check the topology that you have created by executing "net" command in the Mininet console.

```
mininet> net
h1 h1-eth0.200:s1-eth1
h2 h2-eth0.300:s2-eth2
h3 h3-eth0.200:s2-eth3
h4 h4-eth0.300:s2-eth4
h5 h5-eth0.200:s3-eth2
h6 h6-eth0.300:s3-eth3
s1 lo:  s1-eth1:h1-eth0.200 s1-eth2:s2-eth1 s1-eth3:s3-eth1
s2 lo:  s2-eth1:s1-eth2 s2-eth2:h2-eth0.300 s2-eth3:h3-eth0.200 s2-eth4:h4-eth0.300
s3 lo:  s3-eth1:s1-eth3 s3-eth2:h5-eth0.200 s3-eth3:h6-eth0.300
c0
```

Configuration

To test vlan-map, execute REST API provided by VTN Manager as follows.

- Create a virtual tenant named vtn1 by executing the update-vtn RPC.

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn:update-vtn -d '{"input":{"tenant-name":"vtn1
↪"}}'
```

- Create a virtual bridge named vbr1 in the tenant vtn1 by executing the update-vbridge RPC.

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn-vbridge:update-vbridge -d '{"input":{"tenant-
↪name":"vtn1","bridge-name":"vbr1"}}'
```

- Configure a vlan map with vlanid 200 for vBridge vbr1 by executing the add-vlan-map RPC.

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn-vlan-map:add-vlan-map -d '{"input":{"vlan-id
↪":200,"tenant-name":"vtn1","bridge-name":"vbr1"}}'
```

- Create a virtual bridge named vbr2 in the tenant vtn1 by executing the update-vbridge RPC.

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn-vbridge:update-vbridge -d '{"input":{"tenant-
↪name":"vtn1","bridge-name":"vbr2"}}'
```

- Configure a vlan map with vlanid 300 for vBridge vbr2 by executing the add-vlan-map RPC.

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn-vlan-map:add-vlan-map -d '{"input":{"vlan-id
↪":300,"tenant-name":"vtn1","bridge-name":"vbr2"}}'
```

Verification

- Please execute pingall in mininet environment to view the host reachability.

```
mininet> pingall
Ping: testing ping reachability
h1 -> X h3 X h5 X
h2 -> X X h4 X h6
h3 -> h1 X X h5 X
h4 -> X h2 X X h6
h5 -> h1 X h3 X X
h6 -> X h2 X h4 X
```

- You can also verify the configuration by executing the following REST API. It shows all configurations in VTN Manager.

```
curl --user "admin":"admin" -H "Content-type: application/json" -X GET http://
↪localhost:8181/restconf/operational/vtn:vtns
```

- The result of the command should be like this.

```
{
  "vtns": {
    "vtn": [
    {
      "name": "vtn1",
        "vtenant-config": {
```

(continues on next page)

(continued from previous page)

```
          "hard-timeout": 0,
          "idle-timeout": 300,
          "description": "creating vtn"
        },
        "vbridge": [
        {
          "name": "vbr2",
          "vbridge-config": {
            "age-interval": 600,
            "description": "creating vbr2"
          },
          "bridge-status": {
            "state": "UP",
            "path-faults": 0
          },
          "vlan-map": [
          {
            "map-id": "ANY.300",
            "vlan-map-config": {
              "vlan-id": 300
            },
            "vlan-map-status": {
              "active": true
            }
          }
          ]
        },
        {
          "name": "vbr1",
          "vbridge-config": {
            "age-interval": 600,
            "description": "creating vbr1"
          },
          "bridge-status": {
            "state": "UP",
            "path-faults": 0
          },
          "vlan-map": [
          {
            "map-id": "ANY.200",
            "vlan-map-config": {
              "vlan-id": 200
            },
            "vlan-map-status": {
              "active": true
            }
          }
          ]
        }
      ]
    }
    ]
  }
}
```

Cleaning Up

- You can delete the virtual tenant vtn1 by executing the remove-vtn RPC.

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn:remove-vtn -d '{"input":{"tenant-name":"vtn1
↪"}}'
```

How To Configure Service Function Chaining using VTN Manager

Overview

This page explains how to configure VTN Manager for Service Chaining. This page targets Boron release, so the procedure described here does not work in other releases.

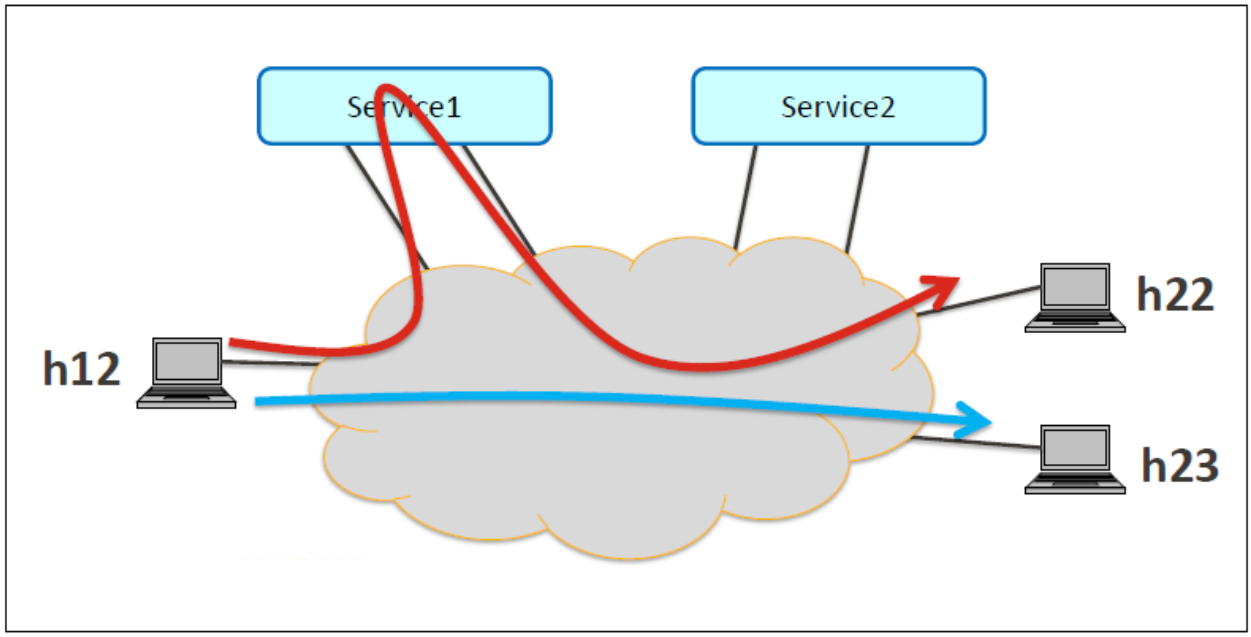

Fig. 130: Service Chaining With One Service

Requirements

- Please refer to the Installation Pages to run ODL with VTN Feature enabled.
- Please ensure Bridge-Utils package is installed in mininet environment before running the mininet script.
- To install Bridge-Utils package run sudo apt-get install bridge-utils (assuming Ubuntu is used to run mininet, If not then this is not required).
- Save the mininet script given below as topo_handson.py and run the mininet script in the mininet environment where Mininet is installed.

Mininet Script

- Script for emulating network with multiple hosts.
- Before executing the mininet script, please confirm Controller is up and running.
- Run the mininet script.
- Replace <path> and <Controller IP> based on your environment

```
sudo mn --controller=remote,ip=<Controller IP> --custom <path>\topo_handson.py --topo␣
↪mytopo2
```

```
mininet> net
h11 h11-eth0:s1-eth1
h12 h12-eth0:s1-eth2
h21 h21-eth0:s2-eth1
h22 h22-eth0:s2-eth2
h23 h23-eth0:s2-eth3
srvc1 srvc1-eth0:s3-eth3 srvc1-eth1:s4-eth3
srvc2 srvc2-eth0:s3-eth4 srvc2-eth1:s4-eth4
s1 lo:  s1-eth1:h11-eth0 s1-eth2:h12-eth0 s1-eth3:s2-eth4 s1-eth4:s3-eth2
s2 lo:  s2-eth1:h21-eth0 s2-eth2:h22-eth0 s2-eth3:h23-eth0 s2-eth4:s1-eth3 s2-eth5:s4-
↪eth1
s3 lo:  s3-eth1:s4-eth2 s3-eth2:s1-eth4 s3-eth3:srvc1-eth0 s3-eth4:srvc2-eth0
s4 lo:  s4-eth1:s2-eth5 s4-eth2:s3-eth1 s4-eth3:srvc1-eth1 s4-eth4:srvc2-eth1
```

Configurations

Mininet

- Please follow the below steps to configure the network in mininet as in the below image:

Configure service nodes

- Please execute the following commands in the mininet console where mininet script is executed.

```
mininet> srvc1 ip addr del 10.0.0.6/8 dev srvc1-eth0
mininet> srvc1 brctl addbr br0
mininet> srvc1 brctl addif br0 srvc1-eth0
mininet> srvc1 brctl addif br0 srvc1-eth1
mininet> srvc1 ifconfig br0 up
mininet> srvc1 tc qdisc add dev srvc1-eth1 root netem delay 200ms
mininet> srvc2 ip addr del 10.0.0.7/8 dev srvc2-eth0
mininet> srvc2 brctl addbr br0
mininet> srvc2 brctl addif br0 srvc2-eth0
mininet> srvc2 brctl addif br0 srvc2-eth1
mininet> srvc2 ifconfig br0 up
mininet> srvc2 tc qdisc add dev srvc2-eth1 root netem delay 300ms
```

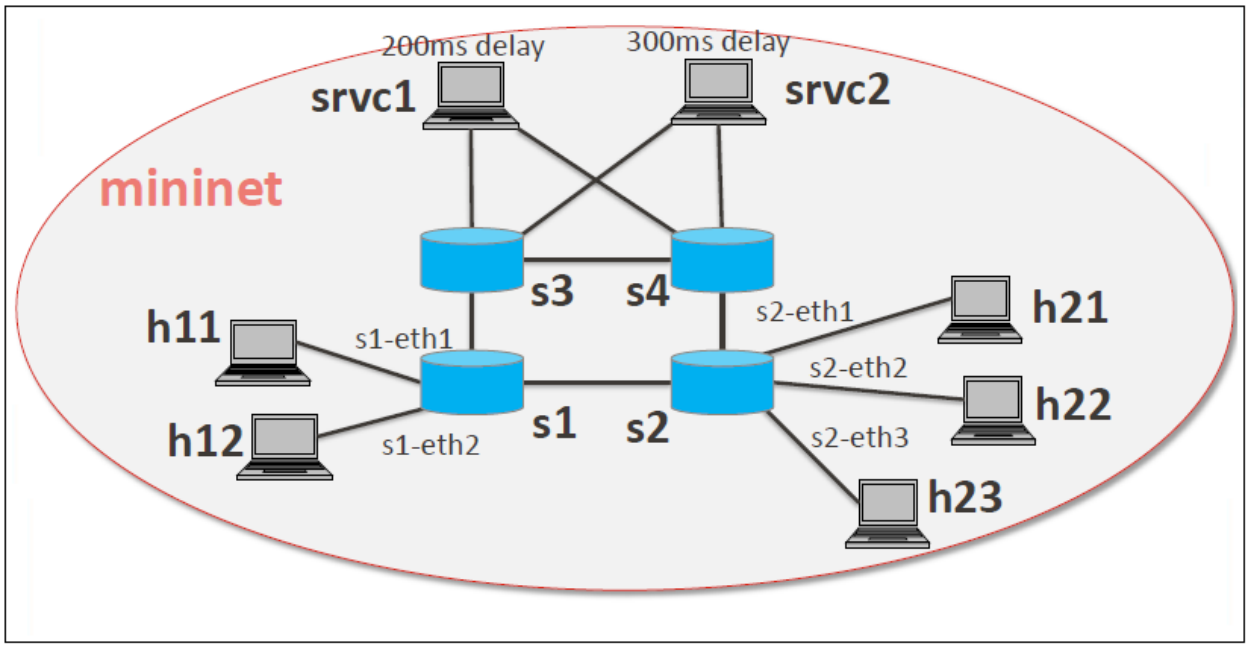

Fig. 131: Mininet Configuration

Controller

Multi-Tenancy

- Please execute the below commands to configure the network topology in the controller as in the below image:

Please execute the below commands in controller

Note: The below commands are for the difference in behavior of Manager in Boron topology. The Link below has the details for this bug: https://bugs.opendaylight.org/show_bug.cgi?id=3818.

```
curl --user admin:admin -H 'content-type: application/json' -H 'ipaddr:127.0.0.1' -X␣
↪PUT http://localhost:8181/restconf/config/vtn-static-topology:vtn-static-topology/
↪static-edge-ports -d '{"static-edge-ports": {"static-edge-port": [ {"port":
↪"openflow:3:3"}, {"port": "openflow:3:4"}, {"port": "openflow:4:3"}, {"port":
↪"openflow:4:4"}]}}'
```

- Create a virtual tenant named vtn1 by executing the update-vtn RPC.

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn:update-vtn -d '{"input":{"tenant-name":"vtn1
↪","update-mode":"CREATE","operation":"SET","description":"creating vtn","idle-
↪timeout":300,"hard-timeout":0}}'
```

- Create a virtual bridge named vbr1 in the tenant vtn1 by executing the update-vbridge RPC.

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn-vbridge:update-vbridge -d '{"input":{"update-
↪mode":"CREATE","operation":"SET","description":"creating vbr","tenant-name":"vtn1",
↪"bridge-name":"vbr1"}}'
```

(continues on next page)

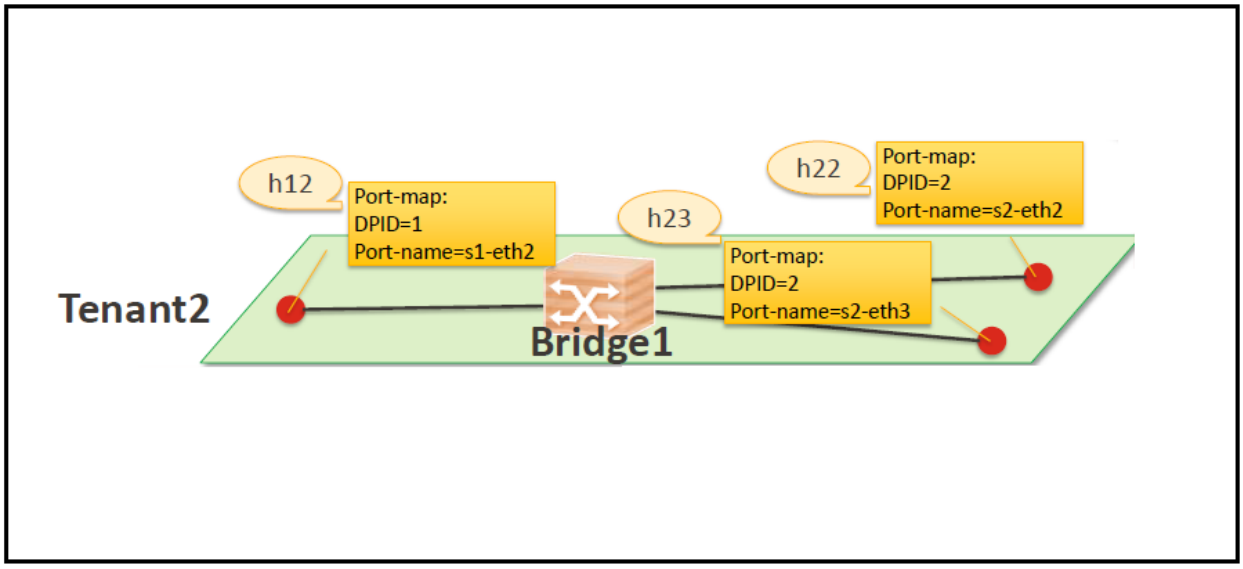

Fig. 132: Tenant2

(continued from previous page)

- Create interface if1 into the virtual bridge vbr1 by executing the update-vinterface RPC.

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn-vinterface:update-vinterface -d '{"input":{
↪"update-mode":"CREATE","operation":"SET","description":"Creating vbrif1 interface",
↪"tenant-name":"vtn1","bridge-name":"vbr1","interface-name":"if1"}}'
```

- Configure port mapping on the interface by executing the set-port-map RPC.
  - The interface if1 of the virtual bridge will be mapped to the port "s1-eth2" of the switch "openflow:1" of the Mininet.
    - The h12 is connected to the port "s1-eth2".

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn-port-map:set-port-map -d '{"input":{"vlan-id
↪":0,"tenant-name":"vtn1","bridge-name":"vbr1","interface-name":"if1","node":
↪"openflow:1","port-name":"s1-eth2"}}'
```

- Create interface if2 into the virtual bridge vbr1 by executing the update-vinterface RPC.

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn-vinterface:update-vinterface -d '{"input":{
↪"update-mode":"CREATE","operation":"SET","description":"Creating vbrif2 interface",
↪"tenant-name":"vtn1","bridge-name":"vbr1","interface-name":"if2"}}'
```

- Configure port mapping on the interface by executing the set-port-map RPC.
  - The interface if2 of the virtual bridge will be mapped to the port "s2-eth2" of the switch "openflow:2" of the Mininet.
    - The h22 is connected to the port "s2-eth2".

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn-port-map:set-port-map -d '{"input":{"vlan-id
↪":0,"tenant-name":"vtn1","bridge-name":"vbr1","interface-name":"if2","node":
↪"openflow:2","port-name":"s2-eth2"}}'
```

- Create interface if3 into the virtual bridge vbr1 by executing the update-vinterface RPC.

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn-vinterface:update-vinterface -d '{"input":{
↪"update-mode":"CREATE","operation":"SET","description":"Creating vbrif3 interface",
↪"tenant-name":"vtn1","bridge-name":"vbr1","interface-name":"if3"}}'
```

- Configure port mapping on the interfaces by executing the set-port-map RPC.
  - The interface if3 of the virtual bridge will be mapped to the port "s2-eth3" of the switch "openflow:2" of the Mininet.
    - The h23 is connected to the port "s2-eth3".

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn-port-map:set-port-map -d '{"input":{"vlan-id
↪":0,"tenant-name":"vtn1","bridge-name":"vbr1","interface-name":"if3","node":
↪"openflow:2","port-name":"s2-eth3"}}'
```

Traffic filtering

- Create flowcondition named cond_1 by executing the set-flow-condition RPC.
  - For option source and destination-network, get inet address of host h12(src) and h22(dst) from mininet.

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn-flow-condition:set-flow-condition -d '{"input
↪":{"operation":"SET","present":"false","name":"cond_1","vtn-flow-match":[{"index":1,
↪"vtn-ether-match":{},"vtn-inet-match":{"source-network":"10.0.0.2/32","destination-
↪network":"10.0.0.4/32"}}]}}'
```

- Flow filter demonstration with DROP action-type. Create Flowfilter in VBR Interface if1 by executing the set-flow-filter RPC.

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn-flow-filter:set-flow-filter -d '{"input":{
↪"output":"false","tenant-name":"vtn1","bridge-name":"vbr1","interface-name":"if1",
↪"vtn-flow-filter":[{"condition":"cond_1","index":10,"vtn-drop-filter":{}}]}}'
```

Service Chaining

With One Service

- Please execute the below commands to configure the network topology which sends some specific traffic via a single service(External device) in the controller as in the below image:
- Create a virtual terminal named vt_srvc1_1 in the tenant vtn1 by executing the update-vterminal RPC.

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn-vterminal:update-vterminal -d '{"input":{
↪"update-mode":"CREATE","operation":"SET","tenant-name":"vtn1","terminal-name":"vt_
↪srvc1_1","description":"Creating vterminal"}}'
```

(continues on next page)

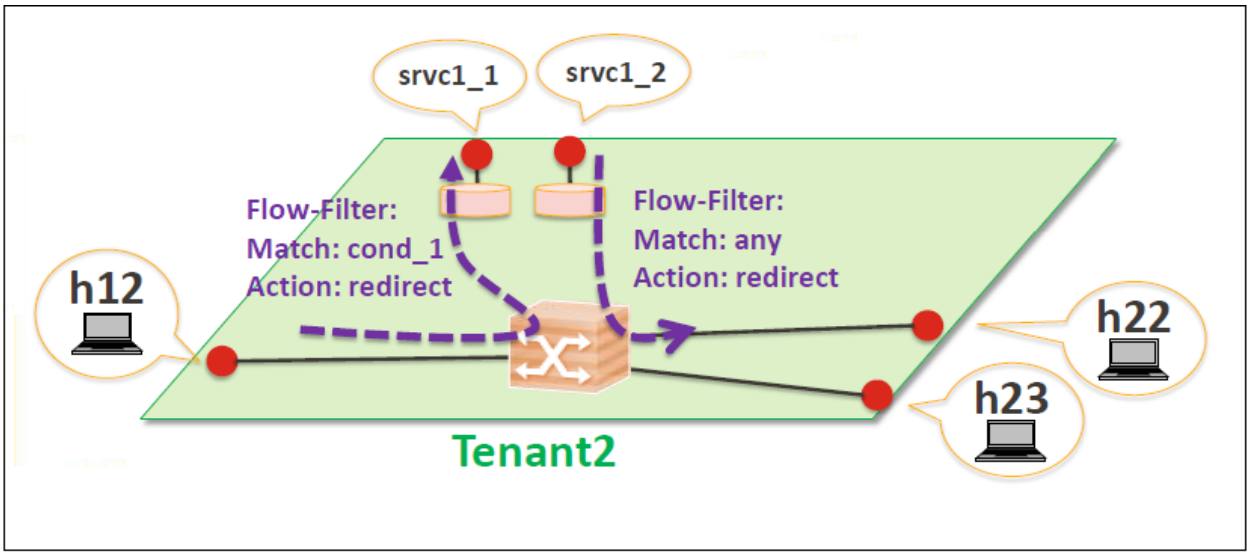

Fig. 133: Service Chaining With One Service LLD

(continued from previous page)

- Create interface IF into the virtual terminal vt_srvc1_1 by executing the update-vinterface RPC.

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn-vinterface:update-vinterface -d '{"input":{
↪"update-mode":"CREATE","operation":"SET","description":"Creating vterminal IF",
↪"enabled":"true","tenant-name":"vtn1","terminal-name":"vt_srvc1_1","interface-name":
↪"IF"}}'
```

- Configure port mapping on the interfaces by executing the set-port-map RPC.
  - The interface IF of the virtual terminal will be mapped to the port "s3-eth3" of the switch "openflow:3" of the Mininet.
    - The h12 is connected to the port "s3-eth3".

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn-port-map:set-port-map -d '{"input":{"tenant-
↪name":"vtn1","terminal-name":"vt_srvc1_1","interface-name":"IF","node":"openflow:3",
↪"port-name":"s3-eth3"}}'
```

- Create a virtual terminal named vt_srvc1_2 in the tenant vtn1 by executing the update-vterminal RPC.

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn-vterminal:update-vterminal -d '{"input":{
↪"update-mode":"CREATE","operation":"SET","tenant-name":"vtn1","terminal-name":"vt_
↪srvc1_2","description":"Creating vterminal"}}'
```

- Create interface IF into the virtual terminal vt_srvc1_2 by executing the update-vinterface RPC.

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn-vinterface:update-vinterface -d '{"input":{
↪"update-mode":"CREATE","operation":"SET","description":"Creating vterminal IF",
↪"enabled":"true","tenant-name":"vtn1","terminal-name":"vt_srvc1_2","interface-name":
↪"IF"}}'
```

(continues on next page)

(continued from previous page)

- Configure port mapping on the interfaces by executing the set-port-map RPC.
  - The interface IF of the virtual terminal will be mapped to the port "s4-eth3" of the switch "openflow:4" of the Mininet.
    - The h22 is connected to the port "s4-eth3".

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn-port-map:set-port-map -d '{"input":{"tenant-
↪name":"vtn1","terminal-name":"vt_srvc1_2","interface-name":"IF","node":"openflow:4",
↪"port-name":"s4-eth3"}}'
```

- Create flowcondition named cond_1 by executing the set-flow-condition RPC.
  - For option source and destination-network, get inet address of host h12(src) and h22(dst) from mininet.

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn-flow-condition:set-flow-condition -d '{"input
↪":{"operation":"SET","present":"false","name":"cond_1","vtn-flow-match":[{"index":1,
↪"vtn-ether-match":{},"vtn-inet-match":{"source-network":"10.0.0.2/32","destination-
↪network":"10.0.0.4/32"}}]}}'
```

- Create flowcondition named cond_any by executing the set-flow-condition RPC.

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn-flow-condition:set-flow-condition -d '{"input
↪":{"operation":"SET","present":"false","name":"cond_any","vtn-flow-match":[{"index
↪":1}]}}'
```

- Flow filter demonstration with redirect action-type. Create Flowfilter in virtual terminal vt_srvc1_2 interface IF by executing the set-flow-filter RPC.
  - Flowfilter redirects vt_srvc1_2 to bridge1-IF2

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn-flow-filter:set-flow-filter -d '{"input":{
↪"output":"false","tenant-name":"vtn1","terminal-name":"vt_srvc1_2","interface-name":
↪"IF","vtn-flow-filter":[{"condition":"cond_any","index":10,"vtn-redirect-filter":{
↪"redirect-destination":{"bridge-name":"vbr1","interface-name":"if2"},"output":"true
↪"}}]}}'
```

- Flow filter demonstration with redirect action-type. Create Flowfilter in vbridge vbr1 interface if1 by executing the set-flow-filter RPC.
  - Flow filter redirects Bridge1-IF1 to vt_srvc1_1

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn-flow-filter:set-flow-filter -d '{"input":{
↪"output":"false","tenant-name":"vtn1","bridge-name":"vbr1","interface-name":"if1",
↪"vtn-flow-filter":[{"condition":"cond_1","index":10,"vtn-redirect-filter":{
↪"redirect-destination":{"terminal-name":"vt_srvc1_1","interface-name":"IF"},"output
↪":"true"}}]}}'
```

Verification

- Ping host12 to host22 to view the host rechability, a delay of 200ms will be taken to reach host22 as below.

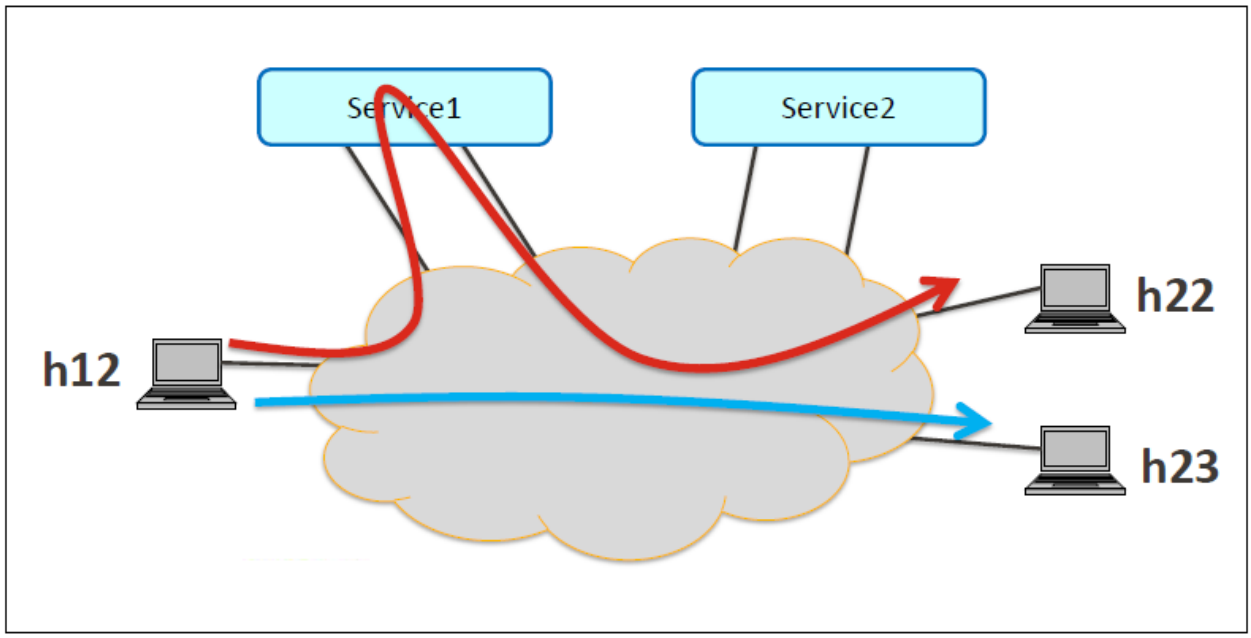

Fig. 134: Service Chaining With One Service

```
mininet> h12 ping h22
PING 10.0.0.4 (10.0.0.4) 56(84) bytes of data.
64 bytes from 10.0.0.4: icmp_seq=35 ttl=64 time=209 ms
64 bytes from 10.0.0.4: icmp_seq=36 ttl=64 time=201 ms
64 bytes from 10.0.0.4: icmp_seq=37 ttl=64 time=200 ms
64 bytes from 10.0.0.4: icmp_seq=38 ttl=64 time=200 ms
```

With two services

- Please execute the below commands to configure the network topology which sends some specific traffic via two services(External device) in the controller as in the below image.

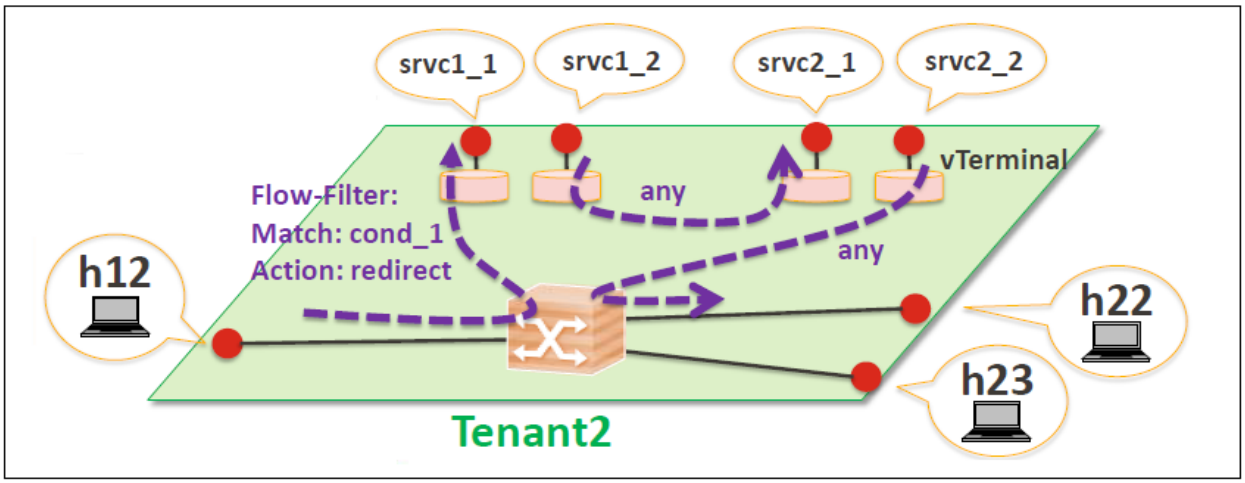

Fig. 135: Service Chaining With Two Services LLD

- Create a virtual terminal named vt_srvc2_1 in the tenant vtn1 by executing the update-vterminal RPC.

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn-vterminal:update-vterminal -d '{"input":{
↪"update-mode":"CREATE","operation":"SET","tenant-name":"vtn1","terminal-name":"vt_
↪srvc2_1","description":"Creating vterminal"}}'
```

- Create interface IF into the virtual terminal vt_srvc2_1 by executing the update-vinterface RPC.

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn-vinterface:update-vinterface -d '{"input":{
↪"update-mode":"CREATE","operation":"SET","description":"Creating vterminal IF",
↪"enabled":"true","tenant-name":"vtn1","terminal-name":"vt_srvc2_1","interface-name":
↪"IF"}}'
```

- Configure port mapping on the interfaces by executing the set-port-map RPC.
  - The interface IF of the virtual terminal will be mapped to the port "s3-eth4" of the switch "openflow:3" of the Mininet.
    - The host h12 is connected to the port "s3-eth4".

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn-port-map:set-port-map -d '{"input":{"tenant-
↪name":"vtn1","terminal-name":"vt_srvc2_1","interface-name":"IF","node":"openflow:3",
↪"port-name":"s3-eth4"}}'
```

- Create a virtual terminal named vt_srvc2_2 in the tenant vtn1 by executing the update-vterminal RPC.

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn-vterminal:update-vterminal -d '{"input":{
↪"update-mode":"CREATE","operation":"SET","tenant-name":"vtn1","terminal-name":"vt_
↪srvc2_2","description":"Creating vterminal"}}'
```

- Create interfaces IF into the virtual terminal vt_srvc2_2 by executing the update-vinterface RPC.

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn-vinterface:update-vinterface -d '{"input":{
↪"update-mode":"CREATE","operation":"SET","description":"Creating vterminal IF",
↪"enabled":"true","tenant-name":"vtn1","terminal-name":"vt_srvc2_2","interface-name":
↪"IF"}}'
```

- Configure port mapping on the interfaces by executing the set-port-map RPC.
  - The interface IF of the virtual terminal will be mapped to the port "s4-eth4" of the switch "openflow:4" of the mininet.
    - The host h22 is connected to the port "s4-eth4".

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn-port-map:set-port-map -d '{"input":{"tenant-
↪name":"vtn1","terminal-name":"vt_srvc2_2","interface-name":"IF","node":"openflow:4",
↪"port-name":"s4-eth4"}}'
```

- Flow filter demonstration with redirect action-type. Create Flowfilter in virtual terminal vt_srvc2_2 interface IF by executing the set-flow-filter RPC.
  - Flow filter redirects vt_srvc2_2 to Bridge1-IF2.

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn-flow-filter:set-flow-filter -d '{"input":{
↪"output":"false","tenant-name":"vtn1","terminal-name":"vt_srvc2_2","interface-name":
↪"IF","vtn-flow-filter":[{"condition":"cond_any","index":10,"vtn-redirect-filter":{
↪"redirect-destination":{"bridge-name":"vbr1","interface-name":"if2"},"output":"true
↪"}}]}}'
```

- Flow filter demonstration with redirect action-type. Create Flowfilter in virtual terminal vt_srvc2_2 interface IF by executing the set-flow-filter RPC.
  - Flow filter redirects vt_srvc1_2 to vt_srvc2_1.

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn-flow-filter:set-flow-filter -d '{"input":{
↪"output":"false","tenant-name":"vtn1","terminal-name":"vt_srvc1_2","interface-name":
↪"IF","vtn-flow-filter":[{"condition":"cond_any","index":10,"vtn-redirect-filter":{
↪"redirect-destination":{"terminal-name":"vt_srvc2_1","interface-name":"IF"},"output
↪":"true"}}]}}'
```

Verification

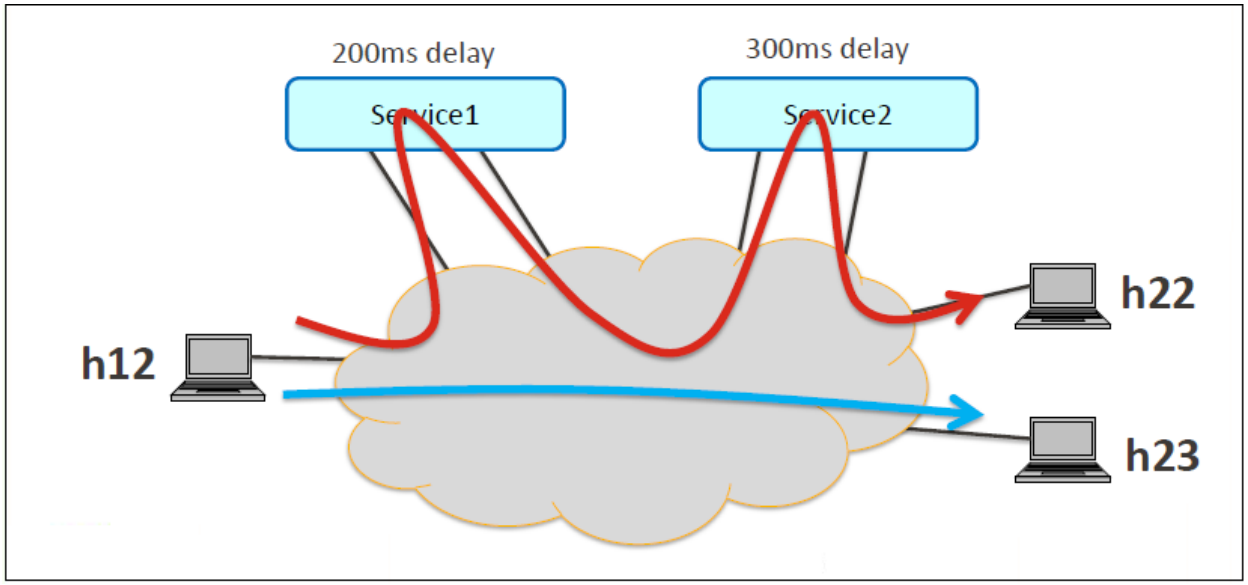

Fig. 136: Service Chaining With Two Service

- Ping host12 to host22 to view the host rechability, a delay of 500ms will be taken to reach host22 as below.

```
mininet> h12 ping h22
PING 10.0.0.4 (10.0.0.4) 56(84) bytes of data.
64 bytes from 10.0.0.4: icmp_seq=1 ttl=64 time=512 ms
64 bytes from 10.0.0.4: icmp_seq=2 ttl=64 time=501 ms
64 bytes from 10.0.0.4: icmp_seq=3 ttl=64 time=500 ms
64 bytes from 10.0.0.4: icmp_seq=4 ttl=64 time=500 ms
```

- You can verify the configuration by executing the following REST API. It shows all configuration in VTN Manager.

```
curl --user "admin":"admin" -H "Content-type: application/json" -X GET http://
↪localhost:8181/restconf/operational/vtn:vtns
```


```
{
  "vtn": [
  {
    "name": "vtn1",
      "vtenant-config": {
        "hard-timeout": 0,
        "idle-timeout": 300,
        "description": "creating vtn"
      },
      "vbridge": [
      {
        "name": "vbr1",
        "vbridge-config": {
          "age-interval": 600,
          "description": "creating vbr"
        },
        "bridge-status": {
          "state": "UP",
          "path-faults": 0
        },
        "vinterface": [
        {
          "name": "if1",
          "vinterface-status": {
            "mapped-port": "openflow:1:2",
            "state": "UP",
            "entity-state": "UP"
          },
          "port-map-config": {
            "vlan-id": 0,
            "node": "openflow:1",
            "port-name": "s1-eth2"
          },
          "vinterface-config": {
            "description": "Creating vbrif1 interface",
            "enabled": true
          },
          "vinterface-input-filter": {
            "vtn-flow-filter": [
            {
              "index": 10,
              "condition": "cond_1",
              "vtn-redirect-filter": {
                "output": true,
                "redirect-destination": {
                  "terminal-name": "vt_srvc1_1",
                  "interface-name": "IF"
                }
              }
            }
            ]
          }
        },
        {
```


(continues on next page)

(continued from previous page)


```
          "name": "if2",
          "vinterface-status": {
            "mapped-port": "openflow:2:2",
            "state": "UP",
            "entity-state": "UP"
          },
          "port-map-config": {
            "vlan-id": 0,
            "node": "openflow:2",
            "port-name": "s2-eth2"
          },
          "vinterface-config": {
            "description": "Creating vbrif2 interface",
            "enabled": true
          }
        },
        {
          "name": "if3",
          "vinterface-status": {
            "mapped-port": "openflow:2:3",
            "state": "UP",
            "entity-state": "UP"
          },
          "port-map-config": {
            "vlan-id": 0,
            "node": "openflow:2",
            "port-name": "s2-eth3"
          },
          "vinterface-config": {
            "description": "Creating vbrif3 interface",
            "enabled": true
          }
        }
        ]
      }
    ],
      "vterminal": [
      {
        "name": "vt_srvc2_2",
        "bridge-status": {
          "state": "UP",
          "path-faults": 0
        },
        "vinterface": [
        {
          "name": "IF",
          "vinterface-status": {
            "mapped-port": "openflow:4:4",
            "state": "UP",
            "entity-state": "UP"
          },
          "port-map-config": {
            "vlan-id": 0,
            "node": "openflow:4",
            "port-name": "s4-eth4"
          },
          "vinterface-config": {
```


(continues on next page)

(continued from previous page)


```
            "description": "Creating vterminal IF",
            "enabled": true
          },
          "vinterface-input-filter": {
            "vtn-flow-filter": [
            {
              "index": 10,
              "condition": "cond_any",
              "vtn-redirect-filter": {
                "output": true,
                "redirect-destination": {
                  "bridge-name": "vbr1",
                  "interface-name": "if2"
                }
              }
            }
            ]
          }
        }
        ],
          "vterminal-config": {
            "description": "Creating vterminal"
          }
      },
      {
        "name": "vt_srvc1_1",
        "bridge-status": {
          "state": "UP",
          "path-faults": 0
        },
        "vinterface": [
        {
          "name": "IF",
          "vinterface-status": {
            "mapped-port": "openflow:3:3",
            "state": "UP",
            "entity-state": "UP"
          },
          "port-map-config": {
            "vlan-id": 0,
            "node": "openflow:3",
            "port-name": "s3-eth3"
          },
          "vinterface-config": {
            "description": "Creating vterminal IF",
            "enabled": true
          }
        }
        ],
          "vterminal-config": {
            "description": "Creating vterminal"
          }
      },
      {
        "name": "vt_srvc1_2",
        "bridge-status": {
          "state": "UP",
```


(continues on next page)

(continued from previous page)


```
          "path-faults": 0
        },
        "vinterface": [
        {
          "name": "IF",
          "vinterface-status": {
            "mapped-port": "openflow:4:3",
            "state": "UP",
            "entity-state": "UP"
          },
          "port-map-config": {
            "vlan-id": 0,
            "node": "openflow:4",
            "port-name": "s4-eth3"
          },
          "vinterface-config": {
            "description": "Creating vterminal IF",
            "enabled": true
          },
          "vinterface-input-filter": {
            "vtn-flow-filter": [
            {
              "index": 10,
              "condition": "cond_any",
              "vtn-redirect-filter": {
                "output": true,
                "redirect-destination": {
                  "terminal-name": "vt_srvc2_1",
                  "interface-name": "IF"
                }
              }
            }
            ]
          }
        }
        ],
          "vterminal-config": {
            "description": "Creating vterminal"
          }
      },
      {
        "name": "vt_srvc2_1",
        "bridge-status": {
          "state": "UP",
          "path-faults": 0
        },
        "vinterface": [
        {
          "name": "IF",
          "vinterface-status": {
            "mapped-port": "openflow:3:4",
            "state": "UP",
            "entity-state": "UP"
          },
          "port-map-config": {
            "vlan-id": 0,
            "node": "openflow:3",
```


(continues on next page)

(continued from previous page)

```
            "port-name": "s3-eth4"
          },
          "vinterface-config": {
            "description": "Creating vterminal IF",
            "enabled": true
          }
        }
        ],
          "vterminal-config": {
            "description": "Creating vterminal"
          }
      }
    ]
  }
  ]
}
```

Cleaning Up

- To clean up both VTN and flowconditions.
- You can delete the virtual tenant vtn1 by executing the remove-vtn RPC.

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn:remove-vtn -d '{"input":{"tenant-name":"vtn1
↪"}}'
```

- You can delete the flowcondition cond_1 and cond_any by executing the remove-flow-condition RPC.

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn-flow-condition:remove-flow-condition -d '{
↪"input":{"name":"cond_1"}}'
```

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn-flow-condition:remove-flow-condition -d '{
↪"input":{"name":"cond_any"}}'
```

How To View Dataflows

Overview

This page explains how to view Dataflows using VTN Manager. This page targets Boron release, so the procedure described here does not work in other releases.

Dataflow feature enables retrieval and display of data flows in the OpenFlow network. The data flows can be retrieved based on an OpenFlow switch or a switch port or a L2 source host.

The flow information provided by this feature are

- Location of virtual node which maps the incoming packet and outgoing packets.
- Location of physical switch port where incoming and outgoing packets is sent and received.
- A sequence of physical route info which represents the packet route in the physical network.

Configuration

- To view Dataflow information, configure with VLAN Mapping https://wiki.opendaylight.org/view/VTN:Mananger:How_to_test_Vlan-map_using_mininet.

Verification

After creating vlan mapping configuration from the above page, execute as below in mininet to get switch details.

```
mininet> net
h1 h1-eth0.200:s1-eth1
h2 h2-eth0.300:s2-eth2
h3 h3-eth0.200:s2-eth3
h4 h4-eth0.300:s2-eth4
h5 h5-eth0.200:s3-eth2
h6 h6-eth0.300:s3-eth3
s1 lo:  s1-eth1:h1-eth0.200 s1-eth2:s2-eth1 s1-eth3:s3-eth1
s2 lo:  s2-eth1:s1-eth2 s2-eth2:h2-eth0.300 s2-eth3:h3-eth0.200 s2-eth4:h4-eth0.300
s3 lo:  s3-eth1:s1-eth3 s3-eth2:h5-eth0.200 s3-eth3:h6-eth0.300
c0
mininet>
```

Please execute ping from h1 to h3 to check hosts reachability.

```
mininet> h1 ping h3
PING 10.0.0.3 (10.0.0.3) 56(84) bytes of data.
64 bytes from 10.0.0.3: icmp_seq=1 ttl=64 time=11.4 ms
64 bytes from 10.0.0.3: icmp_seq=2 ttl=64 time=0.654 ms
64 bytes from 10.0.0.3: icmp_seq=3 ttl=64 time=0.093 ms
```

Parallely execute below Restconf command to get data flow information of node "openflow:1" and its port "s1-eth1".

- Get the Dataflows information by executing the get-data-flow RPC.

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn-flow:get-data-flow -d '{"input":{"tenant-name
↪":"vtn1","mode":"DETAIL","node":"openflow:1","data-flow-port":{"port-id":"1","port-
↪name":"s1-eth1"}}}'
```

```
{
  "output": {
    "data-flow-info": [
    {
      "averaged-data-flow-stats": {
        "packet-count": 1.1998800119988002,
          "start-time": 1455241209151,
          "end-time": 1455241219152,
          "byte-count": 117.58824117588242
      },
        "physical-route": [
        {
          "physical-ingress-port": {
            "port-name": "s2-eth3",
            "port-id": "3"
          },
          "physical-egress-port": {
```

(continues on next page)

(continued from previous page)


```
            "port-name": "s2-eth1",
            "port-id": "1"
          },
          "node": "openflow:2",
          "order": 0
        },
        {
          "physical-ingress-port": {
            "port-name": "s1-eth2",
            "port-id": "2"
          },
          "physical-egress-port": {
            "port-name": "s1-eth1",
            "port-id": "1"
          },
          "node": "openflow:1",
          "order": 1
        }
      ],
        "data-egress-node": {
          "bridge-name": "vbr1",
          "tenant-name": "vtn1"
        },
        "hard-timeout": 0,
        "idle-timeout": 300,
        "data-flow-stats": {
          "duration": {
            "nanosecond": 640000000,
            "second": 362
          },
          "packet-count": 134,
          "byte-count": 12932
        },
        "data-egress-port": {
          "node": "openflow:1",
          "port-name": "s1-eth1",
          "port-id": "1"
        },
        "data-ingress-node": {
          "bridge-name": "vbr1",
          "tenant-name": "vtn1"
        },
        "data-ingress-port": {
          "node": "openflow:2",
          "port-name": "s2-eth3",
          "port-id": "3"
        },
        "creation-time": 1455240855753,
        "data-flow-match": {
          "vtn-ether-match": {
            "vlan-id": 200,
            "source-address": "6a:ff:e2:81:86:bb",
            "destination-address": "26:9f:82:70:ec:66"
          }
        },
        "virtual-route": [
        {
```


(continues on next page)

(continued from previous page)


```
          "reason": "VLANMAPPED",
          "virtual-node-path": {
            "bridge-name": "vbr1",
            "tenant-name": "vtn1"
          },
          "order": 0
        },
        {
          "reason": "FORWARDED",
          "virtual-node-path": {
            "bridge-name": "vbr1",
            "tenant-name": "vtn1"
          },
          "order": 1
        }
      ],
        "flow-id": 16
    },
    {
      "averaged-data-flow-stats": {
        "packet-count": 1.1998800119988002,
        "start-time": 1455241209151,
        "end-time": 1455241219152,
        "byte-count": 117.58824117588242
      },
      "physical-route": [
      {
        "physical-ingress-port": {
          "port-name": "s1-eth1",
          "port-id": "1"
        },
        "physical-egress-port": {
          "port-name": "s1-eth2",
          "port-id": "2"
        },
        "node": "openflow:1",
        "order": 0
      },
      {
        "physical-ingress-port": {
          "port-name": "s2-eth1",
          "port-id": "1"
        },
        "physical-egress-port": {
          "port-name": "s2-eth3",
          "port-id": "3"
        },
        "node": "openflow:2",
        "order": 1
      }
      ],
        "data-egress-node": {
          "bridge-name": "vbr1",
          "tenant-name": "vtn1"
        },
        "hard-timeout": 0,
        "idle-timeout": 300,
```


(continues on next page)

(continued from previous page)


```
        "data-flow-stats": {
          "duration": {
            "nanosecond": 587000000,
            "second": 362
          },
          "packet-count": 134,
          "byte-count": 12932
        },
        "data-egress-port": {
          "node": "openflow:2",
          "port-name": "s2-eth3",
          "port-id": "3"
        },
        "data-ingress-node": {
          "bridge-name": "vbr1",
          "tenant-name": "vtn1"
        },
        "data-ingress-port": {
          "node": "openflow:1",
          "port-name": "s1-eth1",
          "port-id": "1"
        },
        "creation-time": 1455240855747,
        "data-flow-match": {
          "vtn-ether-match": {
            "vlan-id": 200,
            "source-address": "26:9f:82:70:ec:66",
            "destination-address": "6a:ff:e2:81:86:bb"
          }
        },
        "virtual-route": [
        {
          "reason": "VLANMAPPED",
          "virtual-node-path": {
            "bridge-name": "vbr1",
            "tenant-name": "vtn1"
          },
          "order": 0
        },
        {
          "reason": "FORWARDED",
          "virtual-node-path": {
            "bridge-name": "vbr1",
            "tenant-name": "vtn1"
          },
          "order": 1
        }
      ],
        "flow-id": 15
    }
    ]
  }
}
```

How To Create Mac Map In VTN

Overview

- This page demonstrates Mac Mapping. This demonstration aims at enabling communication between two hosts and denying communication of particular host by associating a Vbridge to the hosts and configuring Mac Mapping (mac address) to the Vbridge.
- This page targets Boron release, so the procedure described here does not work in other releases.

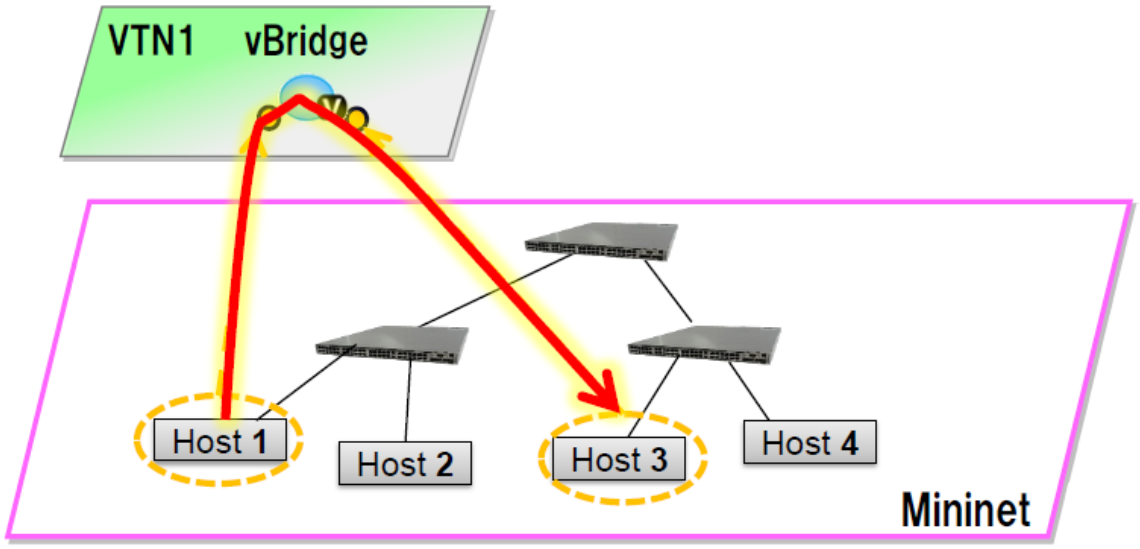

Fig. 137: Single Controller Mapping

Requirement

Configure mininet and create a topology

- Script for emulating network with multiple hosts.
- Before executing the mininet script, please confirm Controller is up and running.
- Run the mininet script.
- Replace <path> and <Controller IP> based on your environment.

```
sudo mn --controller=remote,ip=<Controller IP> --custom <path>\topo_handson.py --topo␣
↪mytopo2
```

```
mininet> net
h11 h11-eth0:s1-eth1
h12 h12-eth0:s1-eth2
h21 h21-eth0:s2-eth1
h22 h22-eth0:s2-eth2
h23 h23-eth0:s2-eth3
srvc1 srvc1-eth0:s3-eth3 srvc1-eth1:s4-eth3
srvc2 srvc2-eth0:s3-eth4 srvc2-eth1:s4-eth4
```

(continues on next page)

(continued from previous page)

```
s1 lo:  s1-eth1:h11-eth0 s1-eth2:h12-eth0 s1-eth3:s2-eth4 s1-eth4:s3-eth2
s2 lo:  s2-eth1:h21-eth0 s2-eth2:h22-eth0 s2-eth3:h23-eth0 s2-eth4:s1-eth3 s2-eth5:s4-
↪eth1
s3 lo:  s3-eth1:s4-eth2 s3-eth2:s1-eth4 s3-eth3:srvc1-eth0 s3-eth4:srvc2-eth0
s4 lo:  s4-eth1:s2-eth5 s4-eth2:s3-eth1 s4-eth3:srvc1-eth1 s4-eth4:srvc2-eth1
```

Configuration

To create Mac Map in VTN, execute REST API provided by VTN Manager as follows. It uses curl command to call REST API.

- Create a virtual tenant named Tenant1 by executing the update-vtn RPC.

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn:update-vtn -d '{"input":{"tenant-name":
↪"Tenant1"}}'
```

- Create a virtual bridge named vBridge1 in the tenant Tenant1 by executing the update-vbridge RPC.

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn-vbridge:update-vbridge -d '{"input":{"tenant-
↪name":"Tenant1","bridge-name":"vBridge1"}}'
```

- Configuring Mac Mappings on the vBridge1 by giving the mac address of host h12 and host h22 as follows to allow the communication by executing the set-mac-map RPC.

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn-mac-map:set-mac-map -d '{"input":{"operation
↪":"SET","allowed-hosts":["de:05:40:c4:96:76@0","62:c5:33:bc:d7:4e@0"],"tenant-name":
↪"Tenant1","bridge-name":"vBridge1"}}'
```

Note: Mac Address of host h12 and host h22 can be obtained with the following command in mininet.

```
mininet> h12 ifconfig
h12-eth0  Link encap:Ethernet  HWaddr 62:c5:33:bc:d7:4e
inet addr:10.0.0.2  Bcast:10.255.255.255  Mask:255.0.0.0
inet6 addr: fe80::60c5:33ff:febc:d74e/64 Scope:Link
```

```
mininet> h22 ifconfig
h22-eth0  Link encap:Ethernet  HWaddr de:05:40:c4:96:76
inet addr:10.0.0.4  Bcast:10.255.255.255  Mask:255.0.0.0
inet6 addr: fe80::dc05:40ff:fec4:9676/64 Scope:Link
```

- MAC Mapping will not be activated just by configuring it, a two end communication needs to be established to activate Mac Mapping.
- Ping host h22 from host h12 in mininet, the ping will not happen between the hosts as only one way activation is enabled.

```
mininet> h12 ping h22
PING 10.0.0.4 (10.0.0.4) 56(84) bytes of data.
From 10.0.0.2 icmp_seq=1 Destination Host Unreachable
From 10.0.0.2 icmp_seq=2 Destination Host Unreachable
```

- Ping host h12 from host h22 in mininet, now the ping communication will take place as the two end communication is enabled.

```
mininet> h22 ping h12
PING 10.0.0.2 (10.0.0.2) 56(84) bytes of data.
64 bytes from 10.0.0.2: icmp_req=1 ttl=64 time=91.8 ms
64 bytes from 10.0.0.2: icmp_req=2 ttl=64 time=0.510 ms
```

- After two end communication enabled, now host h12 can ping host h22

```
mininet> h12 ping h22
PING 10.0.0.4 (10.0.0.4) 56(84) bytes of data.
64 bytes from 10.0.0.4: icmp_req=1 ttl=64 time=0.780 ms
64 bytes from 10.0.0.4: icmp_req=2 ttl=64 time=0.079 ms
```

Verification

- To view the configured Mac Map of allowed host execute the following command.

```
curl --user "admin":"admin" -H "Content-type: application/json" -X GET http://
↪localhost:8181/restconf/operational/vtn:vtns/vtn/Tenant1/vbridge/vBridge1/mac-map
```

```
{
  "mac-map": {
    "mac-map-status": {
      "mapped-host": [
      {
        "mac-address": "c6:44:22:ba:3e:72",
          "vlan-id": 0,
          "port-id": "openflow:1:2"
      },
      {
        "mac-address": "f6:e0:43:b6:3a:b7",
        "vlan-id": 0,
        "port-id": "openflow:2:2"
      }
      ]
    },
      "mac-map-config": {
        "allowed-hosts": {
          "vlan-host-desc-list": [
          {
            "host": "c6:44:22:ba:3e:72@0"
          },
          {
            "host": "f6:e0:43:b6:3a:b7@0"
          }
          ]
        }
      }
  }
}
```

Note: When Deny is configured a broadcast message is sent to all the hosts connected to the vBridge, so a two end communication need not be establihed like allow, the hosts can communicate directly without any two way communi-

cation enabled.

1. To Deny host h23 communication from hosts connected on vBridge1, the following configuration can be applied.

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn-mac-map:set-mac-map -d '{"input":{"operation
↪": "SET", "denied-hosts": ["0a:d3:ea:3d:8f:a5@0"],"tenant-name": "Tenant1","bridge-
↪name": "vBridge1"}}'
```

Cleaning Up

- You can delete the virtual tenant Tenant1 by executing the remove-vtn RPC.

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn:remove-vtn -d '{"input":{"tenant-name":
↪"Tenant1"}}'
```

How To Configure Flowfilters

Overview

- This page explains how to provision flowfilter using VTN Manager. This page targets Boron release, so the procedure described here does not work in other releases.
- The flow-filter function discards, permits, or redirects packets of the traffic within a VTN, according to specified flow conditions. The table below lists the actions to be applied when a packet matches the condition:

Action	Function
Pass	Permits the packet to pass along the determined path. As options, packet transfer priority (set priority) and DSCP change (set ip-dscp) is specified.
Drop	Discards the packet.
Redirect	Redirects the packet to a desired virtual interface. As an option, it is possible to change the MAC address when the packet is transferred.

- Following steps explain flow-filter function:
  - when a packet is transferred to an interface within a virtual network, the flow-filter function evaluates whether the transferred packet matches the condition specifed in the flow-list.
  - If the packet matches the condition, the flow-filter applies the flow-list matching action specified in the flow-filter.

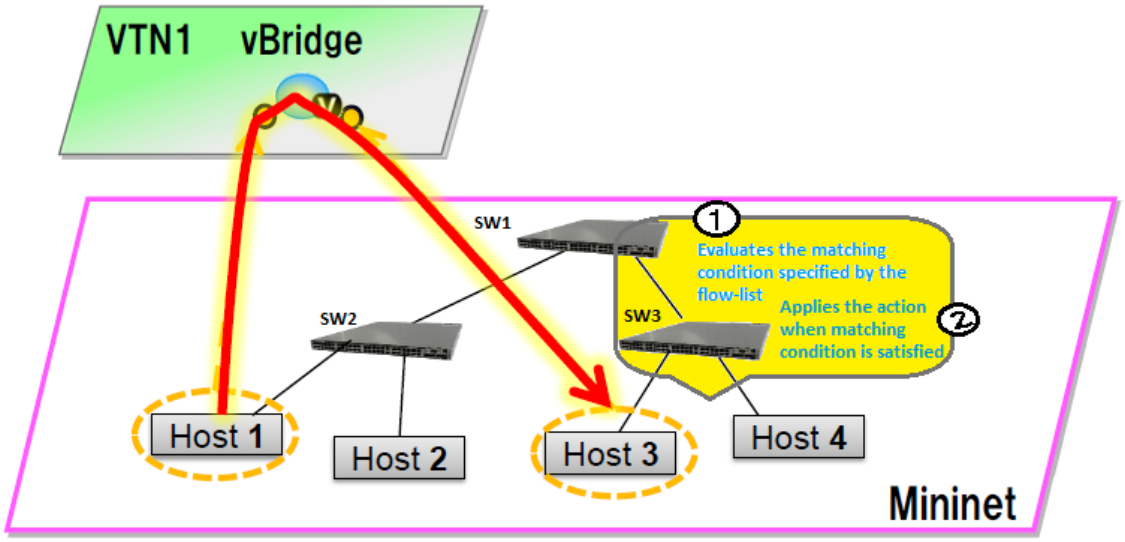

Fig. 138: Flow Filter Example

Requirements

To apply the packet filter, configure the following:

- Create a flow condition.
- Specify where to apply the flow-filter, for example VTN, vBridge, or interface of vBridge.

To provision OpenFlow switches, this page uses Mininet. Mininet details and set-up can be referred at the below page:
https://wiki.opendaylight.org/view/OpenDaylight_Controller:Installation#Using_Mininet

Start Mininet, and create three switches (s1, s2, and s3) and four hosts (h1, h2, h3 and h4) in it.

```
sudo mn --controller=remote,ip=192.168.0.100 --topo tree,2
```

Note: Replace "192.168.0.100" with the IP address of OpenDaylight controller based on your environment.

You can check the topology that you have created by executing "net" command in the Mininet console.

```
mininet> net
h1 h1-eth0:s2-eth1
h2 h2-eth0:s2-eth2
h3 h3-eth0:s3-eth1
h4 h4-eth0:s3-eth2
s1 lo:  s1-eth1:s2-eth3 s1-eth2:s3-eth3
s2 lo:  s2-eth1:h1-eth0 s2-eth2:h2-eth0 s2-eth3:s1-eth1
s3 lo:  s3-eth1:h3-eth0 s3-eth2:h4-eth0 s3-eth3:s1-eth2
```

In this guide, you will provision flowfilters to establish communication between h1 and h3.

Configuration

To provision the virtual L2 network for the two hosts (h1 and h3), execute REST API provided by VTN Manager as follows. It uses curl command to call the REST API.

- Create a virtual tenant named vtn1 by executing the update-vtn RPC.

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn:update-vtn -d '{"input":{"tenant-name":"vtn1
↪"}}'
```

- Create a virtual bridge named vbr1 in the tenant vtn1 by executing the update-vbridge RPC.

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn-vbridge:update-vbridge -d '{"input":{"tenant-
↪name":"vtn1","bridge-name":"vbr1"}}'
```

- Create two interfaces into the virtual bridge by executing the update-vinterface RPC.

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn-vinterface:update-vinterface -d '{"input":{
↪"tenant-name":"vtn1","bridge-name":"vbr1","interface-name":"if1"}}'
```

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn-vinterface:update-vinterface -d '{"input":{
↪"tenant-name":"vtn1","bridge-name":"vbr1","interface-name":"if2"}}'
```

- Configure two mappings on the interfaces by executing the set-port-map RPC.
  - The interface if1 of the virtual bridge will be mapped to the port "s2-eth1" of the switch "openflow:2" of the Mininet.
    - The h1 is connected to the port "s2-eth1".
  - The interface if2 of the virtual bridge will be mapped to the port "s3-eth1" of the switch "openflow:3" of the Mininet.
    - The h3 is connected to the port "s3-eth1".

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn-port-map:set-port-map -d '{"input":{"tenant-
↪name":"vtn1", "bridge-name":"vbr1", "interface-name":"if1", "node":"openflow:2",
↪"port-name":"s2-eth1"}}'
```

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn-port-map:set-port-map -d '{"input":{"tenant-
↪name":"vtn1", "bridge-name":"vbr1", "interface-name":"if2", "node":"openflow:3",
↪"port-name":"s3-eth1"}}'
```

- Create flowcondition named cond_1 by executing the set-flow-condition RPC.
  - For option source and destination-network, get inet address of host h1 and h3 from mininet.

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn-flow-condition:set-flow-condition -d '{"input
↪":{"name":"cond_1", "vtn-flow-match":[{"vtn-ether-match":{},"vtn-inet-match":{
↪"source-network":"10.0.0.1/32","protocol":1,"destination-network":"10.0.0.3/32"},
↪"index":"1"}]}}'
```

- Flowfilter can be applied either in VTN, VBR or VBR Interfaces. Here in this page we provision flowfilter with VBR Interface and demonstrate with action type drop and then pass.
- Flow filter demonstration with DROP action-type. Create Flowfilter in VBR Interface if1 by executing the set-flow-filter RPC.

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn-flow-filter:set-flow-filter -d '{"input": {
↪"tenant-name": "vtn1", "bridge-name": "vbr1","interface-name":"if1","vtn-flow-filter
↪":[{"condition":"cond_1","vtn-drop-filter":{},"vtn-flow-action":[{"order": "1","vtn-
↪set-inet-src-action":{"ipv4-address":"10.0.0.1/32"}},{"order": "2","vtn-set-inet-
↪dst-action":{"ipv4-address":"10.0.0.3/32"}}],"index": "1"}]}}'
```

Verification of the drop filter

- Please execute ping from h1 to h3. As we have applied the action type "drop" , ping should fail with no packet flows between hosts h1 and h3 as below,

```
mininet> h1 ping h3
```

Configuration for pass filter

- Update the flow filter to pass the packets by executing the set-flow-filter RPC.

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn-flow-filter:set-flow-filter -d '{"input": {
↪"tenant-name": "vtn1", "bridge-name": "vbr1","interface-name":"if1","vtn-flow-filter
↪":[{"condition":"cond_1","vtn-pass-filter":{},"vtn-flow-action":[{"order": "1","vtn-
↪set-inet-src-action":{"ipv4-address":"10.0.0.1/32"}},{"order": "2","vtn-set-inet-
↪dst-action":{"ipv4-address":"10.0.0.3/32"}}],"index": "1"}]}}'
```

Verification For Packets Success

- As we have applied action type PASS now ping should happen between hosts h1 and h3.

```
mininet> h1 ping h3
PING 10.0.0.3 (10.0.0.3) 56(84) bytes of data.
64 bytes from 10.0.0.3: icmp_req=1 ttl=64 time=0.984 ms
64 bytes from 10.0.0.3: icmp_req=2 ttl=64 time=0.110 ms
64 bytes from 10.0.0.3: icmp_req=3 ttl=64 time=0.098 ms
```

- You can also verify the configurations by executing the following REST API. It shows all configuration in VTN Manager.

```
curl --user "admin":"admin" -H "Content-type: application/json" -X GET http://
↪localhost:8181/restconf/operational/vtn:vtns/vtn/vtn1
```

```
{
  "vtn": [
  {
    "name": "vtn1",
      "vtenant-config": {
```

(continues on next page)

(continued from previous page)

```
        "hard-timeout": 0,
        "idle-timeout": 300,
        "description": "creating vtn"
      },
      "vbridge": [
      {
        "name": "vbr1",
        "vbridge-config": {
          "age-interval": 600,
          "description": "creating vBridge1"
        },
        "bridge-status": {
          "state": "UP",
          "path-faults": 0
        },
        "vinterface": [
        {
          "name": "if1",
          "vinterface-status": {
            "mapped-port": "openflow:2:1",
            "state": "UP",
            "entity-state": "UP"
          },
          "port-map-config": {
            "vlan-id": 0,
            "node": "openflow:2",
            "port-name": "s2-eth1"
          },
          "vinterface-config": {
            "description": "Creating if1 interface",
            "enabled": true
          },
          "vinterface-input-filter": {
            "vtn-flow-filter": [
            {
              "index": 1,
              "condition": "cond_1",
              "vtn-flow-action": [
              {
                "order": 1,
                "vtn-set-inet-src-action": {
                  "ipv4-address": "10.0.0.1/32"
                }
              },
              {
                "order": 2,
                "vtn-set-inet-dst-action": {
                  "ipv4-address": "10.0.0.3/32"
                }
              }
              ],
                "vtn-pass-filter": {}
            },
            {
              "index": 10,
              "condition": "cond_1",
              "vtn-drop-filter": {}
```

(continues on next page)

(continued from previous page)

```
            }
            ]
          }
        },
        {
          "name": "if2",
          "vinterface-status": {
            "mapped-port": "openflow:3:1",
            "state": "UP",
            "entity-state": "UP"
          },
          "port-map-config": {
            "vlan-id": 0,
            "node": "openflow:3",
            "port-name": "s3-eth1"
          },
          "vinterface-config": {
            "description": "Creating if2 interface",
            "enabled": true
          }
        }
        ]
      }
    ]
  }
  ]
}
```

Cleaning Up

- To clean up both VTN and flowcondition.
- You can delete the virtual tenant vtn1 by executing the remove-vtn RPC.

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn:remove-vtn -d '{"input":{"tenant-name":"vtn1
↪"}}'
```

- You can delete the flowcondition cond_1 by executing the remove-flow-condition RPC.

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn-flow-condition:remove-flow-condition -d '{
↪"input":{"name":"cond_1"}}'
```

How to use VTN to change the path of the packet flow

Overview

- This page explains how to create specific VTN Pathmap using VTN Manager. This page targets Boron release, so the procedure described here does not work in other releases.

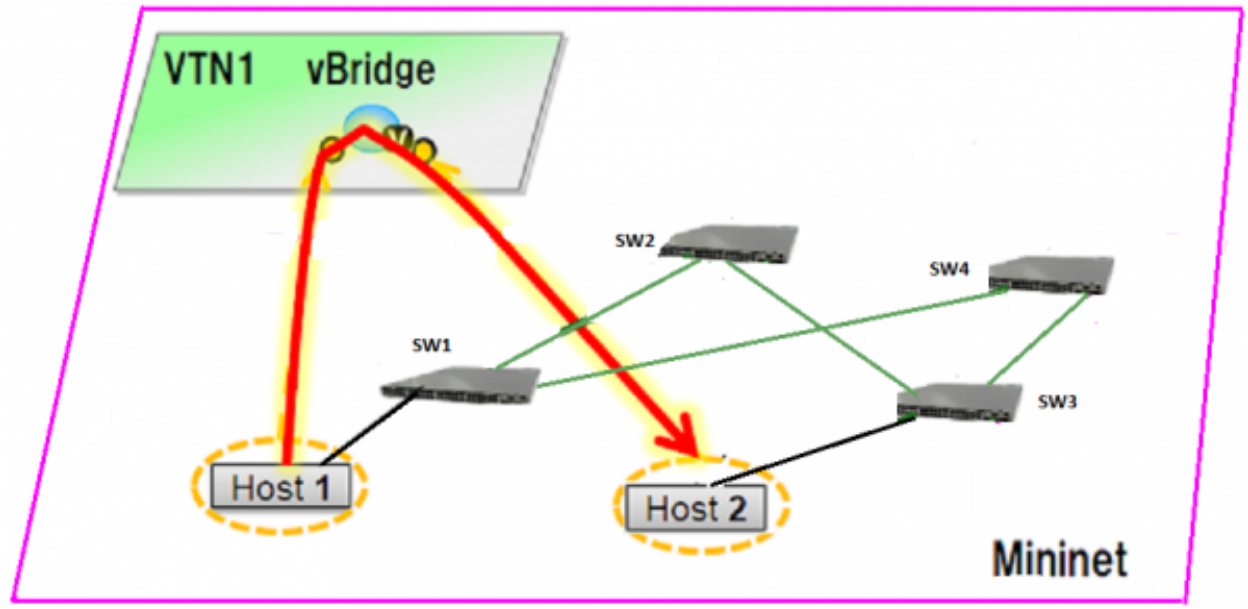

Fig. 139: Pathmap

Requirement

- Save the mininet script given below as pathmap_test.py and run the mininet script in the mininet environment where Mininet is installed.
- Create topology using the below mininet script:

```
from mininet.topo import Topo
class MyTopo( Topo ):
   "Simple topology example."
   def __init__( self ):
       "Create custom topo."
       # Initialize topology
       Topo.__init__( self )
       # Add hosts and switches
       leftHost = self.addHost( 'h1' )
       rightHost = self.addHost( 'h2' )
       leftSwitch = self.addSwitch( 's1' )
       middleSwitch = self.addSwitch( 's2' )
       middleSwitch2 = self.addSwitch( 's4' )
       rightSwitch = self.addSwitch( 's3' )
       # Add links
       self.addLink( leftHost, leftSwitch )
       self.addLink( leftSwitch, middleSwitch )
       self.addLink( leftSwitch, middleSwitch2 )
       self.addLink( middleSwitch, rightSwitch )
       self.addLink( middleSwitch2, rightSwitch )
       self.addLink( rightSwitch, rightHost )
topos = { 'mytopo': ( lambda: MyTopo() ) }
```

- After creating new file with the above script start the mininet as below,

```
sudo mn --controller=remote,ip=10.106.138.124 --custom pathmap_test.py --topo mytopo
```

Note: Replace "10.106.138.124" with the IP address of OpenDaylight controller based on your environment.

```
mininet> net
h1 h1-eth0:s1-eth1
h2 h2-eth0:s3-eth3
s1 lo:  s1-eth1:h1-eth0 s1-eth2:s2-eth1 s1-eth3:s4-eth1
s2 lo:  s2-eth1:s1-eth2 s2-eth2:s3-eth1
s3 lo:  s3-eth1:s2-eth2 s3-eth2:s4-eth2 s3-eth3:h2-eth0
s4 lo:  s4-eth1:s1-eth3 s4-eth2:s3-eth2
c0
```

- Generate traffic by pinging between host h1 and host h2 before creating the portmaps respectively.

```
mininet> h1 ping h2
PING 10.0.0.2 (10.0.0.2) 56(84) bytes of data.
From 10.0.0.1 icmp_seq=1 Destination Host Unreachable
From 10.0.0.1 icmp_seq=2 Destination Host Unreachable
From 10.0.0.1 icmp_seq=3 Destination Host Unreachable
From 10.0.0.1 icmp_seq=4 Destination Host Unreachable
```

Configuration

- To change the path of the packet flow, execute REST API provided by VTN Manager as follows. It uses curl command to call the REST API.
- Create a virtual tenant named vtn1 by executing the update-vtn RPC.

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn:update-vtn -d '{"input":{"tenant-name":"vtn1
↪"}}'
```

- Create a virtual bridge named vbr1 in the tenant vtn1 by executing the update-vbridge RPC.

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn-vbridge:update-vbridge -d '{"input":{"tenant-
↪name":"vtn1","bridge-name":"vbr1"}}'
```

- Create two interfaces into the virtual bridge by executing the update-vinterface RPC.

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn-vinterface:update-vinterface -d '{"input":{
↪"tenant-name":"vtn1","bridge-name":"vbr1","interface-name":"if1"}}'
```

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn-vinterface:update-vinterface -d '{"input":{
↪"tenant-name":"vtn1","bridge-name":"vbr1","interface-name":"if2"}}'
```

- Configure two mappings on the interfaces by executing the set-port-map RPC.
  - The interface if1 of the virtual bridge will be mapped to the port "s2-eth1" of the switch "openflow:1" of the Mininet.

        * The h1 is connected to the port "s1-eth1".
    - The interface if2 of the virtual bridge will be mapped to the port "s3-eth1" of the switch "openflow:3" of the Mininet.
        * The h3 is connected to the port "s3-eth3".

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn-port-map:set-port-map -d '{"input":{"tenant-
↪name":"vtn1", "bridge-name":"vbr1", "interface-name":"if1", "node":"openflow:1",
↪"port-name":"s1-eth1"}}'
```

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn-port-map:set-port-map -d '{"input":{"tenant-
↪name":"vtn1", "bridge-name":"vbr1", "interface-name":"if2", "node":"openflow:3",
↪"port-name":"s3-eth3"}}'
```

- Genarate traffic by pinging between host h1 and host h2 after creating the portmaps respectively.

```
mininet> h1 ping h2
PING 10.0.0.2 (10.0.0.2) 56(84) bytes of data.
64 bytes from 10.0.0.2: icmp_seq=1 ttl=64 time=0.861 ms
64 bytes from 10.0.0.2: icmp_seq=2 ttl=64 time=0.101 ms
64 bytes from 10.0.0.2: icmp_seq=3 ttl=64 time=0.101 ms
```

- Get the Dataflows information by executing the get-data-flow RPC.

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn-flow:get-data-flow -d '{"input":{"tenant-name
↪":"vtn1","mode":"DETAIL","node":"openflow:1","data-flow-port":{"port-id":1,"port-
↪name":"s1-eth1"}}}'
```

- Create flowcondition named cond_1 by executing the set-flow-condition RPC.
    - For option source and destination-network, get inet address of host h1 or host h2 from mininet

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn-flow-condition:set-flow-condition -d '{"input
↪":{"operation":"SET","present":"false","name":"cond_1", "vtn-flow-match":[{"vtn-
↪ether-match":{},"vtn-inet-match":{"source-network":"10.0.0.1/32","protocol":1,
↪"destination-network":"10.0.0.2/32"},"index":"1"}]}}'
```

- Create pathmap with flowcondition cond_1 by executing the set-path-map RPC.

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn-path-map:set-path-map -d '{"input":{"tenant-
↪name":"vtn1","path-map-list":[{"condition":"cond_1","policy":"1","index": "1","idle-
↪timeout":"300","hard-timeout":"0"}]}}'
```

- Create pathpolicy by executing the set-path-policy RPC.

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn-path-policy:set-path-policy -d '{"input":{
↪"operation":"SET","id": "1","default-cost": "10000","vtn-path-cost": [{"port-desc":
↪"openflow:1,3,s1-eth3","cost":"1000"},{"port-desc":"openflow:4,2,s4-eth2","cost":
↪"1000"},{"port-desc":"openflow:3,3,s3-eth3","cost":"100000"}]}}'
```

Verification

- Before applying Path policy get node information by executing get dataflow command.


```
"data-flow-info": [
{
  "physical-route": [
  {
    "physical-ingress-port": {
      "port-name": "s3-eth3",
        "port-id": "3"
    },
      "physical-egress-port": {
        "port-name": "s3-eth1",
        "port-id": "1"
      },
      "node": "openflow:3",
      "order": 0
  },
  {
    "physical-ingress-port": {
      "port-name": "s2-eth2",
      "port-id": "2"
    },
    "physical-egress-port": {
      "port-name": "s2-eth1",
      "port-id": "1"
    },
    "node": "openflow:2",
    "order": 1
  },
  {
    "physical-ingress-port": {
      "port-name": "s1-eth2",
      "port-id": "2"
    },
    "physical-egress-port": {
      "port-name": "s1-eth1",
      "port-id": "1"
    },
    "node": "openflow:1",
    "order": 2
  }
  ],
    "data-egress-node": {
      "interface-name": "if1",
      "bridge-name": "vbr1",
      "tenant-name": "vtn1"
    },
    "data-egress-port": {
      "node": "openflow:1",
      "port-name": "s1-eth1",
      "port-id": "1"
    },
    "data-ingress-node": {
      "interface-name": "if2",
      "bridge-name": "vbr1",
```


(continues on next page)

(continued from previous page)

```
      "tenant-name": "vtn1"
    },
    "data-ingress-port": {
      "node": "openflow:3",
      "port-name": "s3-eth3",
      "port-id": "3"
    },
    "flow-id": 32
  },
}
```

- After applying Path policy get node information by executing get dataflow command.

```
"data-flow-info": [
{
  "physical-route": [
  {
    "physical-ingress-port": {
      "port-name": "s1-eth1",
        "port-id": "1"
    },
      "physical-egress-port": {
        "port-name": "s1-eth3",
        "port-id": "3"
      },
      "node": "openflow:1",
      "order": 0
  },
  {
    "physical-ingress-port": {
      "port-name": "s4-eth1",
      "port-id": "1"
    },
    "physical-egress-port": {
      "port-name": "s4-eth2",
      "port-id": "2"
    },
    "node": "openflow:4",
    "order": 1
  },
  {
    "physical-ingress-port": {
      "port-name": "s3-eth2",
      "port-id": "2"
    },
    "physical-egress-port": {
      "port-name": "s3-eth3",
      "port-id": "3"
    },
    "node": "openflow:3",
    "order": 2
  }
  ],
    "data-egress-node": {
      "interface-name": "if2",
      "bridge-name": "vbr1",
      "tenant-name": "vtn1"
```

(continues on next page)

(continued from previous page)

```
    },
    "data-egress-port": {
      "node": "openflow:3",
      "port-name": "s3-eth3",
      "port-id": "3"
    },
    "data-ingress-node": {
      "interface-name": "if1",
      "bridge-name": "vbr1",
      "tenant-name": "vtn1"
    },
    "data-ingress-port": {
      "node": "openflow:1",
      "port-name": "s1-eth1",
      "port-id": "1"
    },
}
```

Cleaning Up

- To clean up both VTN and flowcondition.
- You can delete the virtual tenant vtn1 by executing the remove-vtn RPC.

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn:remove-vtn -d '{"input":{"tenant-name":"vtn1
↪"}}'
```

- You can delete the flowcondition cond_1 by executing the remove-flow-condition RPC.

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn-flow-condition:remove-flow-condition -d '{
↪"input":{"name":"cond_1"}}'
```

VTN Coordinator Usage Examples

How to configure L2 Network with Single Controller

Overview

This example provides the procedure to demonstrate configuration of VTN Coordinator with L2 network using VTN Virtualization(single controller). Here is the Example for vBridge Interface Mapping with Single Controller using mininet. mininet details and set-up can be referred at below URL: https://wiki.opendaylight.org/view/OpenDaylight_Controller:Installation#Using_Mininet

Requirements

- Configure mininet and create a topology:

```
mininet@mininet-vm:~$ sudo mn --controller=remote,ip=<controller-ip> --topo tree,2
```

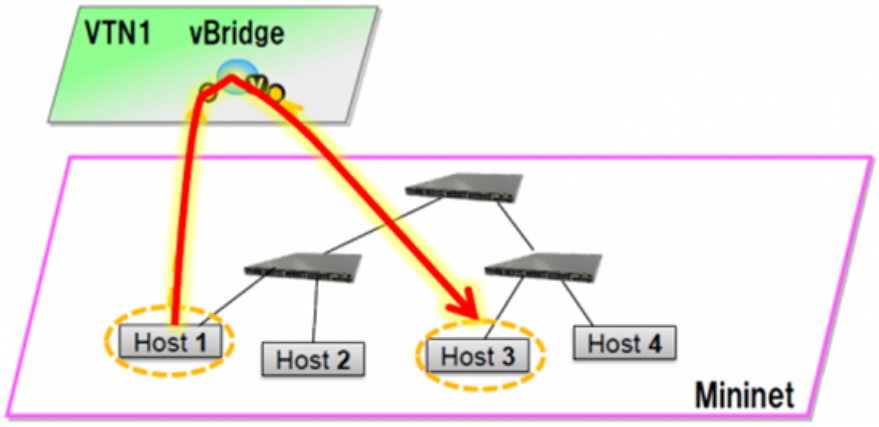

Fig. 140: EXAMPLE DEMONSTRATING SINGLE CONTROLLER

- mininet> net

```
s1 lo:  s1-eth1:h1-eth0 s1-eth2:s2-eth1
s2 lo:  s2-eth1:s1-eth2 s2-eth2:h2-eth0
h1 h1-eth0:s1-eth1
h2 h2-eth0:s2-eth2
```

Configuration

- Create a Controller named controllerone and mention its ip-address in the below create-controller command.

```
curl --user admin:adminpass -H 'content-type: application/json' -X POST -d '{
↪"controller": {"controller_id": "controllerone", "ipaddr":"10.0.0.2", "type": "odc",
↪ "version": "1.0", "auditstatus":"enable"}}' http://127.0.0.1:8083/vtn-webapi/
↪controllers.json
```

- Create a VTN named vtn1 by executing the create-vtn command

```
curl --user admin:adminpass -H 'content-type: application/json' -X POST -d '{"vtn" : {
↪"vtn_name":"vtn1","description":"test VTN" }}' http://127.0.0.1:8083/vtn-webapi/
↪vtns.json
```

- Create a vBridge named vBridge1 in the vtn1 by executing the create-vbr command.

```
curl --user admin:adminpass -H 'content-type: application/json' -X POST -d '{"vbridge
↪" : {"vbr_name":"vBridge1","controller_id":"controllerone","domain_id":"(DEFAULT)" }
↪}' http://127.0.0.1:8083/vtn-webapi/vtns/vtn1/vbridges.json
```

- Create two Interfaces named if1 and if2 into the vBridge1

```
curl --user admin:adminpass -H 'content-type: application/json' -X POST -d '{
↪"interface": {"if_name": "if1","description": "if_desc1"}}' http://127.0.0.1:8083/
↪vtn-webapi/vtns/vtn1/vbridges/vBridge1/interfaces.json
```

```
curl --user admin:adminpass -H 'content-type: application/json' -X POST -d '{
↪"interface": {"if_name": "if2","description": "if_desc2"}}' http://127.0.0.1:8083/
↪vtn-webapi/vtns/vtn1/vbridges/vBridge1/interfaces.json
```

- Get the list of logical ports configured

```
Curl --user admin:adminpass -H 'content-type: application/json' -X GET http://127.0.0.
↪1:8083/vtn-webapi/controllers/controllerone/domains/\(DEFAULT\)/logical_ports.json
```

- Configure two mappings on each of the interfaces by executing the below command.

The interface if1 of the virtual bridge will be mapped to the port "s2-eth1" of the switch "openflow:2" of the Mininet. The h1 is connected to the port "s2-eth1".

The interface if2 of the virtual bridge will be mapped to the port "s3-eth1" of the switch "openflow:3" of the Mininet. The h3 is connected to the port "s3-eth1".

```
curl --user admin:adminpass -H 'content-type: application/json' -X PUT -d '{"portmap":
↪{"logical_port_id": "PP-OF:00:00:00:00:00:00:03-s3-eth1"}}' http://127.0.0.
↪1:8083/vtn-webapi/vtns/vtn1/vbridges/vBridge1/interfaces/if1/portmap.json
curl --user admin:adminpass -H 'content-type: application/json' -X PUT -d '{"portmap":
↪{"logical_port_id": "PP-OF:00:00:00:00:00:00:02-s2-eth1"}}' http://127.0.0.
↪1:8083/vtn-webapi/vtns/vtn1/vbridges/vBridge1/interfaces/if2/portmap.json
```

Verification

Please verify whether the Host1 and Host3 are pinging.

- Send packets from Host1 to Host3

```
mininet> h1 ping h3
PING 10.0.0.3 (10.0.0.3) 56(84) bytes of data.
64 bytes from 10.0.0.3: icmp_req=1 ttl=64 time=0.780 ms
64 bytes from 10.0.0.3: icmp_req=2 ttl=64 time=0.079 ms
```

How to configure L2 Network with Multiple Controllers

- This example provides the procedure to demonstrate configuration of VTN Coordinator with L2 network using VTN Virtualization Here is the Example for vBridge Interface Mapping with Multi-controller using mininet.

Requirements

- Configure multiple controllers using the mininet script given below: https://wiki.opendaylight.org/view/OpenDaylight_Virtual_Tenant_Network_%28VTN%29:Scripts:Mininet#Network_with_multiple_switches_and_OpenFlow_controllers

Configuration

- Create a VTN named vtn3 by executing the create-vtn command

```
curl --user admin:adminpass -H 'content-type: application/json'  -X POST -d '{"vtn" :
↪{"vtn_name":"vtn3"}}' http://127.0.0.1:8083/vtn-webapi/vtns.json
```

- Create two Controllers named odc1 and odc2 with its ip-address in the below create-controller command.

```
curl --user admin:adminpass -H 'content-type: application/json'  -X POST -d '{
↪"controller": {"controller_id": "odc1", "ipaddr":"10.100.9.52", "type": "odc",
↪"version": "1.0", "auditstatus":"enable"}}' http://127.0.0.1:8083/vtn-webapi/
↪controllers.json
```

(continues on next page)

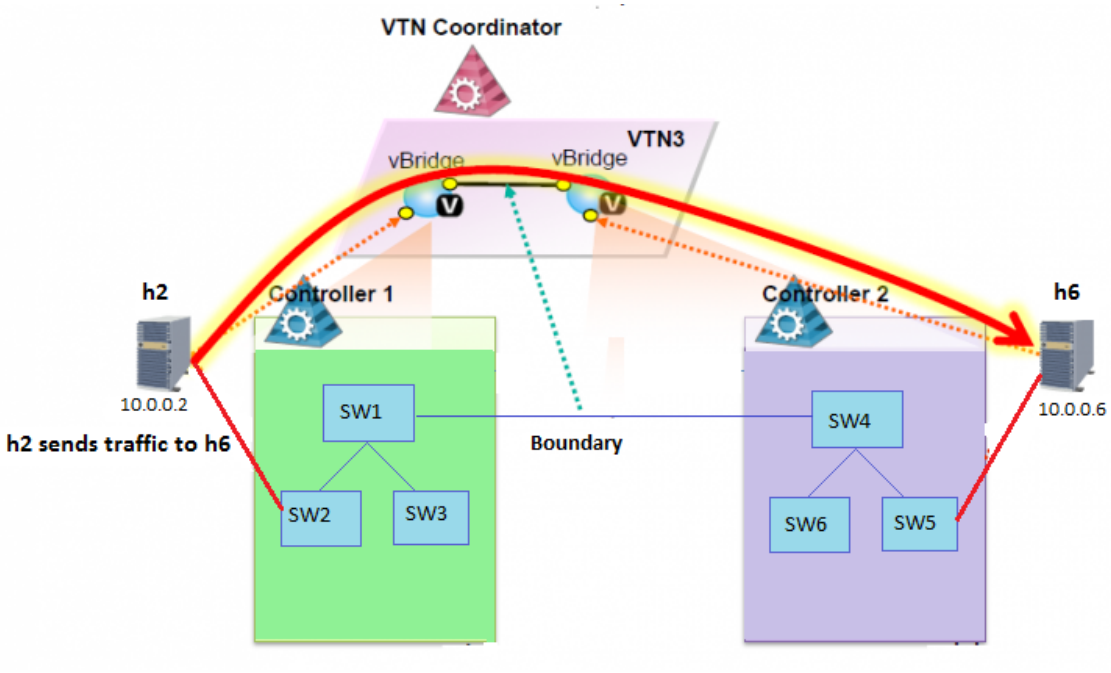

Fig. 141: EXAMPLE DEMONSTRATING MULTIPLE CONTROLLERS

(continued from previous page)

```
curl --user admin:adminpass -H 'content-type: application/json'  -X POST -d '{
↪"controller": {"controller_id": "odc2", "ipaddr":"10.100.9.61", "type": "odc",
↪"version": "1.0", "auditstatus":"enable"}}' http://127.0.0.1:8083/vtn-webapi/
↪controllers.json
```

- Create two vBridges in the VTN like, vBridge1 in Controller1 and vBridge2 in Controller2

```
curl --user admin:adminpass -H 'content-type: application/json' -X POST -d '{"vbridge
↪" : {"vbr_name":"vbr1","controller_id":"odc1","domain_id":"(DEFAULT)" }}' http://
↪127.0.0.1:8083/vtn-webapi/vtns/vtn3/vbridges.json
```

```
curl --user admin:adminpass -H 'content-type: application/json'  -X POST -d '{"vbridge
↪" : {"vbr_name":"vbr2","controller_id":"odc2","domain_id":"(DEFAULT)" }}' http://
↪127.0.0.1:8083/vtn-webapi/vtns/vtn3/vbridges.json
```

- Create two Interfaces if1, if2 for the two vBridges vbr1 and vbr2.

```
curl --user admin:adminpass -H 'content-type: application/json'  -X POST -d '{
↪"interface": {"if_name": "if1"}}' http://127.0.0.1:8083/vtn-webapi/vtns/vtn3/
↪vbridges/vbr1/interfaces.json
```

```
curl --user admin:adminpass -H 'content-type: application/json'  -X POST -d '{
↪"interface": {"if_name": "if2"}}' http://127.0.0.1:8083/vtn-webapi/vtns/vtn3/
↪vbridges/vbr1/interfaces.json
```

```
curl --user admin:adminpass -H 'content-type: application/json'  -X POST -d '{
↪"interface": {"if_name": "if1"}}' http://127.0.0.1:8083/vtn-webapi/vtns/vtn3/
↪vbridges/vbr2/interfaces.json
```

(continues on next page)

(continued from previous page)

```
curl --user admin:adminpass -H 'content-type: application/json'  -X POST -d '{
↪"interface": {"if_name": "if2"}}' http://127.0.0.1:8083/vtn-webapi/vtns/vtn3/
↪vbridges/vbr2/interfaces.json
```

- Get the list of logical ports configured

```
curl --user admin:adminpass -H 'content-type: application/json' -X GET http://127.0.0.
↪1:8083/vtn-webapi/controllers/odc1/domains/\(DEFAULT\)/logical_ports/detail.json
```

- Create boundary and vLink for the two controllers

```
curl --user admin:adminpass -H 'content-type: application/json'   -X POST -d '{
↪"boundary": {"boundary_id": "b1", "link": {"controller1_id": "odc1", "domain1_id":
↪"(DEFAULT)", "logical_port1_id": "PP-OF:00:00:00:00:00:00:01-s1-eth3",
↪"controller2_id": "odc2", "domain2_id": "(DEFAULT)", "logical_port2_id": "PP-
↪OF:00:00:00:00:00:00:04-s4-eth3"}}}' http://127.0.0.1:8083/vtn-webapi/boundaries.
↪json
```

```
curl --user admin:adminpass -H 'content-type: application/json'  -X POST -d '{"vlink
↪": {"vlk_name": "vlink1" , "vnode1_name": "vbr1", "if1_name":"if2", "vnode2_name":
↪"vbr2", "if2_name": "if2", "boundary_map": {"boundary_id":"b1","vlan_id": "50"}}}'␣
↪http://127.0.0.1:8083/vtn-webapi/vtns/vtn3/vlinks.json
```

- Configure two mappings on each of the interfaces by executing the below command.

The interface if1 of the vbr1 will be mapped to the port "s2-eth2" of the switch "openflow:2" of the Mininet. The h2 is connected to the port "s2-eth2".

The interface if2 of the vbr2 will be mapped to the port "s5-eth2" of the switch "openflow:5" of the Mininet. The h6 is connected to the port "s5-eth2".

```
curl --user admin:adminpass -H 'content-type: application/json' -X PUT -d '{"portmap":
↪{"logical_port_id": "PP-OF:00:00:00:00:00:00:02-s2-eth2"}}' http://127.0.0.
↪1:8083/vtn-webapi/vtns/vtn3/vbridges/vbr1/interfaces/if1/portmap.json
```

```
curl --user admin:adminpass -H 'content-type: application/json'  -X PUT -d '{"portmap
↪":{"logical_port_id": "PP-OF:00:00:00:00:00:00:05-s5-eth2"}}' http://127.0.0.
↪1:8083/vtn-webapi/vtns/vtn3/vbridges/vbr2/interfaces/if1/portmap.json
```

Verification

Please verify whether Host h2 and Host h6 are pinging.

- Send packets from h2 to h6

```
mininet> h2 ping h6
```

```
PING 10.0.0.6 (10.0.0.3) 56(84) bytes of data.
64 bytes from 10.0.0.6: icmp_req=1 ttl=64 time=0.780 ms
64 bytes from 10.0.0.6: icmp_req=2 ttl=64 time=0.079 ms
```

How To Test Vlan-Map In Mininet Environment

Overview

This example explains how to test vlan-map in a multi host scenario.

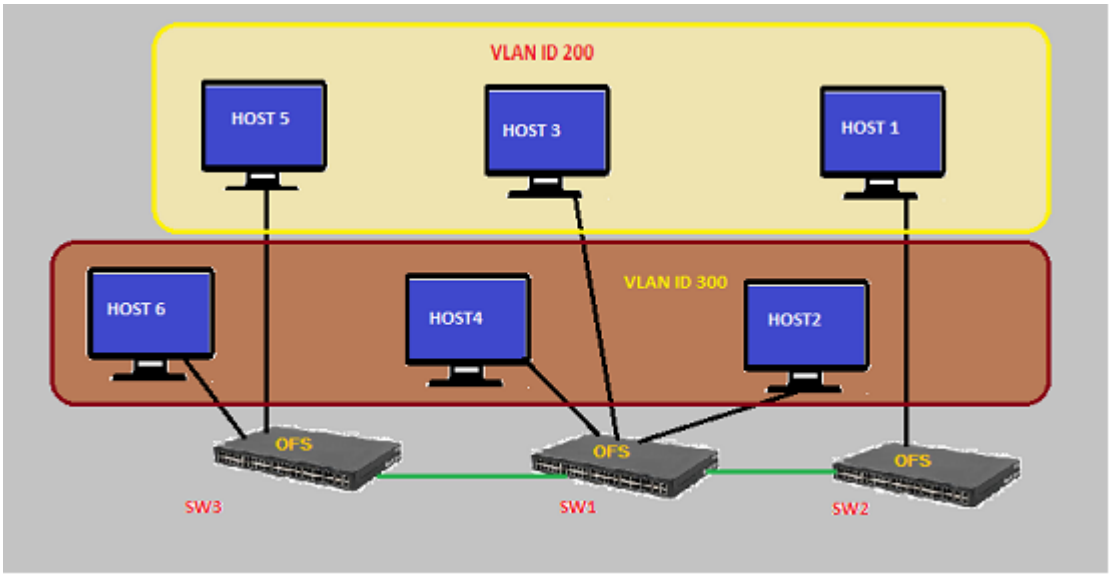

Fig. 142: Example that demonstrates vlanmap testing in Mininet Environment

Requirements

- Save the mininet script given below as vlan_vtn_test.py and run the mininet script in the mininet environment where Mininet is installed.

Mininet Script

https://wiki.opendaylight.org/view/OpenDaylight_Virtual_Tenant_Network_(VTN):Scripts:Mininet#Network_with_hosts_in_different_vlan

- Run the mininet script

```
sudo mn --controller=remote,ip=192.168.64.13 --custom vlan_vtn_test.py --topo mytopo
```

Configuration

Please follow the below steps to test a vlan map using mininet:

- Create a Controller named controllerone and mention its ip-address in the below create-controller command.

```
curl --user admin:adminpass -H 'content-type: application/json'  -X POST -d '{
↪"controller": {"controller_id": "controllerone", "ipaddr":"10.0.0.2", "type": "odc",
↪ "version": "1.0", "auditstatus":"enable"}}' http://127.0.0.1:8083/vtn-webapi/
↪controllers
```

- Create a VTN named vtn1 by executing the create-vtn command

```
curl -X POST -H 'content-type: application/json' -H 'username: admin' -H 'password:␣
↪adminpass' -d '{"vtn" : {"vtn_name":"vtn1","description":"test VTN" }}' http://127.
↪0.0.1:8083/vtn-webapi/vtns.json
```

- Create a vBridge named vBridge1 in the vtn1 by executing the create-vbr command.

```
curl -X POST -H 'content-type: application/json' -H 'username: admin' -H 'password:␣
↪adminpass' -d '{"vbridge" : {"vbr_name":"vBridge1","controller_id":"controllerone",
↪"domain_id":"(DEFAULT)" }}' http://127.0.0.1:8083/vtn-webapi/vtns/vtn1/vbridges.json
```

- Create a vlan map with vlanid 200 for vBridge vBridge1

```
curl -X POST -H 'content-type: application/json' -H 'username: admin' -H 'password:␣
↪adminpass' -d '{"vlanmap" : {"vlan_id": 200 }}' http://127.0.0.1:8083/vtn-webapi/
↪vtns/vtn1/vbridges/vBridge1/vlanmaps.json
```

- Create a vBridge named vBridge2 in the vtn1 by executing the create-vbr command.

```
curl -X POST -H 'content-type: application/json' -H 'username: admin' -H 'password:␣
↪adminpass' -d '{"vbridge" : {"vbr_name":"vBridge2","controller_id":"controllerone",
↪"domain_id":"(DEFAULT)" }}' http://127.0.0.1:8083/vtn-webapi/vtns/vtn1/vbridges.json
```

- Create a vlan map with vlanid 300 for vBridge vBridge2

```
curl -X POST -H 'content-type: application/json' -H 'username: admin' -H 'password:␣
↪adminpass' -d '{"vlanmap" : {"vlan_id": 300 }}' http://127.0.0.1:8083/vtn-webapi/
↪vtns/vtn1/vbridges/vBridge2/vlanmaps.json
```

Verification

Ping all in mininet environment to view the host reachability.

```
mininet> pingall
Ping: testing ping reachability
h1 -> X h3 X h5 X
h2 -> X X h4 X h6
h3 -> h1 X X h5 X
h4 -> X h2 X X h6
h5 -> h1 X h3 X X
h6 -> X h2 X h4 X
```

How To View Specific VTN Station Information.

This example demonstrates on how to view a specific VTN Station information.

Requirement

- Configure mininet and create a topology:

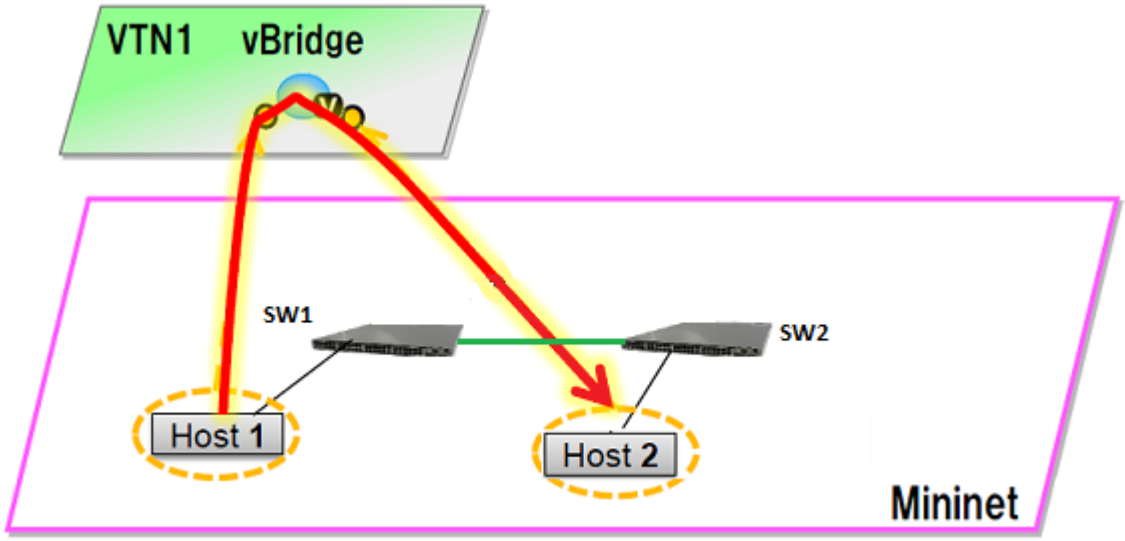

Fig. 143: EXAMPLE DEMONSTRATING VTN STATIONS

```
 $ sudo mn --custom /home/mininet/mininet/custom/topo-2sw-2host.py --
↪controller=remote,ip=10.100.9.61 --topo mytopo
mininet> net

 s1 lo:  s1-eth1:h1-eth0 s1-eth2:s2-eth1
 s2 lo:  s2-eth1:s1-eth2 s2-eth2:h2-eth0
 h1 h1-eth0:s1-eth1
 h2 h2-eth0:s2-eth2
```

- Generate traffic by pinging between hosts h1 and h2 after configuring the portmaps respectively

```
mininet> h1 ping h2
PING 10.0.0.2 (10.0.0.2) 56(84) bytes of data.
64 bytes from 10.0.0.2: icmp_req=1 ttl=64 time=16.7 ms
64 bytes from 10.0.0.2: icmp_req=2 ttl=64 time=13.2 ms
```

Configuration

- Create a Controller named controllerone and mention its ip-address in the below create-controller command

```
curl --user admin:adminpass -H 'content-type: application/json' -X POST -d '{
↪"controller": {"controller_id": "controllerone", "ipaddr":"10.100.9.61", "type":
↪"odc", "version": "1.0", "auditstatus":"enable"}}' http://127.0.0.1:8083/vtn-webapi/
↪controllers.json
```

- Create a VTN named vtn1 by executing the create-vtn command

```
curl --user admin:adminpass -H 'content-type: application/json' -X POST -d '{"vtn" : {
↪"vtn_name":"vtn1","description":"test VTN" }}' http://127.0.0.1:8083/vtn-webapi/
↪vtns.json
```

- Create a vBridge named vBridge1 in the vtn1 by executing the create-vbr command.

```
curl --user admin:adminpass -H 'content-type: application/json' -X POST -d '{"vbridge
↪" : {"vbr_name":"vBridge1","controller_id":"controllerone","domain_id":"(DEFAULT)" }
↪}' http://127.0.0.1:8083/vtn-webapi/vtns/vtn1/vbridges.json
```

- Create two Interfaces named if1 and if2 into the vBridge1

```
curl --user admin:adminpass -H 'content-type: application/json' -X POST -d '{
↪"interface": {"if_name": "if1","description": "if_desc1"}}' http://127.0.0.1:8083/
↪vtn-webapi/vtns/vtn1/vbridges/vBridge1/interfaces.json
curl -v --user admin:adminpass -H 'content-type: application/json' -X POST -d '{
↪"interface": {"if_name": "if2","description": "if_desc2"}}' http://127.0.0.1:8083/
↪vtn-webapi/vtns/vtn1/vbridges/vBridge1/interfaces.json
```

- Configure two mappings on each of the interfaces by executing the below command.

The interface if1 of the virtual bridge will be mapped to the port "s1-eth1" of the switch "openflow:1" of the Mininet. The h1 is connected to the port "s1-eth1".

The interface if2 of the virtual bridge will be mapped to the port "s1-eth2" of the switch "openflow:1" of the Mininet. The h2 is connected to the port "s1-eth2".

```
curl --user admin:adminpass -H 'content-type: application/json' -X PUT -d '{"portmap":
↪{"logical_port_id": "PP-OF:00:00:00:00:00:00:01-s1-eth1"}}' http://127.0.0.
↪1:8083/vtn-webapi/vtns/vtn1/vbridges/vBridge1/interfaces/if1/portmap.json
curl -v --user admin:adminpass -H 'content-type: application/json' -X PUT -d '{
↪"portmap":{"logical_port_id": "PP-OF:00:00:00:00:00:00:02-s2-eth2"}}' http://17.
↪0.0.1:8083/vtn-webapi/vtns/vtn1/vbridges/vBridge1/interfaces/if2/portmap.json
```

- Get the VTN stations information

```
curl -X GET -H 'content-type: application/json' -H 'username: admin' -H 'password:␣
↪adminpass' "http://127.0.0.1:8083/vtn-webapi/vtnstations?controller_
↪id=controllerone&vtn_name=vtn1"
```

Verification

```
curl -X GET -H 'content-type: application/json' -H 'username: admin' -H 'password:␣
↪adminpass' "http://127.0.0.1:8083/vtn-webapi/vtnstations?controller_
↪id=controllerone&vtn_name=vtn1"
{
   "vtnstations": [
       {
           "domain_id": "(DEFAULT)",
           "interface": {},
           "ipaddrs": [
               "10.0.0.2"
           ],
           "macaddr": "b2c3.06b8.2dac",
           "no_vlan_id": "true",
           "port_name": "s2-eth2",
           "station_id": "178195618445172",
           "switch_id": "00:00:00:00:00:00:00:02",
           "vnode_name": "vBridge1",
           "vnode_type": "vbridge",
           "vtn_name": "vtn1"
       },
```

(continues on next page)

(continued from previous page)


```
        {
            "domain_id": "(DEFAULT)",
            "interface": {},
            "ipaddrs": [
                "10.0.0.1"
            ],
            "macaddr": "ce82.1b08.90cf",
            "no_vlan_id": "true",
            "port_name": "s1-eth1",
            "station_id": "206130278144207",
            "switch_id": "00:00:00:00:00:00:00:01",
            "vnode_name": "vBridge1",
            "vnode_type": "vbridge",
            "vtn_name": "vtn1"
        }
    ]
}
```


How To View Dataflows in VTN

This example demonstrates on how to view a specific VTN Dataflow information.

Configuration

The same Configuration as Vlan Mapping Example(https://wiki.opendaylight.org/view/VTN:Coordinator:Beryllium:HowTos:How_To_test_vlanmap_using_mininet)

Verification

Get the VTN Dataflows information

```
curl -X GET -H 'content-type: application/json' --user 'admin:adminpass' "http://127.
↪0.0.1:8083/vtn-webapi/dataflows?controller_id=controllerone&srcmacaddr=924c.e4a3.
↪a743&vlan_id=300&switch_id=openflow:2&port_name=s2-eth1"
```


```
{
    "dataflows": [
        {
            "controller_dataflows": [
                {
                    "controller_id": "controllerone",
                    "controller_type": "odc",
                    "egress_domain_id": "(DEFAULT)",
                    "egress_port_name": "s3-eth3",
                    "egress_station_id": "3",
                    "egress_switch_id": "00:00:00:00:00:00:00:03",
                    "flow_id": "29",
                    "ingress_domain_id": "(DEFAULT)",
                    "ingress_port_name": "s2-eth2",
                    "ingress_station_id": "2",
                    "ingress_switch_id": "00:00:00:00:00:00:00:02",
```


(continues on next page)

(continued from previous page)


```
                    "match": {
                        "macdstaddr": [
                            "4298.0959.0e0b"
                        ],
                        "macsrcaddr": [
                            "924c.e4a3.a743"
                        ],
                        "vlan_id": [
                            "300"
                        ]
                    },
                    "pathinfos": [
                        {
                            "in_port_name": "s2-eth2",
                            "out_port_name": "s2-eth1",
                            "switch_id": "00:00:00:00:00:00:00:02"
                        },
                        {
                            "in_port_name": "s1-eth2",
                            "out_port_name": "s1-eth3",
                            "switch_id": "00:00:00:00:00:00:00:01"
                        },
                        {
                            "in_port_name": "s3-eth1",
                            "out_port_name": "s3-eth3",
                            "switch_id": "00:00:00:00:00:00:00:03"
                        }
                    ]
                }
            ],
            "reason": "success"
        }
    ]
}
```


How To Configure Flow Filters Using VTN

Overview

The flow-filter function discards, permits, or redirects packets of the traffic within a VTN, according to specified flow conditions The table below lists the actions to be applied when a packet matches the condition:

Action	Function
Pass	Permits the packet to pass. As options, packet transfer priority (set priority) and DSCP change (se t ip-dscp) is specified.
Drop	Discards the packet.
Redirect	Redirects the packet to a desired virtual interface. As an option, it is possible to change the MAC address when the packet is transferred.

Following steps explain flow-filter function:

- When a packet is transferred to an interface within a virtual network, the flow-filter function evaluates whether the transferred packet matches the condition specified in the flow-list.

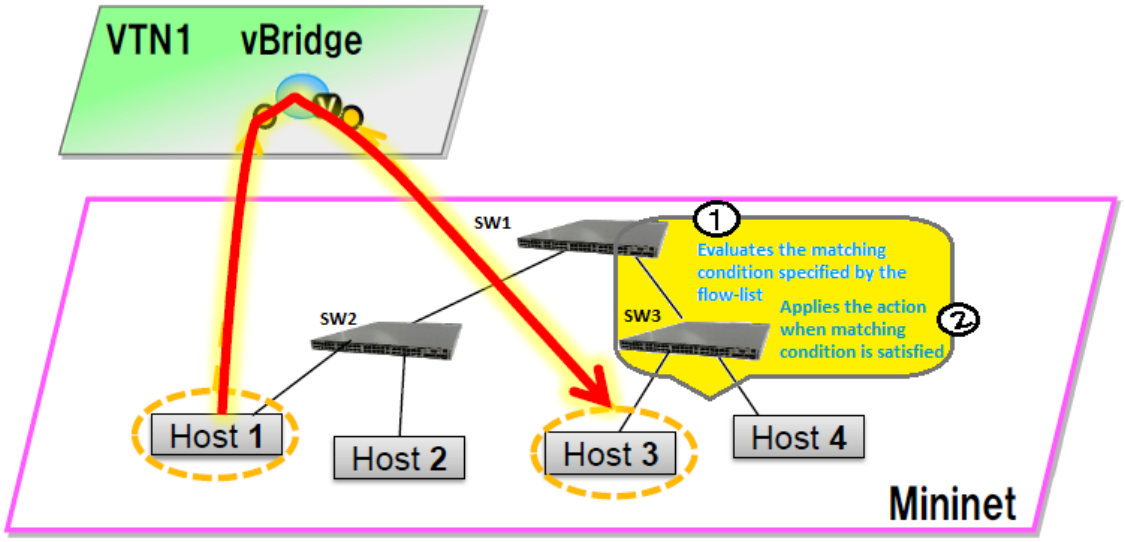

Fig. 144: Flow Filter

- If the packet matches the condition, the flow-filter applies the flow-list matching action specified in the flow-filter.

Requirements

To apply the packet filter, configure the following:

- Create a flow-list and flow-listentry.
- Specify where to apply the flow-filter, for example VTN, vBridge, or interface of vBridge.

Configure mininet and create a topology:

```
$  mininet@mininet-vm:~$ sudo mn --controller=remote,ip=<controller-ip> --topo tree
```

Please generate the following topology

```
$  mininet@mininet-vm:~$ sudo mn --controller=remote,ip=<controller-ip> --topo tree,2
mininet> net
c0
s1 lo:  s1-eth1:s2-eth3 s1-eth2:s3-eth3
s2 lo:  s2-eth1:h1-eth0 s2-eth2:h2-eth0 s2-eth3:s1-eth1
s3 lo:  s3-eth1:h3-eth0 s3-eth2:h4-eth0 s3-eth3:s1-eth2
h1 h1-eth0:s2-eth1
h2 h2-eth0:s2-eth2
h3 h3-eth0:s3-eth1
h4 h4-eth0:s3-eth2
```

Configuration

- Create a Controller named controller1 and mention its ip-address in the below create-controller command.

```
curl --user admin:adminpass -H 'content-type: application/json' -X POST -d '{
↪"controller": {"controller_id": "controller1", "ipaddr":"10.100.9.61", "type": "odc
↪", "version": "1.0", "auditstatus":"enable"}}' http://127.0.0.1:8083/vtn-webapi/
↪controllers
```

- Create a VTN named vtn_one by executing the create-vtn command

```
curl --user admin:adminpass -H 'content-type: application/json' -X POST -d '{"vtn" : {
↪"vtn_name":"vtn_one","description":"test VTN" }}' http://127.0.0.1:8083/vtn-webapi/
↪vtns.json
```

- Create a vBridge named vbr_two in the vtn1 by executing the create-vbr command.

```
curl --user admin:adminpass -H 'content-type: application/json' -X POST -d '{"vbridge
↪" : {"vbr_name":"vbr_one^C"controller_id":"controller1","domain_id":"(DEFAULT)" }}'␣
↪http://127.0.0.1:8083/vtn-webapi/vtns/vtn_one/vbridges.json
curl -v --user admin:adminpass -H 'content-type: application/json' -X POST -d '{
↪"vbridge" :
{"vbr_name":"vbr_two","controller_id":"controller1","domain_id":"(DEFAULT)" }}' http:/
↪/127.0.0.1:8083/vtn-webapi/vtns/vtn_one/vbridges.json
```

- Create two Interfaces named if1 and if2 into the vbr_two

```
curl --user admin:adminpass -H 'content-type: application/json' -X POST -d '{
↪"interface": {"if_name": "if1","description": "if_desc1"}}' http://127.0.0.1:8083/
↪vtn-webapi/vtns/vtn_one/vbridges/vbr_two/interfaces.json
curl -v --user admin:adminpass -H 'content-type: application/json' -X POST -d '{
↪"interface": {"if_name": "if1","description": "if_desc1"}}' http://127.0.0.1:8083/
↪vtn-webapi/vtns/vtn_one/vbridges/vbr_two/interfaces.json
```

- Get the list of logical ports configured

```
curl --user admin:adminpass -H 'content-type: application/json' -X GET  http://127.0.
↪0.1:8083/vtn-webapi/controllers/controllerone/domains/\(DEFAULT\)/logical_ports.json
```

- Configure two mappings on each of the interfaces by executing the below command.

The interface if1 of the virtual bridge will be mapped to the port "s2-eth1" of the switch "openflow:2" of the Mininet. The h1 is connected to the port "s2-eth1".

The interface if2 of the virtual bridge will be mapped to the port "s3-eth1" of the switch "openflow:3" of the Mininet. The h3 is connected to the port "s3-eth1".

```
curl --user admin:adminpass -H 'content-type: application/json' -X PUT -d '{"portmap":
↪{"logical_port_id": "PP-OF:00:00:00:00:00:00:00:03-s3-eth1"}}' http://127.0.0.
↪1:8083/vtn-webapi/vtns/vtn_one/vbridges/vbr_two/interfaces/if1/portmap.json
curl -v --user admin:adminpass -H 'content-type: application/json' -X PUT -d '{
↪"portmap":{"logical_port_id": "PP-OF:00:00:00:00:00:00:00:02-s2-eth1"}}' http://127.
↪0.0.1:8083/vtn-webapi/vtns/vtn_one/vbridges/vbr_two/interfaces/if2/portmap.json
```

- Create Flowlist

```
curl --user admin:adminpass -H 'content-type: application/json' -X POST -d '{"flowlist
↪": {"fl_name": "flowlist1", "ip_version":"IP"}}' http://127.0.0.1:8083/vtn-webapi/
↪flowlists.json
```

- Create Flowlistentry

```
curl --user admin:adminpass -H 'content-type: application/json' -X POST -d '{
↪"flowlistentry": {"seqnum": "233","macethertype": "0x8000","ipdstaddr": "10.0.0.3",
↪"ipdstaddrprefix": "2","ipsrcaddr": "10.0.0.2","ipsrcaddrprefix": "2","ipproto": "17
↪","ipdscp": "55","icmptypenum":"232","icmpcodenum": "232"}}' http://127.0.0.1:8083/
↪vtn-webapi/flowlists/flowlist1/flowlistentries.json
```

- Create vBridge Interface Flowfilter

```
curl --user admin:adminpass -X POST -H 'content-type: application/json' -d '{
↪"flowfilter" : {"ff_type": "in"}}' http://127.0.0.1:8083/vtn-webapi/vtns/vtn_one/
↪vbridges/vbr_two/interfaces/if1/flowfilters.json
```

Flow filter demonstration with DROP action-type

```
curl --user admin:adminpass -X POST -H 'content-type: application/json' -d '{
↪"flowfilterentry": {"seqnum": "233", "fl_name": "flowlist1", "action_type":"drop",
↪"priority":"3", "dscp":"55" }}' http://127.0.0.1:8083/vtn-webapi/vtns/vtn_one/
↪vbridges/vbr_two/interfaces/if1/flowfilters/in/flowfilterentries.json
```

Verification

As we have applied the action type “drop” , ping should fail.

```
mininet> h1 ping h3
PING 10.0.0.2 (10.0.0.2) 56(84) bytes of data.
From 10.0.0.1 icmp_seq=1 Destination Host Unreachable
From 10.0.0.1 icmp_seq=2 Destination Host Unreachable
```

Flow filter demonstration with PASS action-type

```
curl --user admin:adminpass -X PUT -H 'content-type: application/json' -d '{
↪"flowfilterentry": {"seqnum": "233", "fl_name": "flowlist1", "action_type":"pass",
↪"priority":"3", "dscp":"55" }}' http://127.0.0.1:8083/vtn-webapi/vtns/vtn_one/
↪vbridges/vbr_two/interfaces/if1/flowfilters/in/flowfilterentries/233.json
```

Verification

```
mininet> h1 ping h3
PING 10.0.0.3 (10.0.0.3) 56(84) bytes of data.
64 bytes from 10.0.0.3: icmp_req=1 ttl=64 time=0.984 ms
64 bytes from 10.0.0.3: icmp_req=2 ttl=64 time=0.110 ms
64 bytes from 10.0.0.3: icmp_req=3 ttl=64 time=0.098 ms
```

How To Use VTN To Make Packets Take Different Paths

This example demonstrates on how to create a specific VTN Path Map information.

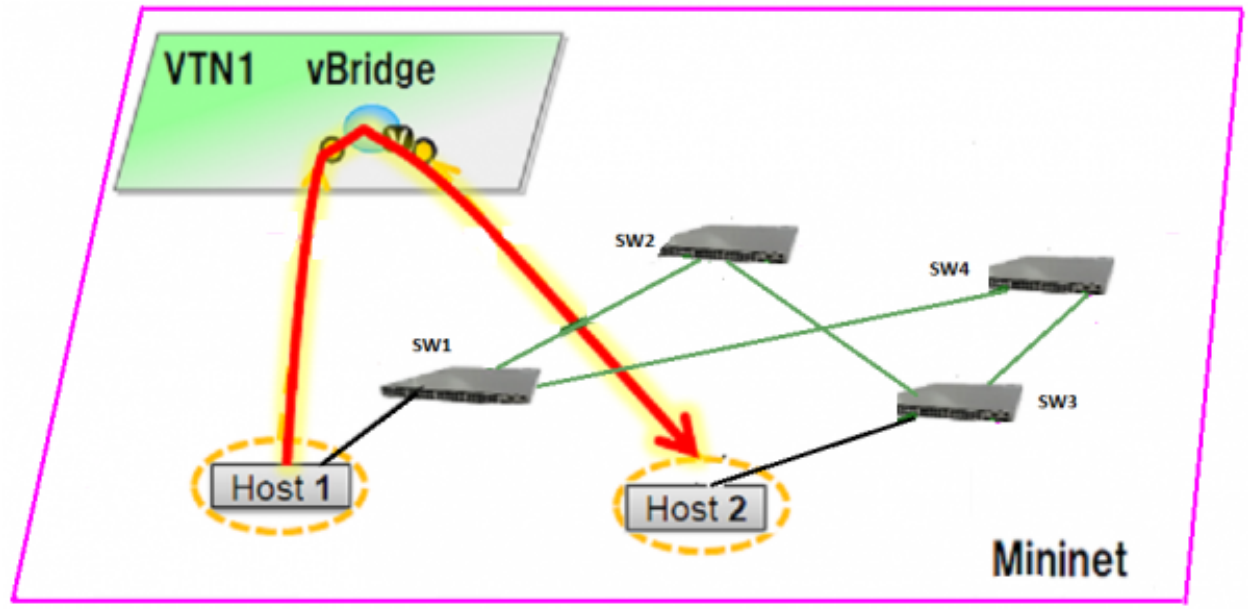

Fig. 145: PathMap

Requirement

- Save the mininet script given below as pathmap_test.py and run the mininet script in the mininet environment where Mininet is installed.
- Create topology using the below mininet script:

```
from mininet.topo import Topo
class MyTopo( Topo ):
   "Simple topology example."
   def __init__( self ):
        "Create custom topo."
        # Initialize topology
        Topo.__init__( self )
        # Add hosts and switches
        leftHost = self.addHost( 'h1' )
        rightHost = self.addHost( 'h2' )
        leftSwitch = self.addSwitch( 's1' )
        middleSwitch = self.addSwitch( 's2' )
        middleSwitch2 = self.addSwitch( 's4' )
        rightSwitch = self.addSwitch( 's3' )
        # Add links
        self.addLink( leftHost, leftSwitch )
        self.addLink( leftSwitch, middleSwitch )
        self.addLink( leftSwitch, middleSwitch2 )
        self.addLink( middleSwitch, rightSwitch )
        self.addLink( middleSwitch2, rightSwitch )
        self.addLink( rightSwitch, rightHost )
topos = { 'mytopo': ( lambda: MyTopo() ) }
```

```
mininet> net
c0
```

(continues on next page)

(continued from previous page)

```
s1 lo:  s1-eth1:h1-eth0 s1-eth2:s2-eth1 s1-eth3:s4-eth1
s2 lo:  s2-eth1:s1-eth2 s2-eth2:s3-eth1
s3 lo:  s3-eth1:s2-eth2 s3-eth2:s4-eth2 s3-eth3:h2-eth0
s4 lo:  s4-eth1:s1-eth3 s4-eth2:s3-eth2
h1 h1-eth0:s1-eth1
h2 h2-eth0:s3-eth3
```

- Generate traffic by pinging between hosts h1 and h2 before creating the portmaps respectively

```
mininet> h1 ping h2
PING 10.0.0.2 (10.0.0.2) 56(84) bytes of data.
From 10.0.0.1 icmp_seq=1 Destination Host Unreachable
From 10.0.0.1 icmp_seq=2 Destination Host Unreachable
From 10.0.0.1 icmp_seq=3 Destination Host Unreachable
From 10.0.0.1 icmp_seq=4 Destination Host Unreachable
```

Configuration

- Create a Controller named controller1 and mention its ip-address in the below create-controller command.

```
curl --user admin:adminpass -H 'content-type: application/json'  -X POST -d '{
↪"controller": {"controller_id": "odc", "ipaddr":"10.100.9.42", "type": "odc",
↪"version": "1.0", "auditstatus":"enable"}}' http://127.0.0.1:8083/vtn-webapi/
↪controllers.json
```

- Create a VTN named vtn1 by executing the create-vtn command

```
curl --user admin:adminpass -H 'content-type: application/json'  -X POST -d '{"vtn" :
↪{"vtn_name":"vtn1","description":"test VTN" }}' http://127.0.0.1:8083/vtn-webapi/
↪vtns.json
```

- Create a vBridge named vBridge1 in the vtn1 by executing the create-vbr command.

```
curl --user admin:adminpass -H 'content-type: application/json'  -X POST -d '{"vbridge
↪" : {"vbr_name":"vBridge1","controller_id":"odc","domain_id":"(DEFAULT)" }}' http://
↪127.0.0.1:8083/vtn-webapi/vtns/vtn1/vbridges.json
```

- Create two Interfaces named if1 and if2 into the vBridge1

```
curl --user admin:adminpass -H 'content-type: application/json'  -X POST -d '{
↪"interface": {"if_name": "if1","description": "if_desc1"}}' http://127.0.0.1:8083/
↪vtn-webapi/vtns/vtn1/vbridges/vBridge1/interfaces.json
curl --user admin:adminpass -H 'content-type: application/json'  -X POST -d '{
↪"interface": {"if_name": "if2","description": "if_desc2"}}' http://127.0.0.1:8083/
↪vtn-webapi/vtns/vtn1/vbridges/vBridge1/interfaces.json
```

- Configure two mappings on each of the interfaces by executing the below command.

The interface if1 of the virtual bridge will be mapped to the port "s1-eth1" of the switch "openflow:1" of the Mininet. The h1 is connected to the port "s1-eth1".

The interface if2 of the virtual bridge will be mapped to the port "s3-eth3" of the switch "openflow:3" of the Mininet. The h2 is connected to the port "s3-eth3".

```
curl --user admin:adminpass -H 'content-type: application/json'  -X PUT -d '{"portmap
↪":{"logical_port_id": "PP-OF:00:00:00:00:00:00:00:01-s1-eth1"}}' http://127.0.0.
↪1:8083/vtn-webapi/vtns/vtn1/vbridges/vBridge1/interfaces/if1/portmap.json
curl --user admin:adminpass -H 'content-type: application/json'  -X PUT -d '{"portmap
↪":{"logical_port_id": "PP-OF:00:00:00:00:00:00:03-s3-eth3"}}' http://127.0.0.
↪1:8083/vtn-webapi/vtns/vtn1/vbridges/vBridge1/interfaces/if2/portmap.json
```

- Generate traffic by pinging between hosts h1 and h2 after creating the portmaps respectively

```
mininet> h1 ping h2
PING 10.0.0.2 (10.0.0.2) 56(84) bytes of data.
64 bytes from 10.0.0.2: icmp_req=1 ttl=64 time=36.4 ms
64 bytes from 10.0.0.2: icmp_req=2 ttl=64 time=0.880 ms
64 bytes from 10.0.0.2: icmp_req=3 ttl=64 time=0.073 ms
64 bytes from 10.0.0.2: icmp_req=4 ttl=64 time=0.081 ms
```

- Get the VTN Dataflows information

```
curl -X GET -H 'content-type: application/json' --user 'admin:adminpass' "http://127.
↪0.0.1:8083/vtn-webapi/dataflows?&switch_id=00:00:00:00:00:00:00:01&port_name=s1-
↪eth1&controller_id=odc&srcmacaddr=de3d.7dec.e4d2&no_vlan_id=true"
```

- Create a Flowcondition in the VTN

(The flowconditions, pathmap and pathpolicy commands have to be executed in the controller).

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn-flow-condition:set-flow-condition -d '{"input
↪":{"operation":"SET","present":"false","name":"cond_1", "vtn-flow-match":[{"vtn-
↪ether-match":{},"vtn-inet-match":{"source-network":"10.0.0.1/32","protocol":1,
↪"destination-network":"10.0.0.2/32"},"index":"1"}]}}'
```

- Create a Pathmap in the VTN

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn-path-map:set-path-map -d '{"input":{"tenant-
↪name":"vtn1","path-map-list":[{"condition":"cond_1","policy":"1","index": "1","idle-
↪timeout":"300","hard-timeout":"0"}]}}'
```

- Get the Path policy information

```
curl --user "admin":"admin" -H "Content-type: application/json" -X POST http://
↪localhost:8181/restconf/operations/vtn-path-policy:set-path-policy -d '{"input":{
↪"operation":"SET","id": "1","default-cost": "10000","vtn-path-cost": [{"port-desc":
↪"openflow:1,3,s1-eth3","cost":"1000"},{"port-desc":"openflow:4,2,s4-eth2","cost":
↪"100000"},{"port-desc":"openflow:3,3,s3-eth3","cost":"10000"}]}}'
```

Verification

- Before applying Path policy information in the VTN

```
{
        "pathinfos": [
            {
              "in_port_name": "s1-eth1",
              "out_port_name": "s1-eth3",
```

(continues on next page)

(continued from previous page)

```
              "switch_id": "openflow:1"
            },
            {
              "in_port_name": "s4-eth1",
              "out_port_name": "s4-eth2",
              "switch_id": "openflow:4"
            },
            {
               "in_port_name": "s3-eth2",
               "out_port_name": "s3-eth3",
               "switch_id": "openflow:3"
            }
                    ]
}
```

- After applying Path policy information in the VTN

```
{
    "pathinfos": [
            {
              "in_port_name": "s1-eth1",
              "out_port_name": "s1-eth2",
              "switch_id": "openflow:1"
            },
            {
              "in_port_name": "s2-eth1",
              "out_port_name": "s2-eth2",
              "switch_id": "openflow:2"
            },
            {
               "in_port_name": "s3-eth1",
               "out_port_name": "s3-eth3",
               "switch_id": "openflow:3"
            }
                    ]
}
```

VTN Coordinator(Troubleshooting HowTo)

Overview

This page demonstrates Installation troubleshooting steps of VTN Coordinator. OpenDaylight VTN provides multi-tenant virtual network functions on OpenDaylight controllers. OpenDaylight VTN consists of two parts:

- VTN Coordinator.
- VTN Manager.

VTN Coordinator orchestrates multiple VTN Managers running in OpenDaylight Controllers, and provides VTN Applications with VTN API. VTN Manager is OSGi bundles running in OpenDaylight Controller. Current VTN Manager supports only OpenFlow switches. It handles PACKET_IN messages, sends PACKET_OUT messages, manages host information, and installs flow entries into OpenFlow switches to provide VTN Coordinator with virtual network functions. The requirements for installing these two are different.Therefore, we recommend that you install VTN Manager and VTN Coordinator in different machines.

List of installation Troubleshooting How to's

- https://wiki.opendaylight.org/view/OpenDaylight_Virtual_Tenant_Network_(VTN):Installation:VTN_Coordinator

After executing db_setup, you have encountered the error "Failed to setup database"?

The error could be due to the below reasons

- Access Restriction

The user who owns /usr/local/vtn/ directory and installs VTN Coordinator, can only start db_setup. Example :

```
The directory should appear as below (assuming the user as "vtn"):
# ls -l /usr/local/
  drwxr-xr-x. 12 vtn  vtn  4096 Mar 14 21:53 vtn
If the user doesnot own /usr/local/vtn/ then, please run the below command (assuming␣
↪the username as vtn),
            chown -R vtn:vtn /usr/local/vtn
```

- Postgres not Present

```
1. In case of Fedora/CentOS/RHEL, please check if /usr/pgsql/<version> directory is␣
↪present and also ensure the commands initdb, createdb,pg_ctl,psql are working. If,␣
↪not please re-install postgres packages
2. In case of Ubuntu, check if /usr/lib/postgres/<version> directory is present and␣
↪check for the commands as in the previous step.
```

- Not enough space to create tables

```
Please check df -k and ensure enough free space is available.
```

- If the above steps do not solve the problem, please refer to the log file for the exact problem

```
/usr/local/vtn/var/dbm/unc_setup_db.log for the exact error.
```

- list of VTN Coordinator processes
- Run the below command ensure the Coordinator daemons are running.

```
   Command:     /usr/local/vtn/bin/unc_dmctl status
   Name              Type           IPC Channel       PID
-----------       -----------      --------------     ------
    drvodcd         DRIVER           drvodcd          15972
    lgcnwd         LOGICAL           lgcnwd           16010
    phynwd         PHYSICAL          phynwd           15996
```

- Issue the curl command to fetch version and ensure the process is able to respond.

How to debug a startup failure?.

The following activities take place in order during startup

- Database server is started after setting virtual memory to required value,Any database startup errors will be reflected in any of the below logs.

```
/usr/local/vtn/var/dbm/unc_db_script.log.
/usr/local/vtn/var/db/pg_log/postgresql-*.log (the pattern will have the date)
```

- uncd daemon is kicked off, The daemon in turn kicks off the rest of the daemons.

```
Any  uncd startup failures will be reflected in /usr/local/vtn/var/uncd/uncd_start.
↪err.
```

After setting up the apache tomcat server, what are the aspects that should be checked.

Please check if catalina is running..

```
The command ps -ef | grep catalina | grep -v grep should list a catalina process
```

If you encounter an erroneous situation where the REST API is always failing..

```
Please ensure the firewall settings for port:8181 (Beryllium release) or port:8083␣
↪(Post Beryllium release) and enable the same.
```

How to debug a REST API returning a failure message?.

Please check the /usr/share/java/apache-tomcat-7.0.39/logs/core/core.log for failure details.

REST API for VTN configuration fails, how to debug?.

The default log level for all daemons is "INFO", to debug the situation TRACE or DEBUG logs may be needed. To increase the log level for individual daemons, please use the commands suggested below

```
/usr/local/vtn/bin/lgcnw_control loglevel trace -- upll daemon log
 /usr/local/vtn/bin/phynw_control loglevel trace -- uppl daemon log
 /usr/local/vtn/bin/unc_control loglevel trace -- uncd daemon log
 /usr/local/vtn/bin/drvodc_control loglevel trace -- Driver daemon log
```

After setting the log levels, the operation can be repeated and the log files can be referred for debugging.

Problems while Installing PostgreSQL due to openssl.

Errors may occur when trying to install postgreSQL rpms. Recently PostgreSQL has upgraded all their binaries to use the latest openssl versions with fix for http://en.wikipedia.org/wiki/Heartbleed Please upgrade the openssl package to the latest version and re-install. For RHEL 6.1/6.4 : If you have subscription, Please use the same and update the rpms. The details are available in the following link https://access.redhat.com/site/solutions/781793 ACCESS-REDHAT

```
rpm -Uvh http://mirrors.kernel.org/centos/6/os/x86_64/Packages/openssl-1.0.1e-15.el6.
↪x86_64.rpm
rpm -ivh http://mirrors.kernel.org/centos/6/os/x86_64/Packages/openssl-devel-1.0.1e-
↪15.el6.x86_64.rpm
```

For other linux platforms, Please do yum update, the public respositroes will have the latest openssl, please install the same.

Support for Microsoft SCVMM 2012 R2 with ODL VTN

Introduction

System Center Virtual Machine Manager (SCVMM) is Microsoft's virtual machine support center for window's based emulations. SCVMM is a management solution for the virtualized data center. You can use it to configure and manage your virtualization host, networking, and storage resources in order to create and deploy virtual machines and services to private clouds that you have created.

The VSEM Provider is a plug-in to bridge between SCVMM and OpenDaylight.

Microsoft Hyper-V is a server virtualization developed by Microsoft, which provides virtualization services through hypervisor-based emulations.

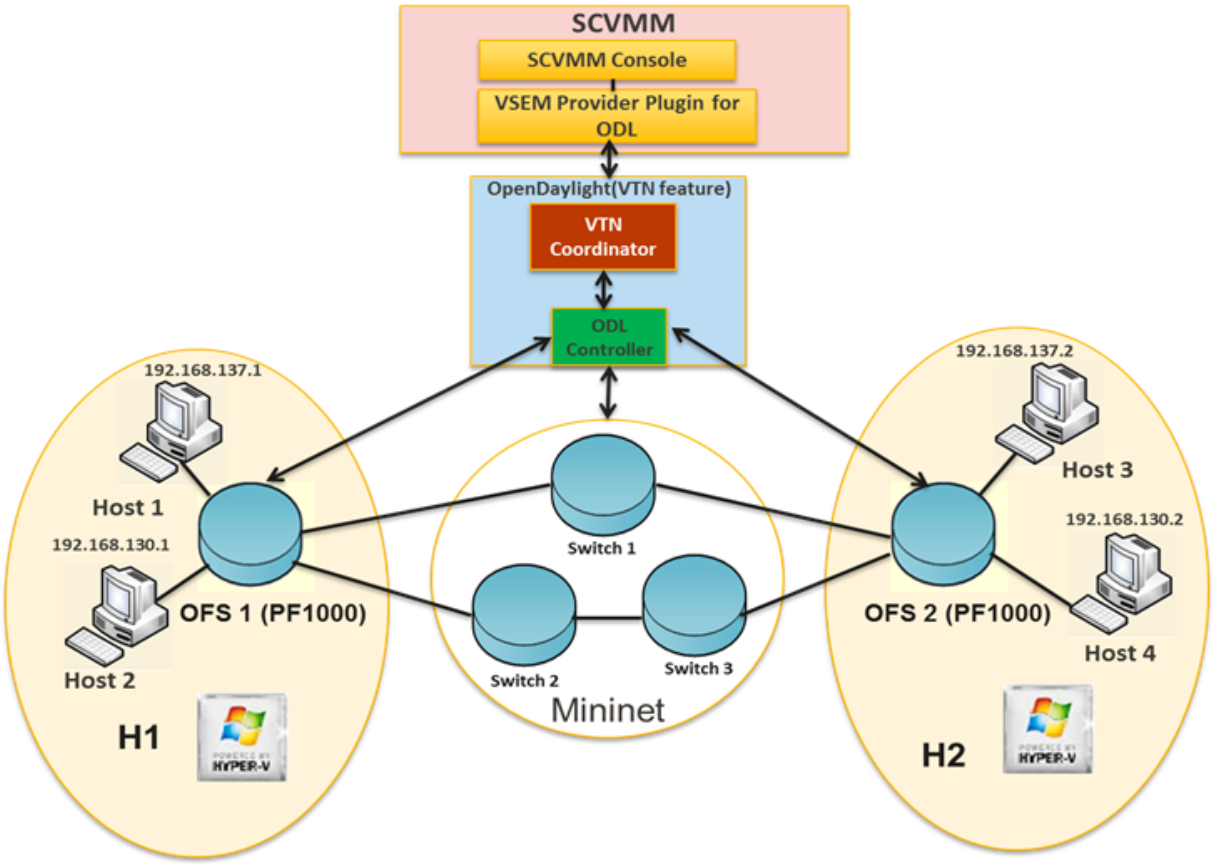

Fig. 146: Set-Up Diagram

The topology used in this set-up is:

- A SCVMM with VSEM Provider installed and a running VTN Coordinator and OpenDaylight with VTN Feature installed.
- PF1000 virtual switch extension has been installed in the two Hyper-V servers as it implements the OpenFlow capability in Hyper-V.
- Three OpenFlow switches simulated using mininet and connected to Hyper-V.
- Four VM's hosted using SCVMM.

It is implemented as two major components:

- SCVMM
- OpenDaylight (VTN Feature)
- VTN Coordinator

VTN Coordinator

OpenDaylight VTN as Network Service provider for SCVMM where VSEM provider is added in the Network Service which will handle all requests from SCVMM and communicate with the VTN Coordinator. It is used to manage the network virtualization provided by OpenDaylight.

Installing HTTPS in VTN Coordinator

- System Center Virtual Machine Manager (SCVMM) supports only https protocol.

Apache Portable Runtime (APR) Installation Steps

- Enter the command "yum install **apr**" in VTN Coordinator installed machine.
- In /usr/bin, create a soft link as "ln –s /usr/bin/apr-1-config /usr/bin/apr-config".
- Extract tomcat under "/usr/share/java" by using the below command "tar -xvf apache-tomcat-8.0.27.tar.gz –C /usr/share/java".

Note: Please go through the bleow link to download apache-tomcat-8.0.27.tar.gz file, https://archive.apache.org/dist/tomcat/tomcat-8/v8.0.27/bin/

- Please go to the directory "cd /usr/share/java/apache-tomcat-8.0.27/bin and unzip tomcat-native.gz using this command "tar -xvf tomcat-native.gz".
- Go to the directory "cd /usr/share/java/apache-tomcat-8.0.27/bin/tomcat-native-1.1.33-src/jni/native".
- Enter the command "./configure –with-os-type=bin –with-apr=/usr/bin/apr-config".
- Enter the command "make" and "make install".
- Apr libraries are successfully installed in "/usr/local/apr/lib".

Enable HTTP/HTTPS in VTN Coordinator

Enter the command "firewall-cmd –zone=public –add-port=8083/tcp –permanent" and "firewall-cmd –reload" to enable firewall settings in server.

Create a CA's private key and a self-signed certificate in server

- Execute the following command "openssl req -x509 -days 365 -extensions v3_ca -newkey rsa:2048 –out /etc/pki/CA/cacert.pem –keyout /etc/pki/CA/private/cakey.pem" in a single line.

Argument	Description
Country Name	Specify the country code. For example, JP
State or Province Name	Specify the state or province. For example, Tokyo
Locality Name	Locality Name For example, Chuo-Ku
Organization Name	Specify the company.
Organizational Unit Name	Specify the department, division, or the like.
Common Name	Specify the host name.
Email Address	Specify the e-mail address.

- Execute the following commands: "touch /etc/pki/CA/index.txt" and "echo 00 > /etc/pki/CA/serial" in server after setting your CA's private key.

Create a private key and a CSR for web server

- Execute the following command "openssl req -new -newkey rsa:2048 -out csr.pem –keyout /usr/local/vtn/tomcat/conf/key.pem" in a single line.
- Enter the PEM pass phrase: Same password you have given in CA's private key PEM pass phrase.

Argument	Description
Country Name	Specify the country code. For example, JP
State or Province Name	Specify the state or province. For example, Tokyo
Locality Name	Locality Name For example, Chuo-Ku
Organization Name	Specify the company.
Organizational Unit Name	Specify the department, division, or the like.
Common Name	Specify the host name.
Email Address	Specify the e-mail address.
A challenge password	Specify the challenge password.
An optional company name	Specify an optional company name.

Create a certificate for web server

- Execute the following command “openssl ca –in csr.pem –out /usr/local/vtn/tomcat/conf/cert.pem –days 365 –batch” in a single line.
- Enter pass phrase for /etc/pki/CA/private/cakey.pem: Same password you have given in CA’s private key PEM pass phrase.
- Open the tomcat file using “vim /usr/local/vtn/tomcat/bin/tomcat”.
- Include the line ” TOMCAT_PROPS=”$TOMCAT_PROPS -Djava.library.path="/usr/local/apr/lib"” ” in 131th line and save the file.

Edit server.xml file and restart the server

- Open the server.xml file using “vim /usr/local/vtn/tomcat/conf/server.xml” and add the below lines.

```
<Connector port="${vtn.port}" protocol="HTTP/1.1" SSLEnabled="true"
maxThreads="150" scheme="https" secure="true"
SSLCertificateFile="/usr/local/vtn/tomcat/conf/cert.pem"
SSLCertificateKeyFile="/usr/local/vtn/tomcat/conf/key.pem"
SSLPassword=same password you have given in CA's private key PEM pass phrase
connectionTimeout="20000" />
```

- Save the file and restart the server.
- To stop vtn use the following command.

```
/usr/local/vtn/bin/vtn_stop
```

- To start vtn use the following command.

```
/usr/local/vtn/bin/vtn_start
```

- Copy the created CA certificate from cacert.pem to cacert.crt by using the following command,

```
openssl x509 -in /etc/pki/CA/cacert.pem -out cacert.crt
```

Checking the HTTP and HTTPS connection from client

- You can check the HTTP connection by using the following command:

```
curl -X GET -H 'contenttype:application/json' -H 'username:admin' -H
↪'password:adminpass' http://<server IP address>:8083/vtn-webapi/api_version.json
```

- You can check the HTTPS connection by using the following command:

```
curl -X GET -H 'contenttype:application/json' -H 'username:admin' -H
↪'password:adminpass' https://<server IP address>:8083/vtn-webapi/api_version.
↪json --cacert /etc/pki/CA/cacert.pem
```

- The response should be like this for both HTTP and HTTPS:

```
{"api_version":{"version":"V1.4"}}
```

Prerequisites to create Network Service in SCVMM machine, Please follow the below steps

1. Please go through the below link to download VSEM Provider zip file, https://nexus.opendaylight.org/content/groups/public/org/opendaylight/vtn/application/vtnmanager-vsemprovider/1.2.0-Boron/vtnmanager-vsemprovider-1.2.0-Boron-bin.zip

2. Unzip the vtnmanager-vsemprovider-1.2.0-Boron-bin.zip file anywhere in your SCVMM machine.
3. Stop SCVMM service from **"service manager→tools→servers→select system center virtual machine manager"** and click stop.
4. Go to **"C:/Program Files"** in your SCVMM machine. Inside **"C:/Program Files"**, create a folder named as **"ODLProvider"**.
5. Inside **"C:/Program Files/ODLProvider"**, create a folder named as "Module" in your SCVMM machine.
6. Inside "C:/Program Files/ODLProvider/Module", Create two folders named as **"Odl.VSEMProvider"** and **"VSEMOdlUI"** in your SCVMM machine.
7. Copy the **"VSEMOdl.dll"** file from **"ODL_SCVMM_PROVIDER/ODL_VSEM_PROVIDER"** to **"C:/Program Files/ODLProvider/Module/Odl.VSEMProvider"** in your SCVMM machine.
8. Copy the **"VSEMOdlProvider.psd1"** file from **"application/vsemprovider/VSEMOdlProvider/VSEMOdlProvider.psd1"** to **"C:/Program Files/ODLProvider/Module/Odl.VSEMProvider"** in your SCVMM machine.
9. Copy the **"VSEMOdlUI.dll"** file from **"ODL_SCVMM_PROVIDER/ODL_VSEM_PROVIDER_UI"** to **"C:/Program Files/ODLProvider/Module/VSEMOdlUI"** in your SCVMM machine.
10. Copy the **"VSEMOdlUI.psd1"** file from **"application/vsemprovider/VSEMOdlUI"** to **"C:/Program Files/ODLProvider/Module/VSEMOdlUI"** in your SCVMM machine.
11. Copy the **"reg_entry.reg"** file from **"ODL_SCVMM_PROVIDER/Register_settings"** to your SCVMM desktop and double click the **"reg_entry.reg"** file to install registry entry in your SCVMM machine.
12. Download **"PF1000.msi"** from this link, https://www.pf-info.com/License/en/index.php?url=index/index_non_buyer and place into **"C:/Program Files/Switch Extension Drivers"** in your SCVMM machine.
13. Start SCVMM service from **"service manager→tools→servers→select system center virtual machine manager"** and click start.

System Center Virtual Machine Manager (SCVMM)

It supports two major features:

- Failover Clustering
- Live Migration

Failover Clustering

A single Hyper-V can host a number of virtual machines. If the host were to fail then all of the virtual machines that are running on it will also fail, thereby resulting in a major outage. Failover clustering treats individual virtual machines as clustered resources. If a host were to fail then clustered virtual machines are able to fail over to a different Hyper-V server where they can continue to run.

Live Migration

Live Migration is used to migrate the running virtual machines from one Hyper-V server to another Hyper-V server without any interruptions. Please go through the below video for more details,

- https://youtu.be/34YMOTzbNJM

SCVMM User Guide

- Please go through the below link for SCVMM user guide: https://wiki.opendaylight.org/images/c/ca/ODL_SCVMM_USER_GUIDE_final.pdf
- Please go through the below links for more details
  - OpenDaylight SCVMM VTN Integration: https://youtu.be/iRt4dxtiz94
  - OpenDaylight Congestion Control with SCVMM VTN: https://youtu.be/34YMOTzbNJM

YANG IDE User Guide

Overview

The YANG IDE project provides an Eclipse plugin that is used to create, view, and edit Yang model files. It currently supports version 1.0 of the Yang specification.

The YANG IDE project uses components from the OpenDaylight project for parsing and verifying Yang model files. The "yangtools" parser in OpenDaylight is generally used for generating Java code associated with Yang models. If you are just using the YANG IDE to view and edit Yang models, you do not need to know any more about this.

Although the YANG IDE plugin is used in Eclipse, it is not necessary to be familiar with the Java programming language to use it effectively.

The YANG IDE also uses the Maven build tool, but you do not have to be a Maven expert to use it, or even know that much about it. Very little configuration of Maven files will have to be done by you. In fact, about the only thing you will likely ever need to change can be done entirely in the Eclipse GUI forms, without even seeing the internal structure of the Maven POM file (Project Object Model).

The YANG IDE plugin provides features that are similar to other programming language plugins in the Eclipse ecosystem.

For instance, you will find support for the following:

- Immediate "as-you-type" display of syntactic and semantic errors
- Intelligent completion of language tokens, limited to only choices valid in the current scope and namespace
- Consistent (and customizable) color-coding of syntactic and semantic symbols
- Provides access to remote Yang models by specifying dependency on Maven artifact containing models (or by manual inclusion in project)
- One-click navigation to referenced symbols in external files
- Mouse hovers display descriptions of referenced components
- Tools for refactoring or renaming components respect namespaces
- Code templates can be entered for common conventions

Forthcoming sections of this manual will step through how to utilize these features.

Creating a Yang Project

After the plugin is installed, the next thing you have to do is create a Yang Project. This is done from the "File" menu, selecting "New", and navigating to the "Yang" section and selecting "YANG Project", and then clicking "Next" for more items to configure.

Some shortcuts for these steps are the following:

- Typically, the key sequence "Ctrl+n" (press "n" while holding down one of the "ctrl" keys) is bound to the "new" function
- In the "New" wizard dialog, the initial focus is in the filter field, where you can enter "yang" to limit the choices to only the functions provided by the YANG IDE plugin
- On the "New" wizard dialog, instead of clicking the "Next" button with your mouse, you can press "Alt+n" (you will see a hint for this with the "N" being underlined)

First Yang Project Wizard Page

After the "Next" button is pressed, it goes to the first wizard page that is specific to creating Yang projects. you will see a subtitle on this page of "YANG Tools Configuration". In almost all cases, you should be able to click "Next" again on this page to go to the next wizard page.

However, some information about the fields on this page would be helpful.

You will see the following labeled fields and sections:

Yang Files Root Directory

This defaults to "src/main/yang". Except when creating your first Yang file, you, you do not even have to know this, as Eclipse presents the same interface to view your Yang files no matter what you set this to.

Source Code Generators

If you do not know what this is, you do not need to know about it. The "yangtools" Yang parser from OpenDaylight uses a "code generator" component to generate specific kinds of Java classes from the Yang models. Again, if you do not need to work with the generated Java code, you do not need to change this.

Create Example YANG File

This is likely the only field you will ever have any reason to change. If this checkbox is set, when the YANG IDE creates the Yang project, it will create a sample "acme-system.yang" file which you can view and edit to demonstrate the features of the tool to yourself. If you do not need this file, then either delete it from the project or uncheck the checkbox to prevent its creation.

When done with the fields on this page, click the "Next" button to go to the next wizard page.

Second Yang Project Wizard Page

This page has a subtitle of "New Maven project". There are several fields on this page, but you will only ever have to see and change the setting of the first field, the "Create a simple project" checkbox. You should always set this ON to avoid the selection of a Maven archetype, which is something you do not need to do for creating a Yang project.

Click "Next" at the bottom of the page to move to the next wizard page.

Third Yang Project Wizard Page

This also has a subtitle of "New Maven project", but with different fields to set. You will likely only ever set the first two fields, and completely ignore everything else.

The first field is labeled "Group id" in the "Artifact" section. It really does not matter what you set this to, but it does have to be set to something. For consistency, you might set this to the name or nickname of your organization. Otherwise, there are no constraints on the value of this field.

The second field is labeled "Artifact id". The value of this field will be used as the name of the project you create, so you will have to think about what you want the project to be called. Also note that this name has to be unique in the Eclipse workspace. You cannot have two projects with the same name.

After you have set this field, you will notice that the "Next" button is insensitive, but now the "Finish" button is sensitive. You can click "Finish" now (or use the keyboard shortcut of "Alt+f"), and the Yang IDE will finally create your project.

Creating a Yang File

Now that you have created your project, it is time to create your first Yang file.

When you created the Yang project, you might have noticed the other option next to "YANG Project", which was "YANG File". That is what you will select now. Click "Next" to go to the first wizard page.

First Yang File Wizard Page

This wizard page lets you specify where the new file will be located, and its name.

You have to select the particular project you want the file to go into, and it needs to go into the "src/main/yang" folder (or a different location if you changed that field when creating the project).

You then enter the desired name of the file in the "File name". The file name should have no spaces or "special characters" in it. You can specify a ".yang" extent if you want. If you do not specify an extent, the YANG IDE will create it with the ".yang" extent.

Click "Next" to go to the next wizard page.

Second Yang File Wizard Page

On this wizard page, you set some metadata about the module that is used to initialize the contents of the Yang file.

It has the following fields:

Module Name

This will default to the "base name" of the file name you created. For instance, if the file name you created was "network-setup.yang", this field will default to "network-setup". You should leave this value as is. There is no good reason to define a model with a name different from the file name.

Namespace

This defaults to "urn:opendaylight:xxx", where "xxx" is the "base name" of the file name you created. You should put a lot of thought into designing a namespace naming scheme that is used throughout your organization. It is quite common for this namespace value to look like a "http" URL, but note that that is just a convention, and will not necessarily imply that there is a web page residing at that HTTP address.

Prefix

This defaults to the "base name" of the file name you created. It mostly does not technically matter what you set this to, as long as it is not empty. Conventionally, it should be a "nickname" that is used to refer to the given namespace in an abbreviated form, when referenced in an "import" statement in another Yang model file.

Revision

This has to be a date value in the form of "yyyy-mm-dd", representing the last modified date of this Yang model. The value will default to the current date.

Revision Description

This is just human-readable text, which will go into the "description" field underneath the Yang "revision" field, which will describe what went into this revision.

When all the fields have the content you want, click the "Finish" button to set the YANG IDE create the file in the specified location. It will then present the new file in the editor view for additional modifications.

Accessing Artifacts for Yang Model Imports

You might be working on Yang models that are "abstract" or are intended to be imported by other Yang models. You might also, and more likely, be working on Yang models that import other "abstract" Yang models.

Assuming you are in that latter more common group, you need to consider for yourself, and for your organization, how you are going to get access to those models that you import.

You could use a very simple and primitive approach of somehow obtaining those models from some source as plain files and just copying them into the "src/main/yang" folder of your project. For a simple demo or a "one-off" very short project, that might be sufficient.

A more robust and maintainable approach would be to reference "coordinates" of the artifacts containing Yang models to import. When you specify unique coordinates associated with that artifact, the Yang IDE can retrieve the artifact in the background and make it available for your "import" statements.

Those "coordinates" that I speak of refer to the Maven concepts of "group id", "artifact id", and "version". you may remember "group id" and "artifact id" from the wizard page for creating a Yang project. It is the same idea. If you ever produce Yang model artifacts that other people are going to import, you will want to think more about what you set those values to when you created the project.

For example, the OpenDaylight project produces several importable artifacts that you can specify to get access to common Yang models.

Turning on Indexing for Maven Repositories

Before we talk about how to add dependencies to Maven artifacts with Yang models for import, I need to explain how to make it easier to find those artifacts.

In the Yang project that you have created, the "pom.xml" file (also called a "POM file") is the file that Maven uses to specify dependencies. We will talk about that in a minute, but first we need to talk about "repositories". These are where artifacts are stored.

We are going to have Eclipse show us the "Maven Repositories" view. In the main menu, select "Window" and then "Show View", and then "Other". Like in the "New" dialog, you can enter "maven" in the filter field to limit the list to views with "maven" in the name. Click on the "Maven Repositories" entry and click OK.

This will usually create the view in the bottom panel of the window.

The view presents an outline view of four principal elements:

- Local Repositories
- Global Repositories
- Project Repositories
- Custom Repositories

For this purpose, the only section you care about is "Project Repositories", being the repositories that are only specified in the POM for the project. There should be a "right-pointing arrow" icon on the line. Click that to expand the entry.

You should see two entries there:

- opendaylight-release
- opendaylight-snapshot

You will also see internet URLs associated with each of those repositories.

For this purpose, you only care about the first one. Right-click on that entry and select "Full Index Enabled". The first time you do this on the first project you create, it will spend several minutes walking the entire tree of artifacts available at that repository and "indexing" all of those components. When this is done, searching for available artifacts in that repository will go very quickly.

Adding Dependencies Containing Yang Models

Double-click the "pom.xml" file in your project. Instead of just bringing up the view of an XML file (although you can see that if you like), it presents a GUI form editor with a handful of tabs.

The first tab, "Overview", shows things like the "Group Id", "Artifact Id", and "Version", which represents the "Maven coordinate" of your project, which I have mentioned before.

Now click on the "Dependencies" tab. You will now see two list components, labeled "Dependencies" and "Dependency Management". You only care about the "Dependencies" section.

In the "Dependencies" section, you should see one dependency for an artifact called "yang-binding". This artifact is part of OpenDaylight, but you do not need to know anything about it.

Now click the "Add" button.

This brings up a dialog titled "Select Dependency". It has three fields at the top labeled "Group Id", "Artifact Id", and "Version", with a "Scope" dropdown. You will never have a need to change the "Scope" dropdown, so ignore it. Despite the fact that you will need to get values into these fields, in general usage, you will never have to manually enter values into them, but you will see values being inserted into these fields by the next steps I describe.

Below those fields is a field labeled "Enter groupId, artifactId . . . ". This is effectively a "filter field", like on the "New" dialog, but instead of limiting the list from a short list of choices, the value you enter there will be matched against all of the artifacts that were indexed in the "opendaylight-release" repository (and others). It will match the string you enter as a substring of any groupId or artifactId.

For all of the entries that match that substring, it will list an entry showing the groupId and artifactId, with an expansion arrow. If you open it by clicking on the arrow, you will see individual entries corresponding to each available version of that artifact, along with some metadata about the artifacts between square brackets, mostly indicating what "type" of artifact is.

For your purposes, you only ever want to use "bundle" or "jar" artifacts.

Let us consider an example that many people will probably be using.

In the filter field, enter "ietf-yang-types". Depending on what versions are available, you should see a small handful of "groupId, artifactId" entries there. One of them should be groupId "org.opendaylight.mdsal.model" and artifactId "ietf-yang-types". Click on the expansion arrow to open that.

What you will see at this point depends on what versions are available. You will likely want to select the newest one (most likely top of the list) that is also either a "bundle" or "jar" type artifact.

If you click on that resulting version entry, you should notice at this point that the "Group Id", "Artifact Id", and "Version" fields at the top of the dialog are now filled in with the values corresponding to this artifact and version.

If this is the version that you want, click OK and this artifact will be added to the dependencies in the POM.

This will now make the Yang models found in that artifact available in "import" statements in Yang models, not to mention the completion choices for that "import" statement.

YANG-PUSH

This section describes how to use the YANG-PUSH feature in OpenDaylight and contains contains configuration, administration, and management sections for the feature.

Overview

YANG PUBSUB project allows applications to place subscriptions upon targeted subtrees of YANG datastores residing on remote devices. Changes in YANG objects within the remote subtree can be pushed to an OpenDaylight MD-SAL and to the application as specified without a requiring the controller to make a continuous set of fetch requests.

YANG-PUSH capabilities available

This module contains the base code which embodies the intent of YANG-PUSH requirements for subscription as defined in {i2rs-pub-sub-requirements} [https://datatracker.ietf.org/doc/draft-ietf-i2rs-pub-sub-requirements/]. The mechanism for delivering on these YANG-PUSH requirements over Netconf transport is defined in {netconf-yang-push} [netconf-yang-push: https://tools.ietf.org/html/draft-ietf-netconf-yang-push-00].

Note that in the current release, not all capabilities of draft-ietf-netconf-yang-push are realized. Currently only implemented is **create-subscription** RPC support from ietf-datastore-push@2015-10-15.yang; and this will be for periodic subscriptions only. There of course is intent to provide much additional functionality in future OpenDaylight releases.

Future YANG-PUSH capabilities

Over time, the intent is to flesh out more robust capabilities which will allow OpenDaylight applications to subscribe to YANG-PUSH compliant devices. Capabilities for future releases will include:

Support for subscription change/delete: **modify-subscription** rpc support for all mountpoint devices or particular mountpoint device **delete-subscription** rpc support for all mountpoint devices or particular mountpoint device

Support for static subscriptions: This will enable the receipt of subscription updates pushed from publishing devices where no signaling from the controller has been used to establish the subscriptions.

Support for additional transports: NETCONF is not the only transport of interest to OpenDaylight or the subscribed devices. Over time this code will support Restconf and HTTP/2 transport requirements defined in {netconf-restconf-yang-push} [https://tools.ietf.org/html/draft-voit-netconf-restconf-yang-push-01]

YANG-PUSH Architecture

The code architecture of Yang push consists of two main elements

YANGPUSH Provider YANGPUSH Listener

YANGPUSH Provider receives create-subscription requests from applications and then establishes/registers the corresponding listener which will receive information pushed by a publisher. In addition, YANGPUSH Provider also invokes an augmented OpenDaylight create-subscription RPC which enables applications to register for notification as per rfc5277. This augmentation adds periodic time period (duration) and subscription-id values to the existing RPC parameters. The Java package supporting this capability is "org.opendaylight.yangpush.impl". YangpushDomProvider is the class which supports this YANGPUSH Provider capability.

The YANGPUSH Listener accepts update notifications from a device after they have been de-encapsulated from the NETCONF transport. The YANGPUSH Listener then passes these updates to MD-SAL. This function is implemented via the YangpushDOMNotificationListener class within the "org.opendaylight.yangpush.listner" Java package. Applications should monitor MD-SAL for the availability of newly pushed subscription updates.

1.3 OpenDaylight with Openstack Guide

1.3.1 Overview

OpenStack is a popular open source Infrastructure as a service project, covering compute, storage and network management. OpenStack can use OpenDaylight as its network management provider through the Modular Layer 2 (ML2) north-bound plug-in. OpenDaylight manages the network flows for the OpenStack compute nodes via the OVSDB south-bound plug-in. This page describes how to set that up, and how to tell when everything is working.

1.3.2 Installing OpenStack

Installing OpenStack is out of scope for this document, but to get started, it is useful to have a minimal multi-node OpenStack deployment.

The reference deployment we will use for this document is a 3 node cluster:

- One control node containing all of the management services for OpenStack (Nova, Neutron, Glance, Swift, Cinder, Keystone)
- Two compute nodes running nova-compute
- Neutron using the OVS back-end and vxlan for tunnels

Once you have installed OpenStack, verify that it is working by connecting to Horizon and performing a few operations. To check the Neutron configuration, create two instances on a private subnet bridging to your public network, and verify that you can connect to them, and that they can see each other.

1.3.3 Installing OpenDaylight

OpenStack with NetVirt

OpenStack with NetVirt

Table of Contents

- *OpenStack with NetVirt*
  - *Installing OpenDaylight on an existing OpenStack*
  - *Installing OpenStack and OpenDaylight using DevStack*
  - *Troubleshooting*
  - *Useful Links*

Prerequisites: OpenDaylight requires Java 1.8.0 and Open vSwitch >= 2.5.0

Installing OpenDaylight on an existing OpenStack

- On the control host, Download the latest OpenDaylight release
- Uncompress it as root, and start OpenDaylight (you can start OpenDaylight by running karaf directly, but exiting from the shell will shut it down):

```
tar xvfz distribution-karaf-0.5.1-Boron-SR1.tar.gz
cd distribution-karaf-0.5.1-Boron-SR1
./bin/start # Start OpenDaylight as a server process
```

- Connect to the Karaf shell, and install the odl-netvirt-openstack bundle, dlux and their dependencies:

```
./bin/client # Connect to OpenDaylight with the client
opendaylight-user@root> feature:install odl-netvirt-openstack odl-dlux-core odl-
↪mdsal-apidocs
```

- If everything is installed correctly, you should now be able to log in to the dlux interface on http://CONTROL_HOST:8181/index.html - the default username and password is "admin/admin" (see screenshot below)

Optional - Advanced OpenDaylight Installation - Configurations and Clustering

- ACL Implementation - Security Groups - Stateful:
  - Default implementation used is stateful, requiring OVS compiled with conntrack modules.
  - This requires using a linux kernel that is >= 4.3
  - To check if OVS is running with conntrack support:

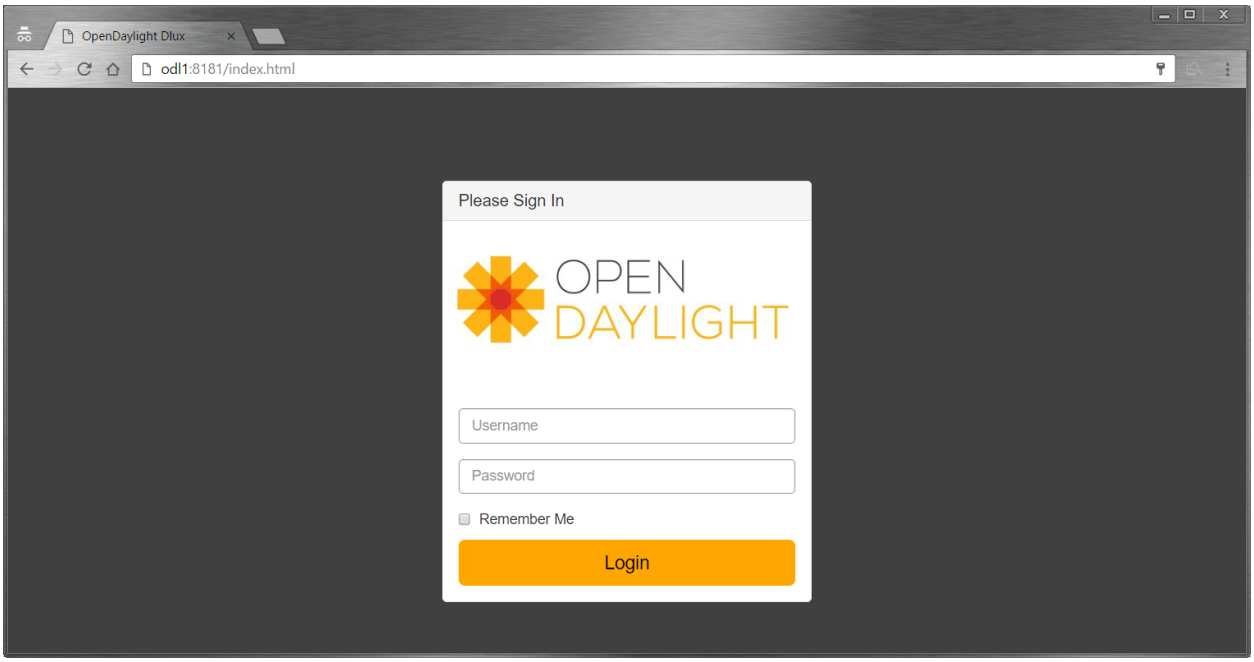


```
root@devstack:~/# lsmod | grep conntrack | grep openvswitch
  nf_conntrack          106496  9 xt_CT,openvswitch,nf_nat,nf_nat_ipv4,xt_
↪conntrack,nf_conntrack_netlink,xt_connmark,nf_conntrack_ipv4,nf_conntrack_
↪ipv6
```

- If the conntrack modules are not installed for OVS, either recompile/install an OVS version with conntrack support, or alternatively configure OpenDaylight to use a non-stateful implementation.
- OpenvSwitch 2.5 with conntrack support can be acquired from this repository for yum based linux distributions:

```
yum install -y http://rdoproject.org/repos/openstack-newton/rdo-release-
↪newton.rpm
yum install -y --nogpgcheck openvswitch
```

- ACL Implementations - Alternative options:
  - "learn" - semi-stateful implementation that does not require conntrack support. This is the most complete non-conntrack implementation.
  - "stateless" - naive security group implementation for TCP connections only. UDP and ICMP packets are allowed by default.
  - "transparent" - no security group support. all traffic is allowed, this is the recommended mode if you don't need to use security groups at all.
  - To configure one of these alternative implementations, the following needs to be done prior to running OpenDaylight:

```
mkdir -p <ODL_FOLDER>/etc/opendaylight/datastore/initial/config/
export CONFFILE=\`find <ODL_FOLDER> -name "\*aclservice\*config.xml"\`
cp \CONFFILE <ODL_FOLDER>/etc/opendaylight/datastore/initial/config/netvirt-
↪aclservice-config.xml
sed -i s/stateful/<learn/transparent>/ <ODL_FOLDER>/etc/opendaylight/
↪datastore/initial/config/netvirt-aclservice-config.xml
```

(continues on next page)

(continued from previous page)

```
cat <ODL_FOLDER>/etc/opendaylight/datastore/initial/config/netvirt-aclservice-
↪config.xml
```

- Running multiple OpenDaylight controllers in a cluster:
  - For redundancy, it is possible to run OpenDaylight in a 3-node cluster.
  - More info on Clustering available here.
  - To configure OpenDaylight in clustered mode, run <ODL_FOLDER>/bin/configure_cluster.sh on each node prior to running OpenDaylight. This script is used to configure cluster parameters on this controller. The user should restart controller to apply changes.

```
Usage: ./configure_cluster.sh <index> <seed_nodes_list>
- index: Integer within 1..N, where N is the number of seed nodes.
- seed_nodes_list: List of seed nodes, separated by comma or space.
```

  - The address at the provided index should belong this controller. When running this script on multiple seed nodes, keep the seed_node_list same, and vary the index from 1 through N.
  - Optionally, shards can be configured in a more granular way by modifying the file "custom_shard_configs.txt" in the same folder as this tool. Please see that file for more details.

Note: OpenDaylight should be restarted after applying any of the above changes via configuration files.

Ensuring OpenStack network state is clean

When using OpenDaylight as the Neutron back-end, OpenDaylight expects to be the only source of truth for Neutron configurations. Because of this, it is necessary to remove existing OpenStack configurations to give OpenDaylight a clean slate.

- Delete instances:

```
nova list
nova delete <instance names>
```

- Remove links from subnets to routers:

```
neutron subnet-list
neutron router-list
neutron router-port-list <router name>
neutron router-interface-delete <router name> <subnet ID or name>
```

- Delete subnets, networks, routers:

```
neutron subnet-delete <subnet name>
neutron net-list
neutron net-delete <net name>
neutron router-delete <router name>
```

- Check that all ports have been cleared - at this point, this should be an empty list:

```
neutron port-list
```

Ensure Neutron is stopped

While Neutron is managing the OVS instances on compute and control nodes, OpenDaylight and Neutron can be in conflict. To prevent issues, we turn off Neutron server on the network controller, and Neutron's Open vSwitch agents on all hosts.

- Turn off neutron-server on control node:

```
systemctl stop neutron-server
systemctl stop neutron-l3-agent
```

- On each node in the cluster, shut down and disable Neutron's agent services to ensure that they do not restart after a reboot:

```
systemctl stop neutron-openvswitch-agent
systemctl disable
neutron-openvswitch-agent
systemctl stop neutron-l3-agent
systemctl disable neutron-l3-agent
```

Configuring Open vSwitch to be managed by OpenDaylight

On each host (both compute and control nodes) we will clear the pre-existing Open vSwitch config and set OpenDaylight to manage the switch:

- Stop the Open vSwitch service, and clear existing OVSDB (OpenDaylight expects to manage vSwitches completely):

```
systemctl stop openvswitch
rm -rf /var/log/openvswitch/*
rm -rf /etc/openvswitch/conf.db
systemctl start openvswitch
```

- At this stage, your Open vSwitch configuration should be empty:

```
[root@odl-compute2 ~]# ovs-vsctl show
9f3b38cb-eefc-4bc7-828b-084b1f66fbfd
    ovs_version: "2.5.1"
```

- Set OpenDaylight as the manager on all nodes:

```
ovs-vsctl set-manager tcp:{CONTROL_HOST}:6640
```

- Set the IP to be used for VXLAN connectivity on all nodes. This IP must correspond to an actual linux interface on each machine.

```
sudo ovs-vsctl set Open_vSwitch . other_config:local_ip=<ip>
```

- You should now see a new section in your Open vSwitch configuration showing that you are connected to the OpenDaylight server via OVSDB, and OpenDaylight will automatically create a br-int bridge that is connected via OpenFlow to the controller:

```
[root@odl-compute2 ~]# ovs-vsctl show
9f3b38cb-eefc-4bc7-828b-084b1f66fbfd
     Manager "tcp:172.16.21.56:6640"
```

(continues on next page)

(continued from previous page)

```
        is_connected: true
    Bridge br-int
        Controller "tcp:172.16.21.56:6633"
            is_connected: true
        fail_mode: secure
        Port br-int
            Interface br-int
    ovs_version: "2.5.1"

[root@odl-compute2 ~]# ovs-vsctl get Open_vSwitch . other_config
{local_ip="10.0.42.161"}
```

- If you do not see the result above (specifically, if you do not see "is_connected: true" in the Manager section or in the Controller section), you may not have a security policies in place to allow Open vSwitch remote administration.

Note:

There might be iptables restrictions - if so the relevant ports should be opened (6640, 6653).

If SELinux is running on your linux, set to permissive mode on all nodes and ensure it stays that way after boot.

```
setenforce 0
sed -i -e 's/SELINUX=enforcing/SELINUX=permissive/g' /etc/selinux/config
```

- Make sure all nodes, including the control node, are connected to OpenDaylight.
- If you reload DLUX, you should now see that all of your Open vSwitch nodes are now connected to OpenDaylight.

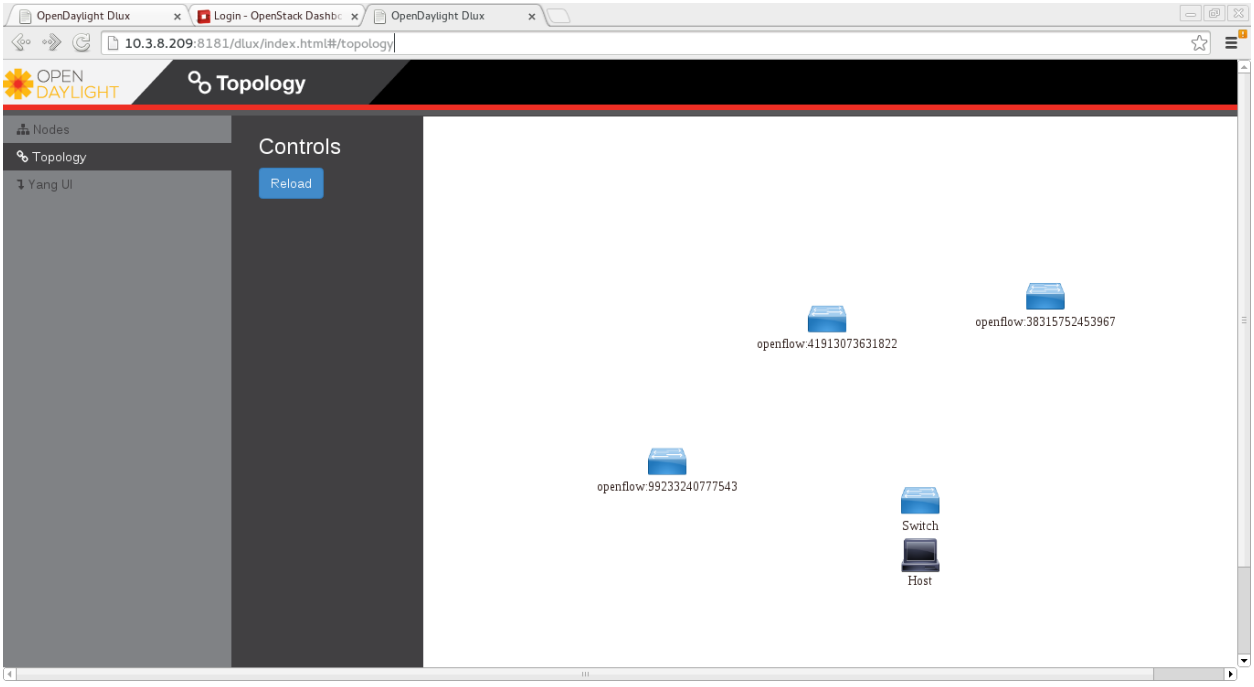

- If something has gone wrong, check `data/log/karaf.log` under the OpenDaylight distribution directory. If you do not see any interesting log entries, set logging for netvirt to TRACE level inside Karaf and try again:

```
log:set TRACE netvirt
```

Configuring Neutron to use OpenDaylight

Once you have configured the vSwitches to connect to OpenDaylight, you can now ensure that OpenStack Neutron is using OpenDaylight.

This requires the neutron networking-odl module to be installed. | `pip install networking-odl`

First, ensure that port 8080 (which will be used by OpenDaylight to listen for REST calls) is available. By default, swift-proxy-service listens on the same port, and you may need to move it (to another port or another host), or disable that service. It can be moved to a different port (e.g. 8081) by editing `/etc/swift/proxy-server.conf` and `/etc/cinder/cinder.conf`, modifying iptables appropriately, and restarting swift-proxy-service. Alternatively, OpenDaylight can be configured to listen on a different port, by modifying the `jetty.port` property value in `etc/jetty.conf`.

```
<Set name="port">
    <Property name="jetty.port" default="8080" />
</Set>
```

- Configure Neutron to use OpenDaylight's ML2 driver:

```
crudini --set /etc/neutron/plugins/ml2/ml2_conf.ini ml2 mechanism_drivers␣
↪opendaylight
crudini --set /etc/neutron/plugins/ml2/ml2_conf.ini ml2 tenant_network_types vxlan

cat <<EOT>> /etc/neutron/plugins/ml2/ml2_conf.ini
[ml2_odl]
url = http://{CONTROL_HOST}:8080/controller/nb/v2/neutron
password = admin
username = admin
EOT
```

- Configure Neutron to use OpenDaylight's odl-router service plugin for L3 connectivity:

```
crudini --set /etc/neutron/plugins/neutron.conf DEFAULT service_plugins odl-router
```

- Configure Neutron DHCP agent to provide metadata services:

```
crudini --set /etc/neutron/plugins/dhcp_agent.ini DEFAULT force_metadata True
```

Note:

If the OpenStack version being used is Newton, this workaround should be applied,
configuring the Neutron DHCP agent to use vsctl as the OVSDB interface:

```
crudini --set /etc/neutron/plugins/dhcp_agent.ini OVS ovsdb_interface vsctl
```

- Reset Neutron's database

```
mysql -e "DROP DATABASE IF EXISTS neutron;"
mysql -e "CREATE DATABASE neutron CHARACTER SET utf8;"
/usr/local/bin/neutron-db-manage --config-file /etc/neutron/neutron.conf --config-
↪file /etc/neutron/plugins/ml2/ml2_conf.ini upgrade head
```

- Restart neutron-server:

```
systemctl start neutron-server
```

Verifying it works

- Verify that OpenDaylight's ML2 interface is working:

```
curl -u admin:admin http://{CONTROL_HOST}:8080/controller/nb/v2/neutron/networks

{
   "networks" : [ ]
}
```

If this does not work or gives an error, check Neutron's log file in `/var/log/neutron/server.log`. Error messages here should give some clue as to what the problem is in the connection with OpenDaylight.

- Create a network, subnet, router, connect ports, and start an instance using the Neutron CLI:

```
neutron router-create router1
neutron net-create private
neutron subnet-create private --name=private_subnet 10.10.5.0/24
neutron router-interface-add router1 private_subnet
nova boot --flavor <flavor> --image <image id> --nic net-id=<network id> test1
nova boot --flavor <flavor> --image <image id> --nic net-id=<network id> test2
```

At this point, you have confirmed that OpenDaylight is creating network end-points for instances on your network and managing traffic to them.

VMs can be reached using Horizon console, or alternatively by issuing `nova get-vnc-console <vm> novnc` Through the console, connectivity between VMs can be verified.

Adding an external network for floating IP connectivity

- In order to connect to the VM using a floating IP, we need to configure external network connectivity, by creating an external network and subnet. This external network must be linked to a physical port on the machine, which will provide connectivity to an external gateway.

```
sudo ovs-vsctl set Open_vSwitch . other_config:provider_mappings=physnet1:eth1
neutron net-create public-net -- --router:external --is-default --
↪provider:network_type=flat --provider:physical_network=physnet1
neutron subnet-create --allocation-pool start=10.10.10.2,end=10.10.10.254 --
↪gateway 10.10.10.1 --name public-subnet public-net 10.10.0.0/16 -- --enable_
↪dhcp=False
neutron router-gateway-set router1 public-net

neutron floatingip-create public-net
nova floating-ip-associate test1 <floating_ip>
```

Installing OpenStack and OpenDaylight using DevStack

The easiest way to load and OpenStack setup using OpenDaylight is by using devstack, which does all the steps mentioned in previous sections. | `git clone https://git.openstack.org/openstack-dev/devstack`

- The following lines need to be added to your local.conf:

```
enable_plugin networking-odl http://git.openstack.org/openstack/networking-odl
↪<branch>
ODL_MODE=allinone
Q_ML2_PLUGIN_MECHANISM_DRIVERS=opendaylight,logger
ODL_GATE_SERVICE_PROVIDER=vpnservice
disable_service q-l3
ML2_L3_PLUGIN=odl-router
ODL_PROVIDER_MAPPINGS={PUBLIC_PHYSICAL_NETWORK}:<external linux interface>
```

- More details on using devstack can be found in the following links:
  - Devstack All-In-One Single Machine Tutorial
  - Devstack networking-odl README

Troubleshooting

VM DHCP Issues

- Trigger DHCP requests - access VM console:
  - View log: `nova console-log <vm>`
  - Access using VNC console: nova get-vnc-console <vm> novnc
  - Trigger DHCP requests: `sudo ifdown eth0 ; sudo ifup eth0`

```
udhcpc (v1.20.1) started
Sending discover...
Sending select for 10.0.123.3...
Lease of 10.0.123.3 obtained, lease time 86400 # This only happens when DHCP␣
↪is properly obtained.
```

- Check if the DHCP requests are reaching the qdhcp agent using the following commands on the OpenStack controller:

```
sudo ip netns
sudo ip netns exec qdhcp-xxxxx ifconfig # xxxx is the neutron network id
sudo ip netns exec qdhcp-xxxxx tcpdump -nei tapxxxxx # xxxxx is the neutron port␣
↪id

# Valid request and response:
15:08:41.684932 fa:16:3e:02:14:bb > ff:ff:ff:ff:ff:ff, ethertype IPv4 (0x0800),␣
↪length 329: 0.0.0.0.68 > 255.255.255.255.67: BOOTP/DHCP, Request from␣
↪fa:16:3e:02:14:bb, length 287
15:08:41.685152 fa:16:3e:79:07:98 > fa:16:3e:02:14:bb, ethertype IPv4 (0x0800),␣
↪length 354: 10.0.123.2.67 > 10.0.123.3.68: BOOTP/DHCP, Reply, length 312
```

- If the requests aren't reaching qdhcp:
  - Verify VXLAN tunnels exist between compute and control nodes by using `ovs-vsctl show`

- Run the following commands to debug the OVS processing of the DHCP request packet:
  `ovs-ofctl -OOpenFlow13 dump-ports-desc br-int` # retrieve VMs ofport and MAC
  `ovs-appctl ofproto/trace br-int in_port=<ofport>,dl_src=<mac>, dl_dst=ff:ff:ff:ff:ff:ff,udp,ip_src=0.0.0.0,ip_dst=255.255.255.255 | grep "Rule\|action"`

```
root@devstack:~# ovs-appctl ofproto/trace br-int in_port=1,dl_
↪src=fe:16:3e:33:8b:d8,dl_dst=ff:ff:ff:ff:ff:ff,udp,ip_src=0.0.0.0,ip_
↪dst=255.255.255.255 | grep "Rule\|action"
    Rule: table=0 cookie=0x8000000 priority=1,in_port=1
    OpenFlow actions=write_metadata:0x20000000001/0xffffff0000000001,goto_
↪table:17
        Rule: table=17 cookie=0x8000001 priority=5,metadata=0x20000000000/
↪0xffffff0000000000
        OpenFlow actions=write_metadata:0xc0000200000222e2/0xfffffffffffffffe,
↪goto_table:19
            Rule: table=19 cookie=0x1080000 priority=0
            OpenFlow actions=resubmit(,17)
                Rule: table=17 cookie=0x8040000 priority=6,
↪metadata=0xc000020000000000/0xffffff0000000000
                OpenFlow actions=write_metadata:0xe00002138a000000/
↪0xffffffffffffffff,goto_table:50
                    Rule: table=50 cookie=0x8050000 priority=0
                    OpenFlow actions=CONTROLLER:65535,goto_table:51
                        Rule: table=51 cookie=0x8030000 priority=0
                        OpenFlow actions=goto_table:52
                            Rule: table=52 cookie=0x870138a priority=5,
↪metadata=0x138a000001/0xffff000001
                            OpenFlow actions=write_actions(group:210003)
    Datapath actions: drop

root@devstack:~# ovs-ofctl -OOpenFlow13 dump-groups br-int | grep 'group_
↪id=210003'
    group_id=210003,type=all
```

- If the requests are reaching qdhcp, but the response isn't arriving to the VM:
  - Locate the compute the VM is residing on (can use `nova show <vm>`).
    - If the VM is on the same node as the qdhcp namespace, `ofproto/trace` can be used to track the packet:
      `ovs-appctl ofproto/trace br-int in_port=<dhcp_ofport>,dl_src=<dhcp_port_mac>,dl_dst=<vm_port_mac>, udp,ip_src=<dhcp_port_ip>,ip_dst=<vm_port_ip> | grep "Rule\|action"`

```
root@devstack:~# ovs-appctl ofproto/trace br-int in_port=2,dl_
↪src=fa:16:3e:79:07:98,dl_dst=fa:16:3e:02:14:bb,udp,ip_src=10.0.123.2,ip_
↪dst=10.0.123.3 | grep "Rule\|action"
    Rule: table=0 cookie=0x8000000 priority=4,in_port=2
    OpenFlow actions=write_metadata:0x10000000000/0xffffff0000000001,goto_
↪table:17
        Rule: table=17 cookie=0x8000001 priority=5,metadata=0x10000000000/
↪0xffffff0000000000
        OpenFlow actions=write_metadata:0x60000100000222e0/
↪0xfffffffffffffffe,goto_table:19
            Rule: table=19 cookie=0x1080000 priority=0
```

(continues on next page)

(continued from previous page)

```
            OpenFlow actions=resubmit(,17)
                Rule: table=17 cookie=0x8040000 priority=6,
↪metadata=0x6000010000000000/0xffffff0000000000
                OpenFlow actions=write_metadata:0x7000011389000000/
↪0xfffffffffffffffe,goto_table:50
                    Rule: table=50 cookie=0x8051389 priority=20,
↪metadata=0x11389000000/0xffffffffff000000,dl_src=fa:16:3e:79:07:98
                    OpenFlow actions=goto_table:51
                        Rule: table=51 cookie=0x8031389 priority=20,
↪metadata=0x1389000000/0xffff000000,dl_dst=fa:16:3e:02:14:bb
                        OpenFlow actions=load:0x300->NXM_NX_REG6[],
↪resubmit(,220)
                            Rule: table=220 cookie=0x8000007 priority=7,
↪reg6=0x300
                            OpenFlow actions=output:3
```

* If the VM isn't on the same node as the qdhcp namepsace:
  - Check if the packet is arriving via VXLAN by running `tcpdump -nei <vxlan_port> port 4789`
  - If it is arriving via VXLAN, the packet can be tracked on the compute node rules, using `ofproto/trace` in a similiar manner to the previous section. Note that packets arriving from a tunnels have a unique tunnel_id (VNI) that should be used as well in the trace, due to the special processing of packets arriving from a VXLAN tunnel.

Floating IP Issues

- If you have assigned an external network and associated a floating IP to a VM but there is still no connectivity:
  - Verify the external gateway IP is reachable through the provided provider network port.
  - Verify OpenDaylight has successfully resolved the MAC address of the external gateway IP. This can be verified by searching for the line "Installing ext-net group" in the karaf.log.
  - Locate the compute the VM is residing on (can use `nova show <vm>`).
  - Run a ping to the VM floating IP.
  - If the ping fails, execute a flow dump of br-int, and search for the flows that are relevant to the VM's floating IP address: `ovs-ofctl -OOpenFlow13 dump-flows br-int | grep "<floating_ip>"`
    * Are there packets on the incoming flow (matching dst_ip=<floating_ip>)?
      If not this probably means the provider network has not been set up properly, verify provider_mappings configuration and the configured external network physical_network value match. Also verify that the Flat/VLAN network configured is actually reachable via the configured port.
    * Are there packets on the outgoing flow (matching src_ip=<floating_ip>)?
      If not, this probably means that OpenDaylight is failing to resolve the MAC of the provided external gateway, required for forwarding packets to the external network.
    * Are there packets being sent on the external network port?
      This can be checked using `tcpdump <port>` or by viewing the appropriate OpenFlow rules. The mapping between the OpenFlow port number and the linux interface can be acquired using `ovs-ofctl dump-ports-desc br-int`

```
ovs-ofctl -OOpenFlow13 dump-flows br-int | grep "<floating_ip>"
cookie=0x8000003, duration=436.710s, table=21, n_packets=190, n_
↪bytes=22602, priority=42,ip,metadata=0x222e2/0xfffffffe,nw_dst=10.64.98.
↪17 actions=goto_table:25
cookie=0x8000004, duration=436.739s, table=25, n_packets=190, n_
↪bytes=22602, priority=10,ip,nw_dst=10.64.98.17 actions=set_field:10.0.
↪123.3->ip_dst,write_metadata:0x222e0/0xfffffffe,goto_table:27
cookie=0x8000004, duration=436.730s, table=26, n_packets=120, n_
↪bytes=15960, priority=10,ip,metadata=0x222e0/0xfffffffe,nw_src=10.0.123.
↪3 actions=set_field:10.64.98.17->ip_src,write_metadata:0x222e2/
↪0xfffffffe,goto_table:28
cookie=0x8000004, duration=436.728s, table=28, n_packets=120, n_
↪bytes=15960, priority=10,ip,metadata=0x222e2/0xfffffffe,nw_src=10.64.98.
↪17 actions=set_field:fa:16:3e:ec:a8:84->eth_src,group:200000
```

Useful Links

- NetVirt Tables Pipeline
- NetVirt Wiki Page
- NetVirt Basic Tutorial (OpenDaylight Summit 2016)
- NetVirt Advanced Tutorial (OpenDaylight Summit 2016)
- Other OpenDaylight Documentation

OpenStack with GroupBasedPolicy

This section is for Application Developers and Network Administrators who are looking to integrate Group Based Policy with OpenStack.

To enable the **GBP** Neutron Mapper feature, at the karaf console:

```
feature:install odl-groupbasedpolicy-neutronmapper
```

Neutron Mapper has the following dependencies that are automatically loaded:

```
odl-neutron-service
```

Neutron Northbound implementing REST API used by OpenStack

```
odl-groupbasedpolicy-base
```

Base **GBP** feature set, such as policy resolution, data model etc.

```
odl-groupbasedpolicy-ofoverlay
```

For this release, **GBP** has one renderer, hence this is loaded by default.

REST calls from OpenStack Neutron are by the Neutron NorthBound project.

GBP provides the implementation of the Neutron V2.0 API.

Features

List of supported Neutron entities:

- Port
- Network
  - Standard Internal
  - External provider L2/L3 network
- Subnet
- Security-groups
- Routers
  - Distributed functionality with local routing per compute
  - External gateway access per compute node (dedicated port required)
  - Multiple routers per tenant
- FloatingIP NAT
- IPv4/IPv6 support

The mapping of Neutron entities to **GBP** entities is as follows:

Neutron Port

The Neutron port is mapped to an endpoint.

The current implementation supports one IP address per Neutron port.

An endpoint and L3-endpoint belong to multiple EndpointGroups if the Neutron port is in multiple Neutron Security Groups.

The key for endpoint is L2-bridge-domain obtained as the parent of L2-flood-domain representing Neutron network. The MAC address is from the Neutron port. An L3-endpoint is created based on L3-context (the parent of the L2-bridge-domain) and IP address of Neutron Port.

Neutron Network

A Neutron network has the following characteristics:

- defines a broadcast domain
- defines a L2 transmission domain
- defines a L2 name space.

To represent this, a Neutron Network is mapped to multiple **GBP** entities. The first mapping is to an L2 flood-domain to reflect that the Neutron network is one flooding or broadcast domain. An L2-bridge-domain is then associated as the parent of L2 flood-domain. This reflects both the L2 transmission domain as well as the L2 addressing namespace.

The third mapping is to L3-context, which represents the distinct L3 address space. The L3-context is the parent of L2-bridge-domain.

Neutron Subnet

Neutron subnet is associated with a Neutron network. The Neutron subnet is mapped to a *GBP* subnet where the parent of the subnet is L2-flood-domain representing the Neutron network.

Neutron Security Group

GBP entity representing Neutron security-group is EndpointGroup.

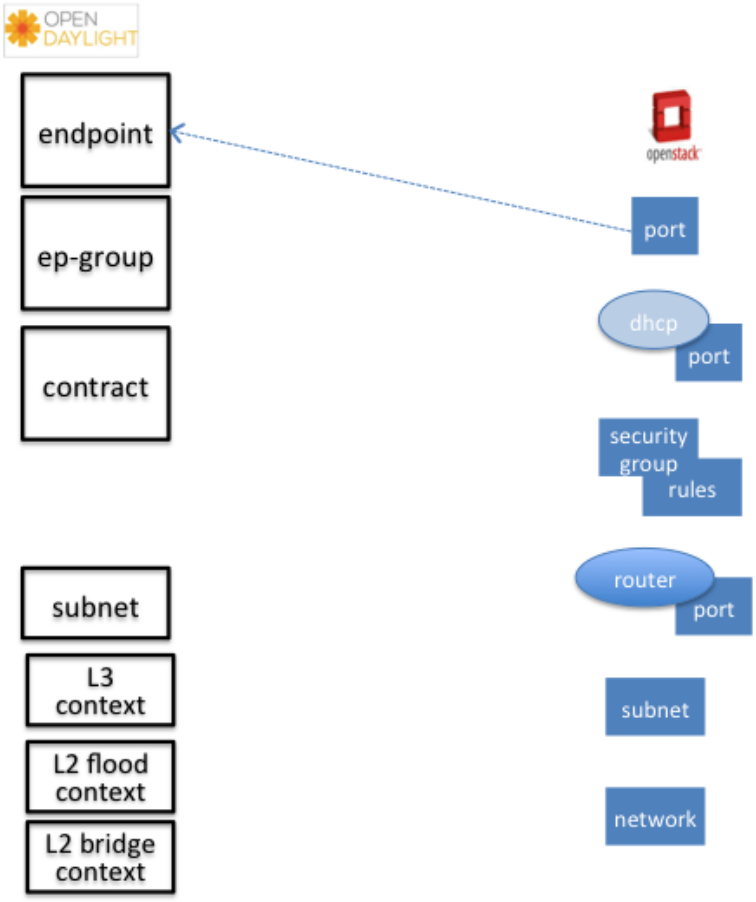

Fig. 147: Neutron Port

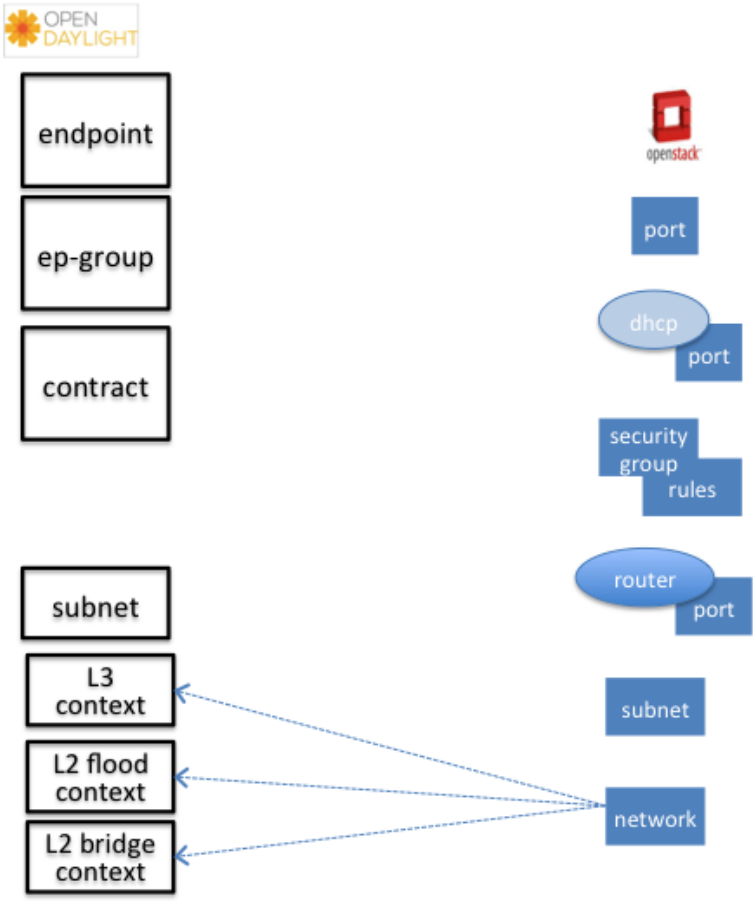

Fig. 148: Neutron Network

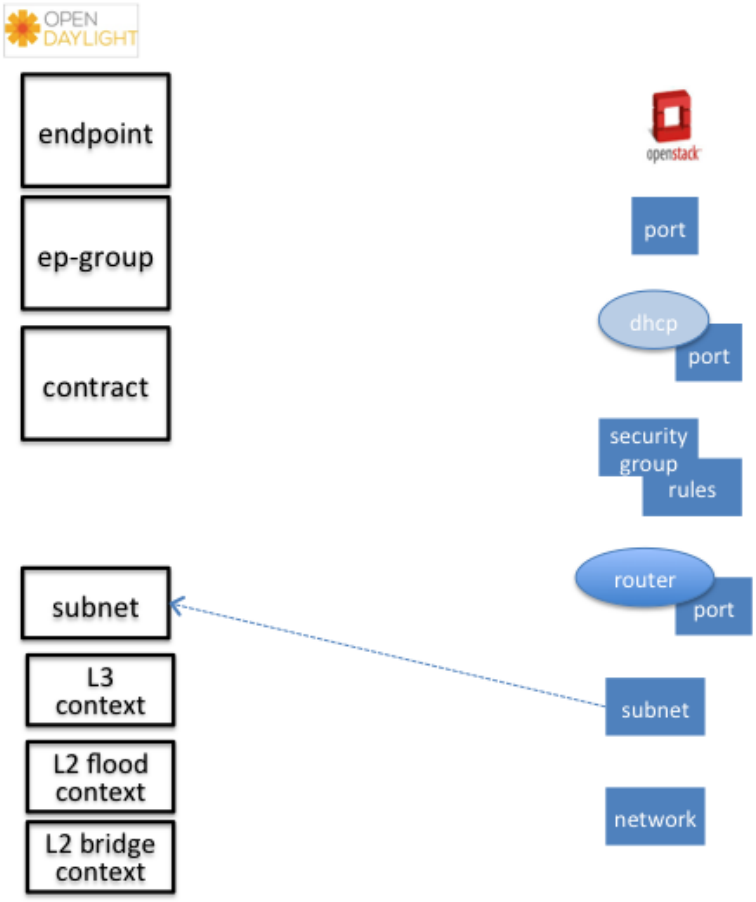

Fig. 149: Neutron Subnet

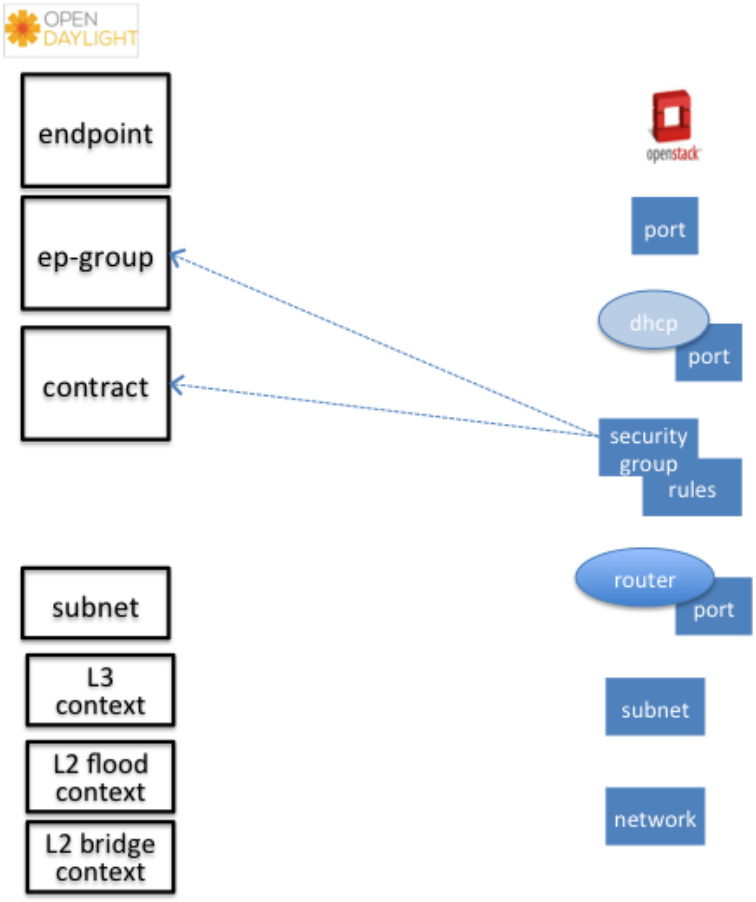

Fig. 150: Neutron Security Group and Rules

Infrastructure EndpointGroups

Neutron-mapper automatically creates EndpointGroups to manage key infrastructure items such as:

- DHCP EndpointGroup - contains endpoints representing Neutron DHCP ports
- Router EndpointGroup - contains endpoints representing Neutron router interfaces
- External EndpointGroup - holds L3-endpoints representing Neutron router gateway ports, also associated with FloatingIP ports.

Neutron Security Group Rules

This mapping is most complicated among all others because Neutron security-group-rules are mapped to contracts with clauses, subjects, rules, action-refs, classifier-refs, etc. Contracts are used between endpoint groups representing Neutron Security Groups. For simplification it is important to note that Neutron security-group-rules are similar to a **GBP** rule containing:

- classifier with direction
- action of *allow*.

Neutron Routers

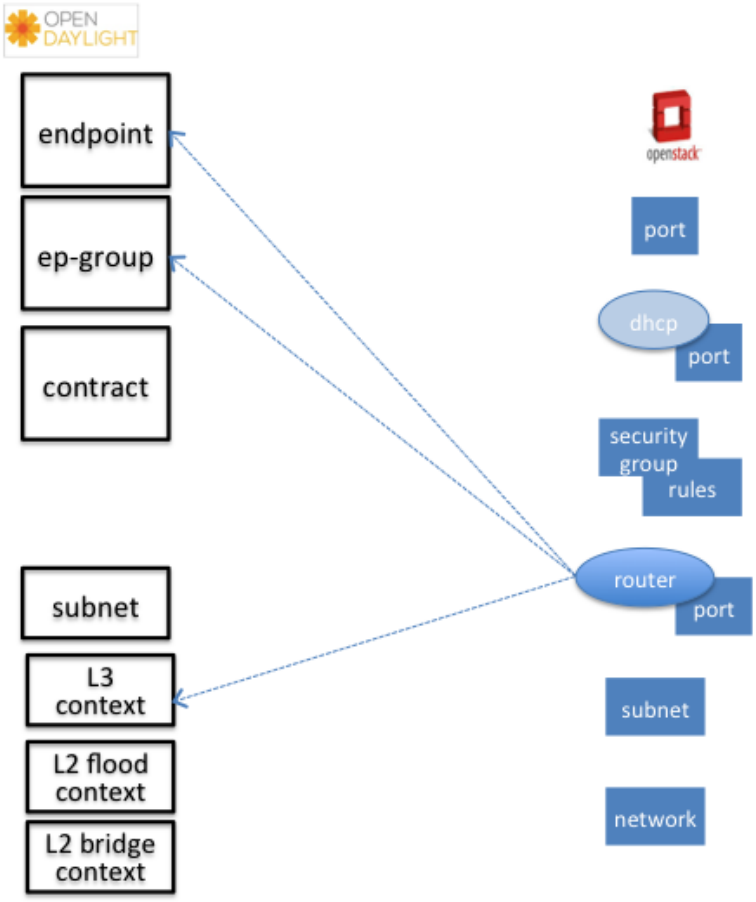

Fig. 151: Neutron Router

Neutron router is represented as a L3-context. This treats a router as a Layer3 namespace, and hence every network attached to it a part of that Layer3 namespace.

This allows for multiple routers per tenant with complete isolation.

The mapping of the router to an endpoint represents the router's interface or gateway port.

The mapping to an EndpointGroup represents the internal infrastructure EndpointGroups created by the **GBP** Neutron Mapper

When a Neutron router interface is attached to a network/subnet, that network/subnet and its associated endpoints or Neutron Ports are seamlessly added to the namespace.

Neutron FloatingIP

When associated with a Neutron Port, this leverages the *GBP* OfOverlay renderer's NAT capabilities.

A dedicated *external* interface on each Nova compute host allows for disitributed external access. Each Nova instance associated with a FloatingIP address can access the external network directly without having to route via the Neutron controller, or having to enable any form of Neutron distributed routing functionality.

Assuming the gateway provisioned in the Neutron Subnet command for the external network is reachable, the combination of *GBP* Neutron Mapper and OfOverlay renderer will automatically ARP for this default gateway, requiring no user intervention.

Troubleshooting within GBP

Logging level for the mapping functionality can be set for package org.opendaylight.groupbasedpolicy.neutron.mapper. An example of enabling TRACE logging level on karaf console:

```
log:set TRACE org.opendaylight.groupbasedpolicy.neutron.mapper
```

Neutron mapping example

As an example for mapping can be used creation of Neutron network, subnet and port. When a Neutron network is created 3 **GBP** entities are created: l2-flood-domain, l2-bridge-domain, l3-context.

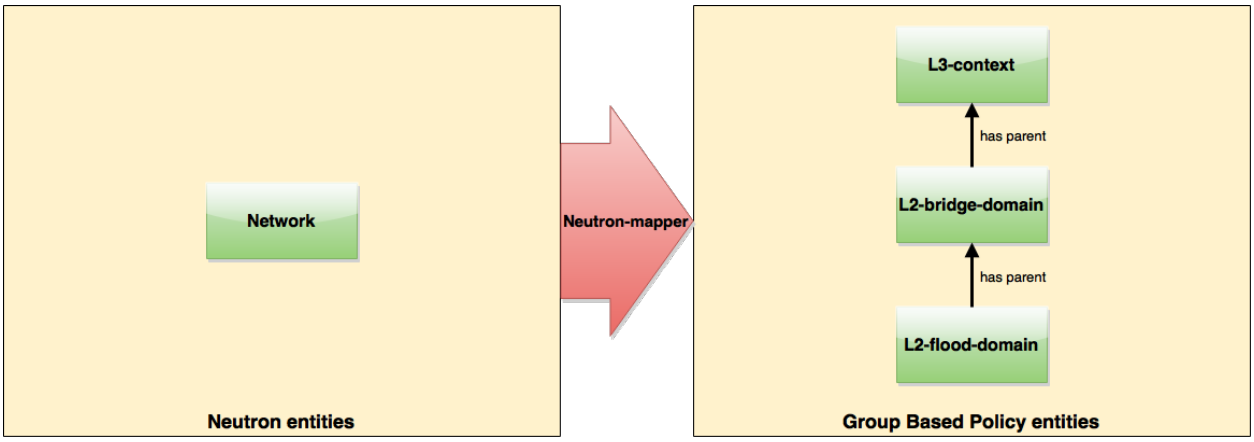

Fig. 152: Neutron network mapping

After an subnet is created in the network mapping looks like this.

If an Neutron port is created in the subnet an endpoint and l3-endpoint are created. The endpoint has key composed from l2-bridge-domain and MAC address from Neutron port. A key of l3-endpoint is compesed from l3-context and IP address. The network containment of endpoint and l3-endpoint points to the subnet.

Configuring GBP Neutron

No intervention passed initial OpenStack setup is required by the user.

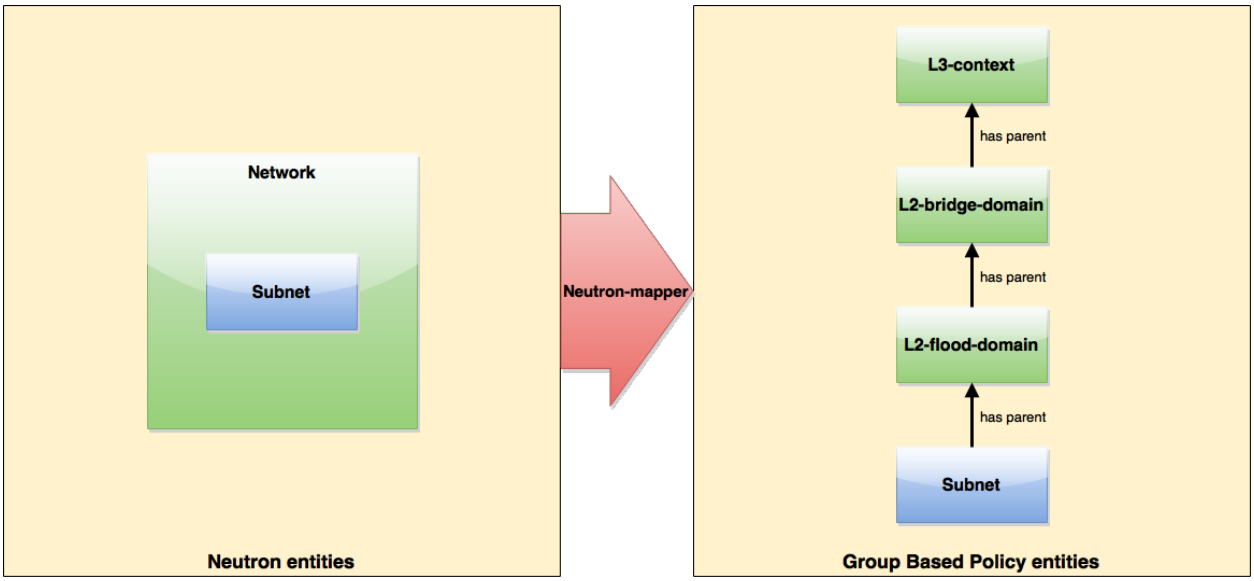

Fig. 153: Neutron subnet mapping

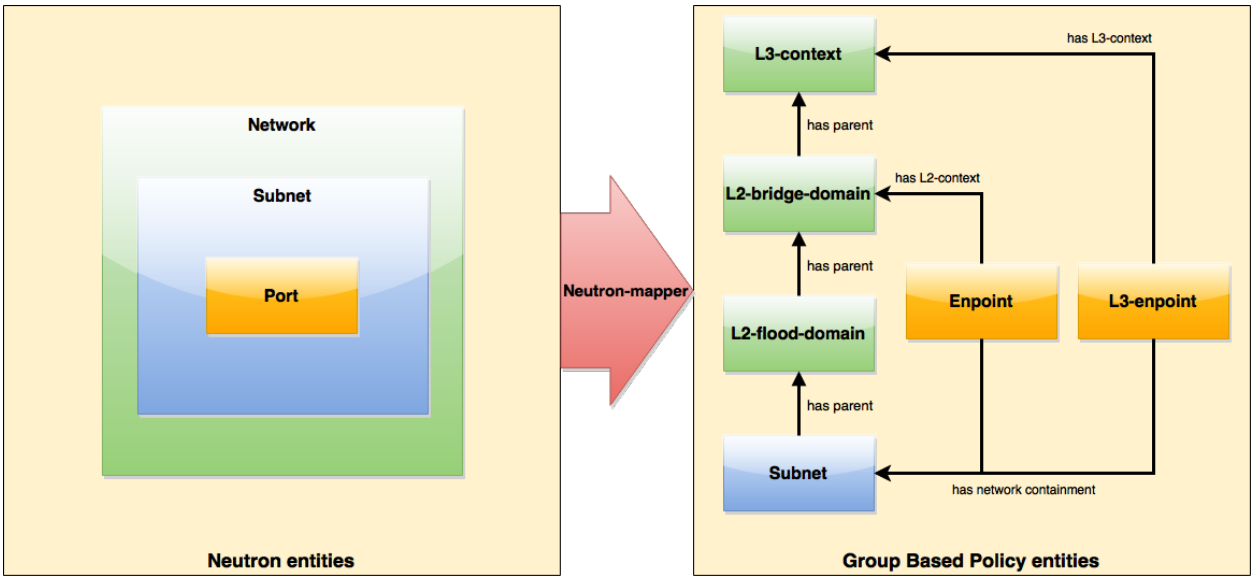

Fig. 154: Neutron port mapping

More information about configuration can be found in our DevStack demo environment on the GBP wiki.

Administering or Managing GBP Neutron

For consistencies sake, all provisioning should be performed via the Neutron API. (CLI or Horizon).

The mapped policies can be augmented via the **GBP** UX,UX, to:

- Enable Service Function Chaining
- Add endpoints from outside of Neutron i.e. VMs/containers not provisioned in OpenStack
- Augment policies/contracts derived from Security Group Rules
- Overlay additional contracts or groupings

Tutorials

A DevStack demo environment can be found on the GBP wiki.

Using Groupbasedpolicy's Neutron VPP Mapper

Overview

Neutron VPP Mapper implements features for support policy-based routing for OpenStack Neutron interface involving VPP devices. It allows using of policy-based schemes defined in GBP controller in a network consisting of OpenStack-provided nodes routed by a VPP node.

Architecture

Neutron VPP Mapper listens to Neutron data store change events, as well as being able to access directly the store. If the data changed match certain criteria (see *Processing Neutron Configuration*), Neutron VPP Mapper converts Neutron data specifically required to render a VPP node configuration with a given End Point, e.g., the virtual host interface name assigned to a `vhostuser` socket. Then the mapped data is stored in the VPP info data store.

Administering Neutron VPP Mapper

To use the Neutron VPP Mapper in Karaf, at least the following Karaf features must be installed:

- odl-groupbasedpolicy-neutron-vpp-mapper
- odl-vbd-ui

Initial pre-requisites

A topology should exist in config datastore, it is necessary to define a node with a particular `node-id`. Later, `node-id` will be used as a physical location reference in VPP renderer's bridge domain:

```
GET http://localhost:8181/restconf/config/network-topology:network-topology/

{
    "network-topology":{
       "topology":[
            {
                "topology-id":"datacentre",
                "node":[
                    {
                       "node-id":"dut2",
                       "vlan-tunnel:super-interface":"GigabitEthernet0/9/0",
                       "termination-point":[
                            {
                                "tp-id":"GigabitEthernet0/9/0",
                                "neutron-provider-topology:physical-interface":{
                                    "interface-name":"GigabitEthernet0/9/0"
                                }
                            }
                        ]
                    }
                ]
            }
        ]
    }
}
```

Processing Neutron Configuration

`NeutronListener` listens to the changes in Neutron datatree in config datastore. It filters the changes, processing only `network` and `port` entities.

For a `network` entity it is checked that it has `physical-network` parameter set (i.e., it is backed-up by a physical network), and that `network-type` is `vlan-network` or `"flat"`, and if this check has passed, a related bridge domain is created in VPP Renderer config datastore (`http://{{controller}}:{{port}}/restconf/config/vpp-renderer:config`), referenced to network by `vlan` field.

In case of `"vlan-network"`, the `vlan` field contains the same value as `neutron-provider-ext:segmentation-id` of network created by Neutron.

In case of `"flat"`, the VLAN specific parameters are not filled out.

Note: In case of VXLAN network (i.e. `network-type` is `"vxlan-network"`), no information is actually written into VPP Renderer datastore, as VXLAN is used for tenant-network (so no packets are going outside). Instead, VPP Renderer looks up GBP flood domains corresponding to existing VPP bridge domains trying to establish a VXLAN tunnel between them.

For a `port` entity it is checked that `vif-type` contains `"vhostuser"` substring, and that `device-owner` contains a specific substring, namely `"compute"`, `"router"` or `"dhcp"`.

In case of `"compute"` substring, a `vhost-user` is written to VPP Renderer config datastore.

In case of `"dhcp"` or `"router"`, a `tap` is written to VPP Renderer config datastore.

Input/output examples

OpenStack is creating network, and these data are being put into the data store:

```
PUT http://{{controller}}:{{port}}/restconf/config/neutron:neutron/networks

{
    "networks": {
        "network": [
            {
                "uuid": "43282482-a677-4102-87d6-90708f30a115",
                "tenant-id": "94836b88-0e56-4150-aaa7-60f1c2b67faa",
                "neutron-provider-ext:segmentation-id": "2016",
                "neutron-provider-ext:network-type": "neutron-networks:network-type-
↪vlan",
                "neutron-provider-ext:physical-network": "datacentre",
                "neutron-L3-ext:external": true,
                "name": "drexternal",
                "shared": false,
                "admin-state-up": true,
                "status": "ACTIVE"
            }
        ]
    }
}
```

Checking bridge domain in VPP Renderer config data store. Note that `physical-location-ref` is referring to `"dut2"`, paired by `neutron-provider-ext:physical-network` -> `topology-id`:

```
GET http://{{controller}}:{{port}}/restconf/config/vpp-renderer:config

{
  "config": {
    "bridge-domain": [
      {
        "id": "43282482-a677-4102-87d6-90708f30a115",
        "type": "vpp-renderer:vlan-network",
        "description": "drexternal",
        "vlan": 2016,
        "physical-location-ref": [
          {
            "node-id": "dut2",
            "interface": [
              "GigabitEthernet0/9/0"
            ]
          }
        ]
      }
    ]
  }
}
```

Port (compute):

```
PUT http://{{controller}}:{{port}}/restconf/config/neutron:neutron/ports

{
```

(continues on next page)

(continued from previous page)

```
    "ports": {
        "port": [
            {
                "uuid": "3d5dff96-25f5-4d4b-aa11-dc03f7f8d8e0",
                "tenant-id": "94836b88-0e56-4150-aaa7-60f1c2b67faa",
                "device-id": "dhcp58155ae3-f2e7-51ca-9978-71c513ab02ee-a91437c0-8492-
↪47e2-b9d0-25c44aef6cda",
                "neutron-binding:vif-details": [
                    {
                        "details-key": "somekey"
                    }
                ],
                "neutron-binding:host-id": "devstack-control",
                "neutron-binding:vif-type": "vhostuser",
                "neutron-binding:vnic-type": "normal",
                "mac-address": "fa:16:3e:4a:9f:c0",
                "name": "",
                "network-id": "a91437c0-8492-47e2-b9d0-25c44aef6cda",
                "neutron-portsecurity:port-security-enabled": false,
                "device-owner": "network:compute",
                "fixed-ips": [
                    {
                        "subnet-id": "0a5834ed-ed31-4425-832d-e273cac26325",
                        "ip-address": "10.1.1.3"
                    }
                ],
                "admin-state-up": true
            }
        ]
    }
}

GET http://{{controller}}:{{port}}/restconf/config/vpp-renderer:config

{
  "config": {
    "vpp-endpoint": [
      {
        "context-type": "l2-l3-forwarding:l2-bridge-domain",
        "context-id": "a91437c0-8492-47e2-b9d0-25c44aef6cda",
        "address-type": "l2-l3-forwarding:mac-address-type",
        "address": "fa:16:3e:4a:9f:c0",
        "vpp-node-path": "/network-topology:network-topology/network-
↪topology:topology[network-topology:topology-id='topology-netconf']/network-
↪topology:node[network-topology:node-id='devstack-control']",
        "vpp-interface-name": "neutron_port_3d5dff96-25f5-4d4b-aa11-dc03f7f8d8e0",
        "socket": "/tmp/socket_3d5dff96-25f5-4d4b-aa11-dc03f7f8d8e0",
        "description": "neutron port"
      }
    ]
  }
}
```

Port (dhcp):

```
PUT http://{{controller}}:{{port}}/restconf/config/neutron:neutron/ports
```

(continues on next page)

(continued from previous page)

```
{
    "ports": {
        "port": [
            {
                "uuid": "3d5dff96-25f5-4d4b-aa11-dc03f7f8d8e0",
                "tenant-id": "94836b88-0e56-4150-aaa7-60f1c2b67faa",
                "device-id": "dhcp58155ae3-f2e7-51ca-9978-71c513ab02ee-a91437c0-8492-
↪47e2-b9d0-25c44aef6cda",
                "neutron-binding:vif-details": [
                    {
                        "details-key": "somekey"
                    }
                ],
                "neutron-binding:host-id": "devstack-control",
                "neutron-binding:vif-type": "vhostuser",
                "neutron-binding:vnic-type": "normal",
                "mac-address": "fa:16:3e:4a:9f:c0",
                "name": "",
                "network-id": "a91437c0-8492-47e2-b9d0-25c44aef6cda",
                "neutron-portsecurity:port-security-enabled": false,
                "device-owner": "network:dhcp",
                "fixed-ips": [
                    {
                        "subnet-id": "0a5834ed-ed31-4425-832d-e273cac26325",
                        "ip-address": "10.1.1.3"
                    }
                ],
                "admin-state-up": true
            }
        ]
    }
}

GET http://{{controller}}:{{port}}/restconf/config/vpp-renderer:config

{
  "config": {
    "vpp-endpoint": [
      {
        "context-type": "l2-l3-forwarding:l2-bridge-domain",
        "context-id": "a91437c0-8492-47e2-b9d0-25c44aef6cda",
        "address-type": "l2-l3-forwarding:mac-address-type",
        "address": "fa:16:3e:4a:9f:c0",
        "vpp-node-path": "/network-topology:network-topology/network-
↪topology:topology[network-topology:topology-id='topology-netconf']/network-
↪topology:node[network-topology:node-id='devstack-control']",
        "vpp-interface-name": "neutron_port_3d5dff96-25f5-4d4b-aa11-dc03f7f8d8e0",
        "physical-address": "fa:16:3e:4a:9f:c0",
        "name": "tap3d5dff96-25",
        "description": "neutron port"
      }
    ]
  }
}
```

OpenStack with Virtual Tenant Network

This section describes using OpenDaylight with the VTN manager feature providing network service for OpenStack. VTN manager utilizes the OVSDB southbound service and Neutron for this implementation. The below diagram depicts the communication of OpenDaylight and two virtual networks connected by an OpenFlow switch using this implementation.

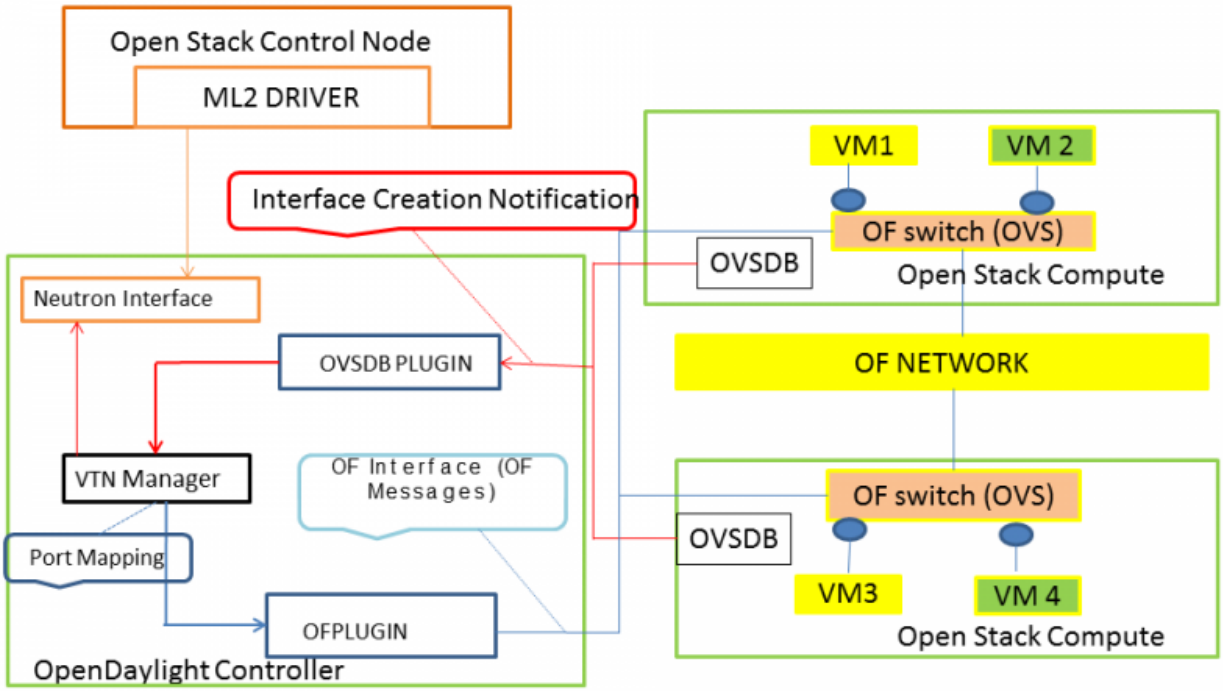

Fig. 155: OpenStack Architecture

Configure OpenStack to work with OpenDaylight(VTN Feature) using PackStack

Prerequisites to install OpenStack using PackStack

- Fresh CentOS 7.1 minimal install
- Use the below commands to disable and remove Network Manager in CentOS 7.1,

```
systemctl stop NetworkManager
systemctl disable NetworkManager
```

- To make SELINUX as permissive, please open the file "/etc/sysconfig/selinux" and change it as "SELINUX=permissive".

- After making selinux as permissive, please restart the CentOS 7.1 machine.

Steps to install OpenStack PackStack in CentOS 7.1

- To install OpenStack juno, use the following command,

```
yum update -y
yum -y install https://repos.fedorapeople.org/repos/openstack/openstack-juno/rdo-
↪release-juno-1.noarch.rpm
```

- To install the packstack installer, please use the below command,

```
yum -y install openstack-packstack
```

- To create all-in-one setup, please use the below command,

```
packstack --allinone --provision-demo=n --provision-all-in-one-ovs-bridge=n
```

- This will end up with Horizon started successfully message.

Steps to install and deploy OpenDaylight in CentOS 7.1

- Download the latest Boron distribution code in the below link,

```
wget https://nexus.opendaylight.org/content/groups/public/org/opendaylight/
↪integration/distribution-karaf/0.5.0-Boron/distribution-karaf-0.5.0-Boron.zip
```

- Unzip the Boron distribution code by using the below command,

```
unzip distribution-karaf-0.5.0-Boron.zip
```

- Please do the below steps in the OpenDaylight to change jetty port,
  - Change the jetty port from 8080 to something else as swift proxy of OpenStack is using it.
  - Open the file "etc/jetty.xml" and change the jetty port from 8080 to 8910 (we have used 8910 as jetty port you can use any other number).
  - Start VTN Manager and install odl-vtn-manager-neutron in it.
  - Ensure all the required ports(6633/6653,6640 and 8910) are in the listen mode by using the command "netstat -tunpl" in OpenDaylight.

Steps to reconfigure OpenStack in CentOS 7.1

- Steps to stop Open vSwitch Agent and clean up ovs

```
sudo systemctl stop neutron-openvswitch-agent
sudo systemctl disable neutron-openvswitch-agent
sudo systemctl stop openvswitch
sudo rm -rf /var/log/openvswitch/*
sudo rm -rf /etc/openvswitch/conf.db
sudo systemctl start openvswitch
sudo ovs-vsctl show
```

- Stop Neutron Server

```
systemctl stop neutron-server
```

- Verify that OpenDaylight's ML2 interface is working:

```
curl -v admin:admin http://{CONTROL_HOST}:{PORT}/controller/nb/v2/neutron/networks

{
   "networks" : [ ]
}
```

If this does not work or gives an error, check Neutron's log file in */var/log/neutron/server.log*. Error messages here should give some clue as to what the problem is in the connection with OpenDaylight

- Configure Neutron to use OpenDaylight's ML2 driver:

```
sudo crudini --set /etc/neutron/plugins/ml2/ml2_conf.ini ml2 mechanism_drivers␣
↪opendaylight
sudo crudini --set /etc/neutron/plugins/ml2/ml2_conf.ini ml2 tenant_network_types␣
↪local
sudo crudini --set /etc/neutron/plugins/ml2/ml2_conf.ini ml2 type_drivers local
sudo crudini --set /etc/neutron/dhcp_agent.ini DEFAULT ovs_use_veth True

cat <<EOT | sudo tee -a /etc/neutron/plugins/ml2/ml2_conf.ini > /dev/null
  [ml2_odl]
  password = admin
  username = admin
  url = http://{CONTROL_HOST}:{PORT}/controller/nb/v2/neutron
  EOT
```

- Reset Neutron's ML2 database

```
sudo mysql -e "drop database if exists neutron_ml2;"
sudo mysql -e "create database neutron_ml2 character set utf8;"
sudo mysql -e "grant all on neutron_ml2.* to 'neutron'@'%';"
sudo neutron-db-manage --config-file /usr/share/neutron/neutron-dist.conf --config-
↪file /etc/neutron/neutron.conf --config-file /etc/neutron/plugin.ini upgrade head
```

- Start Neutron Server

```
sudo systemctl start neutron-server
```

- Restart the Neutron DHCP service

```
system restart neutron-dhcp-agent.service
```

- At this stage, your Open vSwitch configuration should be empty:

```
[root@dneary-odl-compute2 ~]# ovs-vsctl show
686989e8-7113-4991-a066-1431e7277e1f
    ovs_version: "2.3.1"
```

- Set OpenDaylight as the manager on all nodes

```
ovs-vsctl set-manager  tcp:127.0.0.1:6640
```

- You should now see a section in your Open vSwitch configuration showing that you are connected to the OpenDaylight server, and OpenDaylight will automatically create a br-int bridge:

```
[root@dneary-odl-compute2 ~]# ovs-vsctl show
686989e8-7113-4991-a066-1431e7277e1f
    Manager "tcp:127.0.0.1:6640"
        is_connected: true
    Bridge br-int
        Controller "tcp:127.0.0.1:6633"
            is_connected: true
        fail_mode: secure
        Port "ens33"
            Interface "ens33"
    ovs_version: "2.3.1"
```

- Add the default flow to OVS to forward packets to controller when there is a table-miss,

```
ovs-ofctl --protocols=OpenFlow13 add-flow br-int priority=0,actions=output:CONTROLLER
```

- Please see the VTN OpenStack PackStack support guide on the wiki to create VM's from OpenStack Horizon GUI.

Implementation details

VTN Manager

Install **odl-vtn-manager-neutron** feature which provides the integration with Neutron interface.

```
feature:install odl-vtn-manager-neutron
```

It subscribes to the events from Open vSwitch and also implements the Neutron requests received by OpenDaylight.

Functional Behavior

StartUp

- The ML2 implementation for OpenDaylight will ensure that when Open vSwitch is started, the ODL_IP_ADDRESS configured will be set as manager.
- When OpenDaylight receives the update of the Open vSwitch on port 6640 (manager port), VTN Manager handles the event and adds a bridge with required port mappings to the Open vSwitch at the OpenStack node.
- When Neutron starts up, a new network create is POSTed to OpenDaylight, for which VTN Manager creates a Virtual Tenant Network.
- *Network and Sub-Network Create:* Whenever a new sub network is created, VTN Manager will handle the same and create a vbridge under the VTN.
- *VM Creation in OpenStack:* The interface mentioned as integration bridge in the configuration file will be added with more interfaces on creation of a new VM in OpenStack and the network is provisioned for it by the VTN Neutron feature. The addition of a new port is captured by the VTN Manager and it creates a vbridge interface with port mapping for the particular port. When the VM starts to communicate with other VMs, the VTN Manger will install flows in the Open vSwitch and other OpenFlow switches to facilitate communication between them.

Note: To use this feature, VTN feature should be installed

Reference

https://wiki.opendaylight.org/images/5/5c/Integration_of_vtn_and_ovsdb_for_helium.pdf

CHAPTER 2

Content for OpenDaylight Developers

The Following content is intended for developers building applications or code on top of OpenDaylight, but who do not plan to modify OpenDaylight code itself.

2.1 Developer Guide

2.1.1 Overview

Getting started with Git and Gerrit

Overview of Git and Gerrit

Git is an opensource distributed version control system (dvcs) written in the C language and originally developed by Linus Torvalds and others to manage the Linux kernel. In Git, there is no central copy of the repository. After you have cloned the repository, you have a functioning copy of the source code with all the branches and tagged releases, in your local repository.

Gerrit is an opensource web-based collaborative code review tool that integrates with Git. It was developed at Google by Shawn Pearce. Gerrit provides a framework for reviewing code commits before they are accepted into the code base. Changes can be uploaded to Gerrit by any user. However, the changes are not made a part of the project until a code review is completed. Gerrit is also a good collaboration tool for storing the conversations that occur around the code commits.

The OpenDaylight source code is hosted in a repository in Git. Developers must use Gerrit to commit code to the OpenDaylight repository.

Note: For more information on Git, see http://git-scm.com/. For more information on Gerrit, see https://code.google.com/p/gerrit/.

Setting up a Gerrit account

1. Using a Google Chrome or Mozilla Firefox browser, go to https://git.opendaylight.org/gerrit

The main page shows existing Gerrit requests. These are patches that have been pushed to the repository and not yet verified, reviewed, and merged.

Note: If you already have an OpenDaylight account, you can click **Sign In** in the top right corner of the page and follow the instructions to enter the OpenDaylight page.

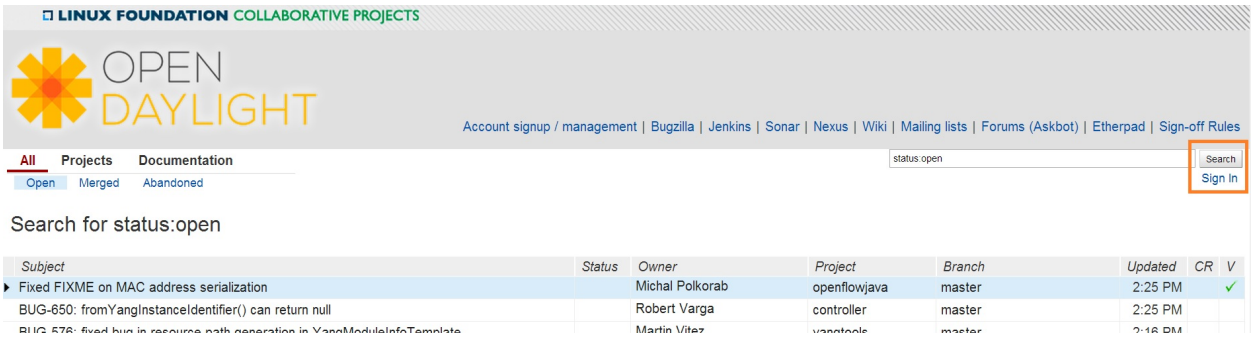

Fig. 1: Signing in to OpenDaylight account

1. If you do not have an existing OpenDaylight account, click **Account signup/management** on the top bar of the main Gerrit page.

The **WS02 Identity Server** page is displayed.

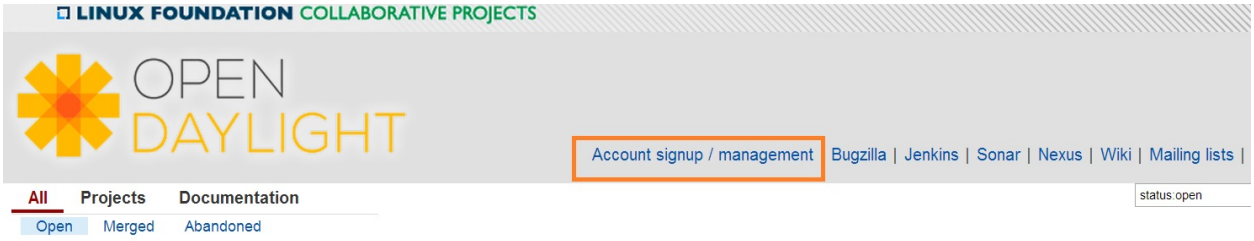

Fig. 2: Gerrit Account signup/management link

1. In the **WS02 Identity Server** page, click **Sign-up** in the left pane.

There is also an option to authenticate your sign in with OpenID. This option is not described in this document.

1. Click on the **Sign-up with User Name/Password** image on the right pane to continue to the actual sign-up page.

1. Fill out the details in the account creation form and then click **Submit**.

You now have an OpenDaylight account that can be used with Gerrit to pull the OpenDaylight code.

Generating SSH keys for your system

You must have SSH keys for your system to register with your Gerrit account. The method for generating SSH keys is different for different types of operating systems.

Fig. 3: Sign-up link for Gerrit account

Fig. 4: Sign-up with User Name/Password Image

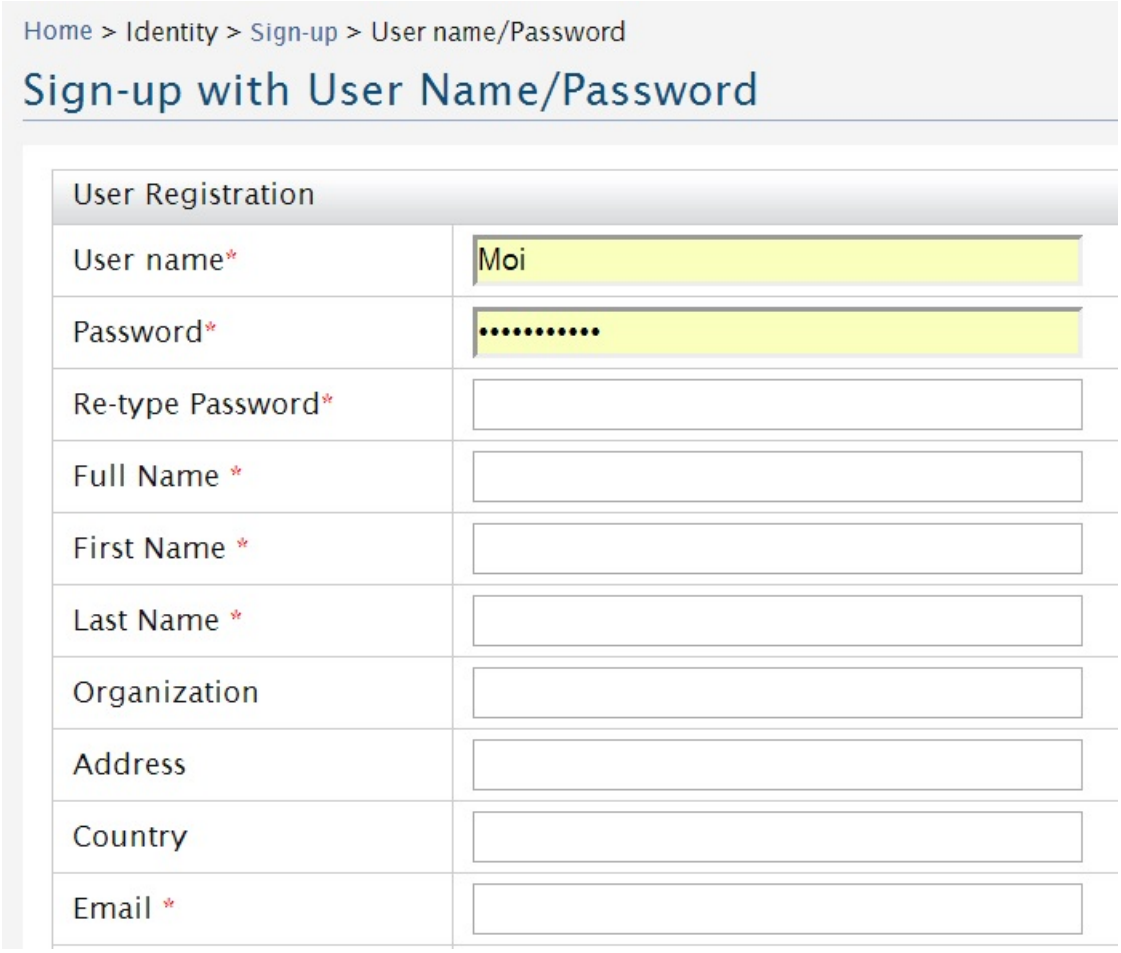

Fig. 5: Filling out the details

The key you register with Gerrit must be identical to the one you will use later to pull or edit the code. For example, if you have a development VM which has a different UID login and keygen than that of your laptop, the SSH key you generate for the VM is different from the laptop. If you register the SSH key generated on your VM with Gerrit and do not reuse it on your laptop when using Git on the laptop, the pull fails.

Note: For more information on SSH keys for Ubuntu, see https://help.ubuntu.com/community/SSH/OpenSSH/Keys. For generating SSH keys for Windows, see https://help.github.com/articles/generating-ssh-keys.

For a system running Ubuntu operating system, follow the steps below:

1. Run the following command:

```
mkdir ~/.ssh
chmod 700 ~/.ssh
ssh-keygen -t rsa
```

1. You are prompted for a location to save the keys, and a passphrase for the keys.

This passphrase protects your private key while it is stored on the hard drive. You must use the passphrase to use the keys every time you need to login to a key-based system:

```
Generating public/private rsa key pair.
Enter file in which to save the key (/home/b/.ssh/id_rsa):
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
Your identification has been saved in /home/b/.ssh/id_rsa.
Your public key has been saved in /home/b/.ssh/id_rsa.pub.
```

Your public key is now available as **.ssh/id_rsa.pub** in your home folder.

Registering your SSH key with Gerrit

1. Using a Google Chrome or Mozilla Firefox browser, go to https://git.opendaylight.org/gerrit.

1. Click **Sign In** to access the OpenDaylight repository.

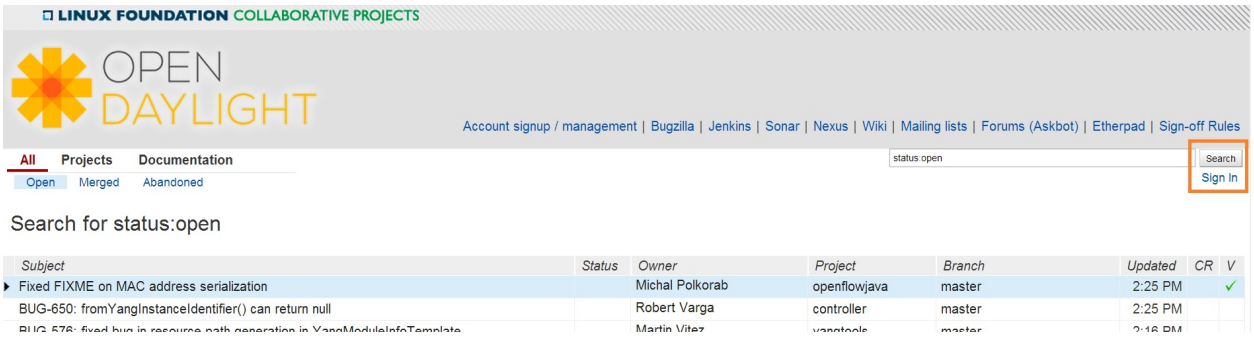

Fig. 6: Signin in to OpenDaylight repository

1. Click your name in the top right corner of the window and then click **Settings**.

The **Settings** page is displayed.

1. Click **SSH Public Keys** under **Settings**.
2. Click **Add Key**.

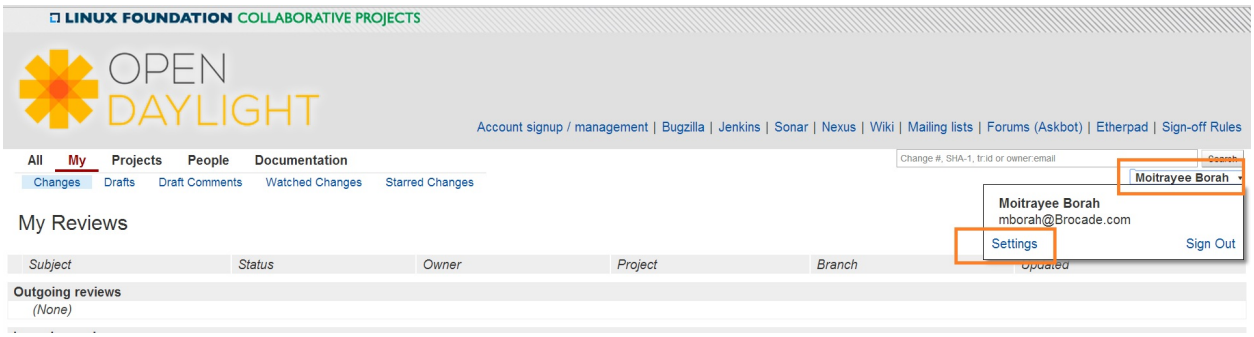

Fig. 7: Settings page for your Gerrit account

3. In the **Add SSH Public Key** text box, paste the contents of your **id_rsa.pub** file and then click **Add**.

Fig. 8: Adding your SSH key

To verify your SSH key is working correctly, try using an SSH client to connect to Gerrit's SSHD port:

```
$ ssh -p 29418 <sshusername>@git.opendaylight.org
Enter passphrase for key '/home/cisco/.ssh/id_rsa':
****    Welcome to Gerrit Code Review    ****
Hi <user>, you have successfully connected over SSH.
Unfortunately, interactive shells are disabled.
To clone a hosted Git repository, use: git clone ssh://<user>@git.opendaylight.
↪org:29418/REPOSITORY_NAME.git
Connection to git.opendaylight.org closed.
```

You can now proceed to either Pulling, Hacking, and Pushing the Code from the CLI or Pulling, Hacking, and Pushing the Code from Eclipse depending on your implementation.

Pulling and Pushing the Code from the CLI

OpenDayligh is a collection of projects, each with their own code repository. This section provides a general guide for to pulling, hacking, and pushing the code for each project. For project specific detail, refer to the project's section in

this guide.

Code reviews are enabled through Gerrit. For setting up Gerrit see the section on Getting started with Git and Gerrit.

Note: You will need to perform the Gerrit Setup before you can access git via ssh as described below.

Pulling code via Git CLI

Pull the code by cloning the project's repository.

```
git clone ssh://<username>@git.opendaylight.org:29418/<project_repo_name>.git
```

where <username> is your OpenDaylight username, and <project_repo_name> is the name of the repository for project you are trying to pull. Here is the current list of project repository names:

aaa, affinity, bgpcep, controller, defense4all, dlux, docs, groupbasedpolicy, integration, l2switch, lispflowmapping, odlparent, opendove, openflowjava, openflowplugin, opflex, ovsdb, packetcable, reservation, sdninterfaceapp, sfc, snbi, snmp4sdn, toolkit, ttp, vtn, yangtools.

For an anonymous git clone, you can use:

```
git clone https://git.opendaylight.org/gerrit/p/<project_repo_name>.git
```

Setting up Gerrit Change-id Commit Message Hook

- This command inserts a unique Change-Id tag in the footer of a commit message. This step is optional but highly recommended for tracking changes.

```
cd <project_repo_name>
scp -p -P 29418 <username>@git.opendaylight.org:hooks/commit-msg .git/hooks/
chmod 755 .git/hooks/commit-msg
```

- Install and setup Git-review. Git-review is a great tool to simplify the hassle of using several git commands to submit a patch for review. Refer to How to install and push codes with git-review for instructions. After initializing git-review, both commit-msg hook and a remote repo named gerrit will be created and a patch can be submitted to Gerrit with a single "git review" command.
- Now you can start making your code changes.

Building the code

While you are in the <project_repo_name> directory, run

```
mvn clean install
```

To run without unitests you can skip building those tests running the following:

```
mvn clean install -DskipTests
/* instead of "mvn clean install" */
```

Runing OpenDaylight from local build

Change to the karaf distribution sub-directory, and run

```
./target/assembly/bin/karaf
```

At this point the OpenDaylight controller is running. You can now open a web browser and point your browser at http://localhost:8080/

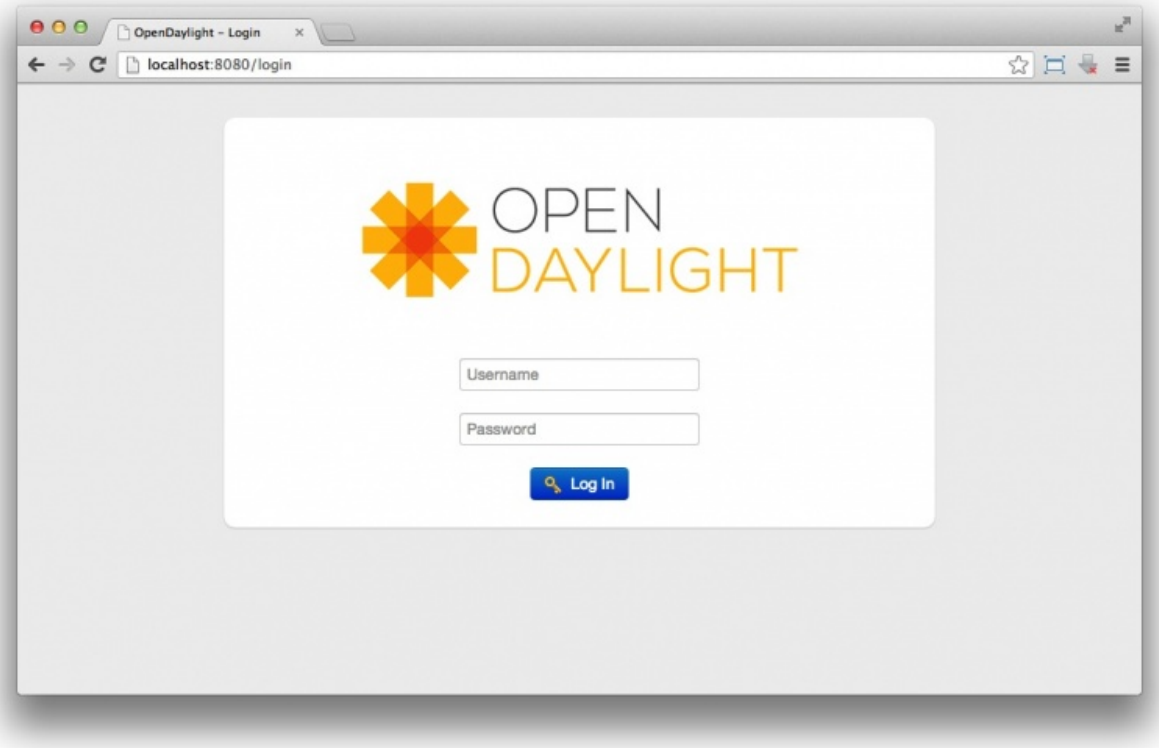

Fig. 9: OpenDaylight Main Page

Commit the code using Git CLI

Note: To be accepted, all code mustcome with a developer certificate of origin as expressed by having a Signed-off-by. This means that you are asserting that you have made the change and you understand that the work was done as part of an open-source license.

```
Developer's Certificate of Origin 1.1

        By making a contribution to this project, I certify that:

        (a) The contribution was created in whole or in part by me and I
            have the right to submit it under the open source license
            indicated in the file; or
```

(continues on next page)

(continued from previous page)

```
        (b) The contribution is based upon previous work that, to the best
            of my knowledge, is covered under an appropriate open source
            license and I have the right under that license to submit that
            work with modifications, whether created in whole or in part
            by me, under the same open source license (unless I am
            permitted to submit under a different license), as indicated
            in the file; or

        (c) The contribution was provided directly to me by some other
            person who certified (a), (b) or (c) and I have not modified
            it.

        (d) I understand and agree that this project and the contribution
            are public and that a record of the contribution (including all
            personal information I submit with it, including my sign-off) is
            maintained indefinitely and may be redistributed consistent with
            this project or the open source license(s) involved.
```

Mechanically you do it this way:

```
git commit --signoff
```

You will be prompted for a commit message. If you are fixing a buzilla bug you can add the associated bug number to your commit message and it will get linked from Gerrit:

For Example:.

```
Fix for bug 2.

Signed-off-by: Ed Warnicke <eaw@cisco.com>
# Please enter the commit message for your changes. Lines starting
# with '#' will be ignored, and an empty message aborts the commit.
# On branch develop
# Changes to be committed:
#   (use "git reset HEAD <file>..." to unstage)
#
#       modified:   README
#
```

Pulling the Code changes via Git CLI

Pull the latest changes from the remote repository

```
git remote update
git rebase origin/<project_main_branch_name>
```

where <project_main_branch_name> is the the branch you want to commit to. For most projects this is master branch. For some projects such as lispflowmapping, a different branch name (develop in the case of lispflowmapping) should be used.

Pushing the Code via Git CLI

Use git review to push your changes back to the remote repository using:

```
git review
```

You can set a topic for your patch by:

```
git review -t <topic>
```

You will get a message pointing you to your gerrit request like:

```
=========================
remote: Resolving deltas: 100% (2/2) +
remote: Processing changes: new: 1, refs: 1, done    +
remote: +
remote: New Changes: +
remote:   http://git.opendaylight.org/gerrit/64 +
remote: +
=========================
```

The Jenkins Controller User will verify your code and post the result on the your gerrit request.

Viewing your Changes in Gerrit

Follow the link you got above to see your commit in Gerrit:

Note that the Jenkins Controller User has verified your code and at the bottom is a link to the Jenkins build.

Once your code has been reviewed and submitted by a committer it will be merged into the authoritative repo, which would look like this:

Troubleshooting

1. **What to do if your Firewall blocks port 29418**

There have been reports that many corporate firewalls block port 29418. If that's the case, please follow the Setting up HTTP in Gerrit instructions and use git URL:

```
git clone https://<your_username>@git.opendaylight.org/gerrit/p/<project_repo_name>.
↪git
```

You will be prompted for the password you generated in Setting up HTTP in Gerrit.

All other instructions on this page remain unchanged.

To download pre-built images with ODP bootstraps see the following Github project:

Pre-Built OpenDaylight VM Images

Developing Apps on the OpenDaylight controller

This section provides information that is required to develop apps on the OpenDaylight controller.

You can either develop apps within the controller using the model-driven SAL (MD-SAL) archetype or develop external apps and use the RESTCONF to communicate with the controller.

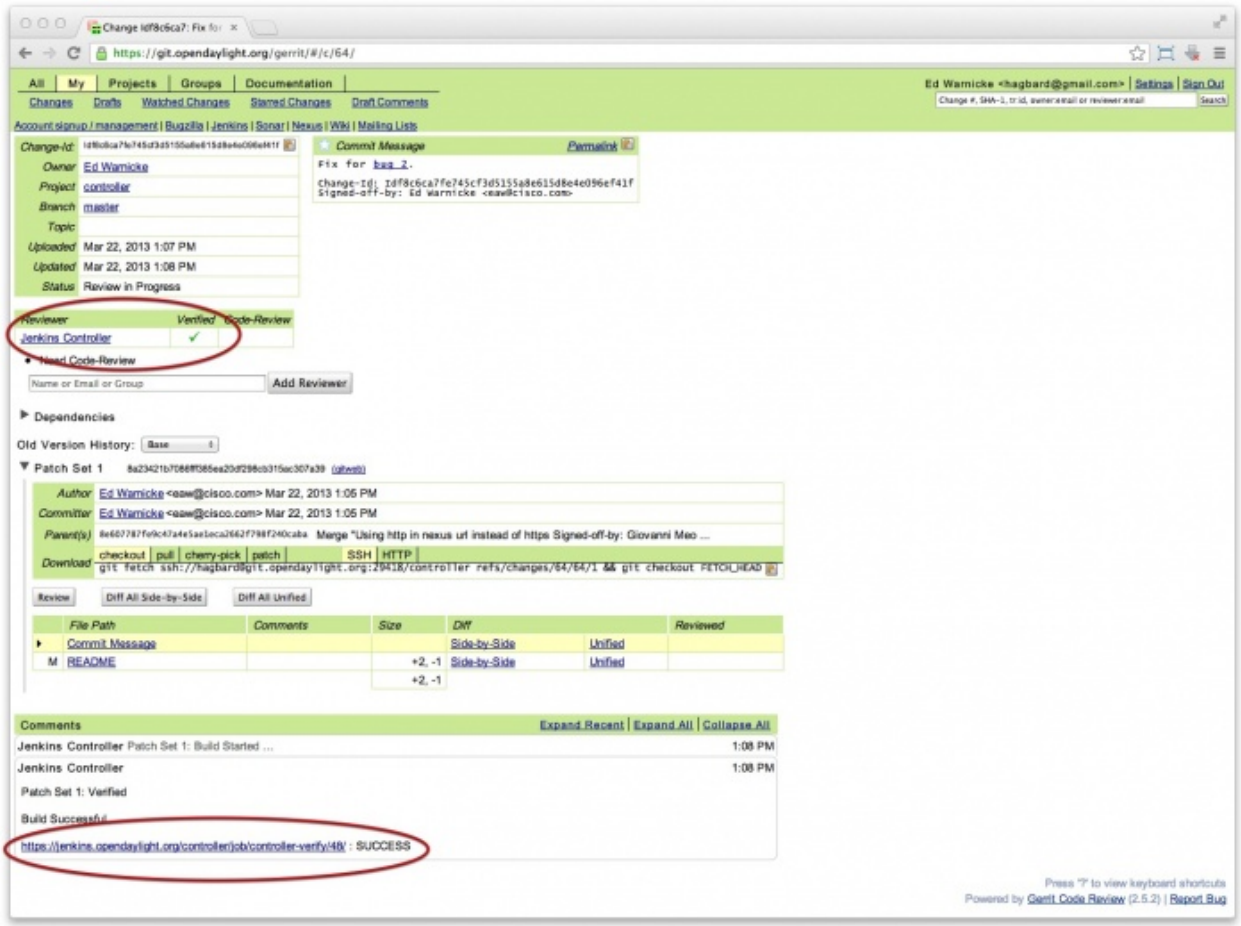

Fig. 10: Gerritt Code Review Sample

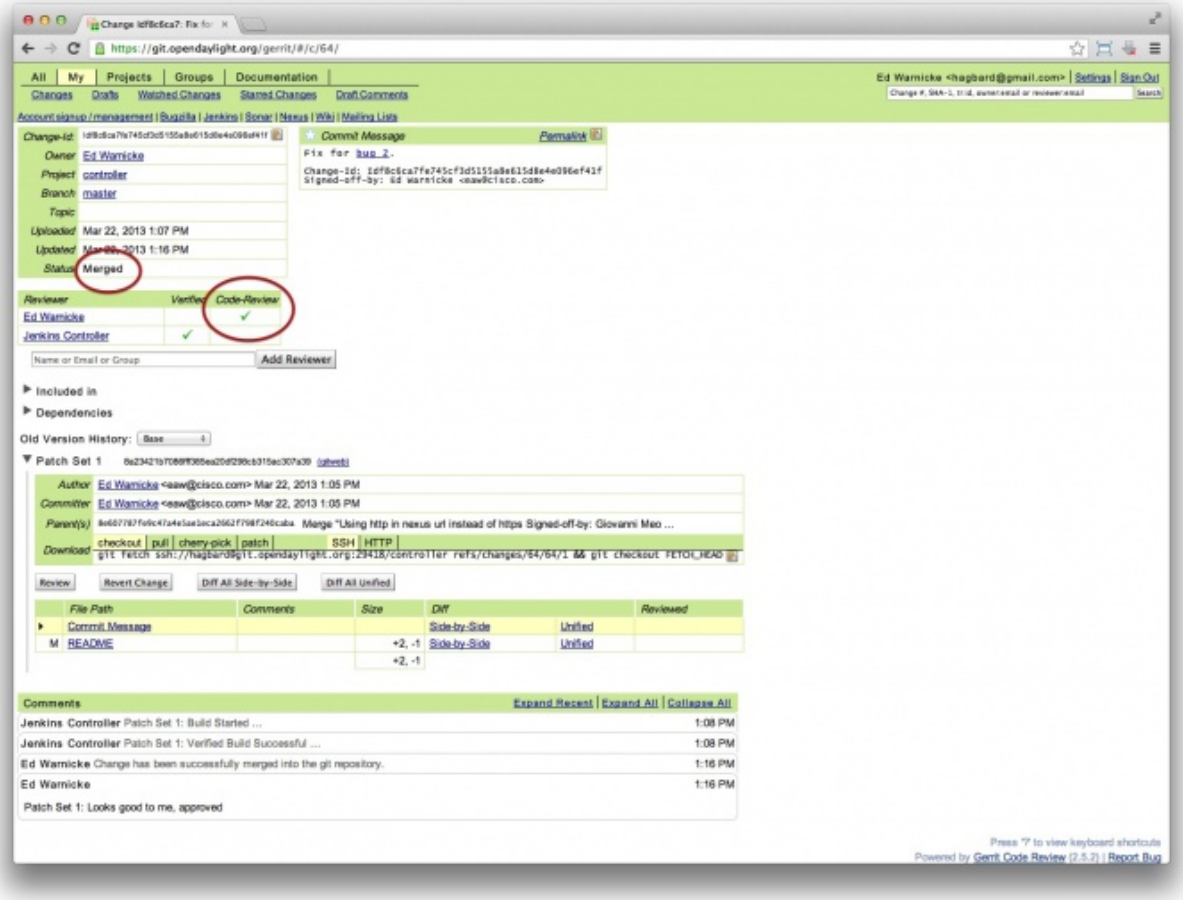

Fig. 11: Gerritt Code Merge Sample

Overview

This section enables you to get started with app development within the OpenDaylight controller. In this example, you perform the following steps to develop an app.

1. Create a local repository for the code using a simple build process.
2. Start the OpenDaylight controller.
3. Test a simple remote procedure call (RPC) which you have created based on the principle of *hello world*.

Pre requisites

This example requires the following.

- A development environment with following set up and working correctly from the shell:
  - Maven 3.1.1 or later
  - Java 7- or Java 8-compliant JDK
  - An appropriate Maven settings.xml file. A simple way to get the default OpenDaylight settings.xml file is:

```
cp -n ~/.m2/settings.xml{,.orig} ; \wget -q -O - https://raw.
↪githubusercontent.com/opendaylight/odlparent/stable/boron/settings.xml > ~/.
↪m2/settings.xml
```

Note: If you are using Linux or Mac OS X as your development OS, your local repository is ~/.m2/repository. For other platforms the local repository location will vary.

Building an example module

To develop an app perform the following steps.

1. Create an *Example* project using Maven and an archetype called the *opendaylight-startup-archetype*. If you are downloading this project for the first time, then it will take sometime to pull all the code from the remote repository.

```
mvn archetype:generate -DarchetypeGroupId=org.opendaylight.controller -
↪DarchetypeArtifactId=opendaylight-startup-archetype \
-DarchetypeRepository=https://nexus.opendaylight.org/content/repositories/public/␣
↪\
-DarchetypeCatalog=https://nexus.opendaylight.org/content/repositories/public/
↪archetype-catalog.xml
```

2. Update the properties values as follows. Ensure that the groupid and the artifactid is lower case.

```
Define value for property 'groupId': : org.opendaylight.example
Define value for property 'artifactId': : example
Define value for property 'version':  1.0-SNAPSHOT: : 1.0.0-SNAPSHOT
Define value for property 'package':  org.opendaylight.example: :
Define value for property 'classPrefix':  ${artifactId.substring(0,1).
↪toUpperCase()}${artifactId.substring(1)}
Define value for property 'copyright': : Copyright (c) 2015 Yoyodyne, Inc.
```

3. Accept the default value of classPrefix that is, `(${artifactId.substring(0,1).toUpperCase()}${artifactId.substring(1)})`. The classPrefix creates a Java Class Prefix by capitalizing the first character of the artifactId.

   Note: In this scenario, the classPrefix used is "Example". Create a top-level directory for the archetype.

```
${artifactId}/
example/
cd example/
api/
artifacts/
features/
impl/
karaf/
pom.xml
```

4. Build the *example* project.

   Note: Depending on your development machine's specification this might take a little while. Ensure that you are in the project's root directory, example/, and then issue the build command, shown below.

```
mvn clean install
```

5. Start the *example* project for the first time.

```
cd karaf/target/assembly/bin
ls
./karaf
```

6. Wait for the karaf cli that appears as follows. Wait for OpenDaylight to fully load all the components. This can take a minute or two after the prompt appears. Check the CPU on your dev machine, specifically the Java process to see when it calms down.

```
opendaylight-user@root>
```

7. Verify if the "example" module is built and search for the log entry which includes the entry *ExampleProvider Session Initiated*.

```
log:display | grep Example
```

8. Shutdown the OpenDaylight through the console by using the following command.

```
shutdown -f
```

Defining a Simple Hello World RPC

1. Run the maven archetype *opendaylight-startup-archetype*, and create the *hello* project.

```
mvn archetype:generate -DarchetypeGroupId=org.opendaylight.controller -
↪DarchetypeArtifactId=opendaylight-startup-archetype \
-DarchetypeRepository=http://nexus.opendaylight.org/content/repositories/
↪opendaylight.snapshot/ \
```

(continues on next page)

(continued from previous page)

```
-DarchetypeCatalog=http://nexus.opendaylight.org/content/repositories/
↪opendaylight.snapshot/archetype-catalog.xml
```

2. Update the Properties values as follows.

```
Define value for property 'groupId': : org.opendaylight.hello
Define value for property 'artifactId': : hello
Define value for property 'version':  1.0-SNAPSHOT: : 1.0.0-SNAPSHOT
Define value for property 'package':  org.opendaylight.hello: :
Define value for property 'classPrefix':  ${artifactId.substring(0,1).
↪toUpperCase()}${artifactId.substring(1)}
Define value for property 'copyright': : Copyright(c) Yoyodyne, Inc.
```

3. View the *hello* project.

```
cd hello/
ls -1
api
artifacts
features
impl
karaf
pom.xml
```

4. Build *hello* project by using the following command.

```
mvn clean install
```

5. Verify that the project is functioning by executing karaf.

```
cd karaf/target/assembly/bin
./karaf
```

6. The karaf cli appears as follows.
   NOTE: Remember to wait for OpenDaylight to load completely. Verify that the Java process CPU has stabilized.+

```
opendaylight-user@root>
```

7. Verify that the *hello* module is loaded by checking the log.

```
log:display | grep Hello
```

8. Shutdown karaf.

```
shutdown -f
```

9. Return to the top of the directory structure:

```
cd ../../../../
```

10. View the entry point to understand where the log line came from. The entry point is in the impl project:

```
impl/src/main/java/org/opendaylight/hello/impl/HelloProvider.java
```

11. Add any new things that you are doing in your implementation by using the HelloProvider.onSessionInitiate method. Its analogous to an Activator.

```
@Override
    public void onSessionInitiated(ProviderContext session) {
        LOG.info("HelloProvider Session Initiated");
    }
```

Add a simple HelloWorld RPC API

1. Navigate to the file.

```
Edit
api/src/main/yang/hello.yang
```

2. Edit this file as follows. In the following example, we are adding the code in a YANG module to define the *hello-world* RPC:

3. Return to the hello/api directory and build your API as follows.

```
cd ../../../
mvn clean install
```

Implement the HelloWorld RPC API

1. Define the HelloService, which is invoked through the *hello-world* API.

```
cd ../impl/src/main/java/org/opendaylight/hello/impl/
```

2. Create a new file called HelloWorldImpl.java and add in the code below.

```
package org.opendaylight.hello.impl;
import java.util.concurrent.Future;
import org.opendaylight.yang.gen.v1.urn.opendaylight.params.xml.ns.yang.hello.
↪rev150105.HelloService;
import org.opendaylight.yang.gen.v1.urn.opendaylight.params.xml.ns.yang.hello.
↪rev150105.HelloWorldInput;
import org.opendaylight.yang.gen.v1.urn.opendaylight.params.xml.ns.yang.hello.
↪rev150105.HelloWorldOutput;
import org.opendaylight.yang.gen.v1.urn.opendaylight.params.xml.ns.yang.hello.
↪rev150105.HelloWorldOutputBuilder;
import org.opendaylight.yangtools.yang.common.RpcResult;
import org.opendaylight.yangtools.yang.common.RpcResultBuilder;
public class HelloWorldImpl implements HelloService {
    @Override
    public Future<RpcResult<HelloWorldOutput>> helloWorld(HelloWorldInput input) {
        HelloWorldOutputBuilder helloBuilder = new HelloWorldOutputBuilder();
        helloBuilder.setGreating("Hello " + input.getName());
        return RpcResultBuilder.success(helloBuilder.build()).buildFuture();
    }
}
```

3. The HelloProvider.java file is in the current directory. Register the RPC that you created in the *hello.yang* file in the HelloProvider.java file. You can either edit the HelloProvider.java to match what is below or you can simple replace it with the code below.

```
/*
 * Copyright(c) Yoyodyne, Inc. and others.  All rights reserved.
 *
 * This program and the accompanying materials are made available under the
 * terms of the Eclipse Public License v1.0 which accompanies this distribution,
 * and is available at http://www.eclipse.org/legal/epl-v10.html
 */
package org.opendaylight.hello.impl;

import org.opendaylight.controller.sal.binding.api.BindingAwareBroker.
↪ProviderContext;
import org.opendaylight.controller.sal.binding.api.BindingAwareBroker.
↪RpcRegistration;
import org.opendaylight.controller.sal.binding.api.BindingAwareProvider;
import org.opendaylight.yang.gen.v1.urn.opendaylight.params.xml.ns.yang.hello.
↪rev150105.HelloService;
import org.slf4j.Logger;
import org.slf4j.LoggerFactory;

public class HelloProvider implements BindingAwareProvider, AutoCloseable {
    private static final Logger LOG = LoggerFactory.getLogger(HelloProvider.
↪class);
    private RpcRegistration<HelloService> helloService;
    @Override
    public void onSessionInitiated(ProviderContext session) {
        LOG.info("HelloProvider Session Initiated");
        helloService = session.addRpcImplementation(HelloService.class, new␣
↪HelloWorldImpl());
    }
    @Override
    public void close() throws Exception {
        LOG.info("HelloProvider Closed");
        if (helloService != null) {
            helloService.close();
        }
    }
}
```

4. Optionally, you can also build the Java classes which will register the new RPC. This is useful to test the edits you have made to HelloProvider.java and HelloWorldImpl.java.

```
cd ../../../../../../../
mvn clean install
```

5. Return to the top level directory

```
cd ../
```

6. Build the entire *hello* again, which will pickup the changes you have made and build them into your project:

```
mvn clean install
```

Execute the *hello* project for the first time

1. Run karaf

```
cd ../karaf/target/assembly/bin
./karaf
```

2. Wait for the project to load completely. Then view the log to see the loaded *Hello* Module:

```
log:display | grep Hello
```

Test the *hello-world* RPC via REST

There are a lot of ways to test your RPC. Following are some examples.

1. Using the API Explorer through HTTP
2. Using a browser REST client

Using the API Explorer through HTTP

1. Navigate to apidoc UI with your web browser.
   NOTE: In the URL mentioned above, Change *localhost* to the IP/Host name to reflect your development machine's network address.
2. Select

```
hello(2015-01-05)
```

3. Select

```
POST /operations/hello:hello-world
```

4. Provide the required value.

```
{"hello:input": { "name":"Your Name"}}
```

5. Click the button.
6. Enter the username and password, by default the credentials are admin/admin.
7. In the response body you should see.

```
{
  "output": {
    "greating": "Hello Your Name"
  }
}
```

Using a browser REST client

For example, use the following information in the Firefox plugin *RESTClient*
[https://github.com/chao/RESTClient\protect\T1\textbraceright

```
POST: http://192.168.1.43:8181/restconf/operations/hello:hello-world
```

Header:

```
application/json
```

Body:

```
{"input": {
    "name": "Andrew"
  }
}
```

Troubleshooting

If you get a response code 501 while attempting to POST /operations/hello:hello-world, check the file: HelloProvider.java and make sure the helloService member is being set. By not invoking "session.addRpcImplementation()" the REST API will be unable to map /operations/hello:hello-world url to HelloWorldImpl.

2.1.2 Project-specific Developer Guides

ALTO Developer Guide

Overview

The topics of this guide are:

1. How to add alto projects as dependencies;
2. How to put/fetch data from ALTO;
3. Basic API and DataType;
4. How to use customized service implementations.

Adding ALTO Projects as Dependencies

Most ALTO packages can be added as dependencies in Maven projects by putting the following code in the *pom.xml* file.

```
<dependency>
    <groupId>org.opendaylight.alto</groupId>
    <artifactId>${THE_NAME_OF_THE_PACKAGE_YOU_NEED}</artifactId>
    <version>${ALTO_VERSION}</version>
</dependency>
```

The current stable version for ALTO is `0.3.0-Boron`.

Putting/Fetching data from ALTO

Using RESTful API

There are two kinds of RESTful APIs for ALTO: the one provided by `alto-northbound` which follows the formats defined in RFC 7285, and the one provided by RESTCONF whose format is defined by the YANG model proposed in

this draft.

One way to get the URLs for the resources from `alto-northbound` is to visit the IRD service first where there is a `uri` field for every entry. However, the IRD service is not yet implemented so currently the developers have to construct the URLs themselves. The base URL is `/alto` and below is a list of the specific paths defined in `alto-core/standard-northbound-route` using Jersey `@Path` annotation:

- `/ird/{rid}`: the path to access *IRD* services;
- `/networkmap/{rid}[/{tag}]`: the path to access *Network Map* and *Filtered Network Map* services;
- `/costmap/{rid}[/{tag}[/{mode}/{metric}]]`: the path to access *Cost Map* and *Filtered Cost Map* services;
- `/endpointprop`: the path to access *Endpoint Property* services;
- `/endpointcost`: the path to access *Endpoint Cost* services.

Note: The segments in brackets are optional.

If you want to fetch the data using RESTCONF, it is highly recommended to take a look at the `apidoc` page (http://\protect\T1\textbraceleftcontroller_ip\protect\T1\textbraceright:8181/apidoc/explorer/index.html) after installing the `odl-alto-release` feature in karaf.

It is also worth pointing out that `alto-northbound` only supports `GET` and `POST` operations so it is impossible to manipulate the data through its RESTful APIs. To modify the data, use `PUT` and `DELETE` methods with RESTCONF.

Note: The current implementation uses the `configuration` data store and that enables the developers to modify the data directly through RESTCONF. In the future this approach might be disabled in the core packages of ALTO but may still be available as an extension.

Using MD-SAL

You can also fetch data from the datastore directly.

First you must get the access to the datastore by registering your module with a data broker.

Then an `InstanceIdentifier` must be created. Here is an example of how to build an `InstanceIdentifier` for a *network map*:

```
import org.opendaylight...alto...Resources;
import org.opendaylight...alto...resources.NetworkMaps;
import org.opendaylight...alto...resources.network.maps.NetworkMap;
import org.opendaylight...alto...resources.network.maps.NetworkMapKey;
...
protected
InstanceIdentifier<NetworkMap> getNetworkMapIID(String resource_id) {
  ResourceId rid = ResourceId.getDefaultInstance(resource_id);
  NetworkMapKey key = new NetworkMapKey(rid);
  InstanceIdentifier<NetworkMap> iid = null;
  iid = InstanceIdentifier.builder(Resources.class)
                          .child(NetworkMaps.class)
                          .child(NetworkMap.class, key)
                          .build();
  return iid;
```

(continues on next page)

(continued from previous page)

```
}
...
```

With the `InstanceIdentifier` you can use `ReadOnlyTransaction`, `WriteTransaction` and `ReadWriteTransaction` to manipulate the data accordingly. The `simple-impl` package, which provides some of the AD-SAL APIs mentioned above, is using this method to get data from the datastore and then convert them into RFC7285-compatible objects.

Basic API and DataType

1. alto-basic-types: Defines basic types of ALTO protocol.
2. alto-service-model-api: Includes the YANG models for the five basic ALTO services defined in RFC 7285.
3. alto-resourcepool: Manages the meta data of each ALTO service, including capabilities and versions.
4. alto-northbound: Provides the root of RFC7285-compatible services at http://localhost:8080/alto.
5. alto-northbound-route: Provides the root of the network map resources at http://localhost:8080/alto/networkmap/.

How to customize service

Define new service API

Add a new module in `alto-core/standard-service-models`. For example, we named our service model module as `model-example`.

Implement service RPC

Add a new module in `alto-basic` to implement a service RPC in `alto-core`.

Currently `alto-core/standard-service-models/model-base` has defined a template of the service RPC. You can define your own RPC using `augment` in YANG. Here is an example in `alto-simpleird`.

Register northbound route

If necessary, you can add a northbound route module in `alto-core/standard-northbound-routes`.

Atrium Developer Guide

Overview

Project Atrium is an open source SDN distribution - a vertically integrated set of open source components which together form a complete SDN stack. It's goals are threefold:

- Close the large integration-gap of the elements that are needed to build an SDN stack - while there are multiple choices at each layer, there are missing pieces with poor or no integration.

- Overcome a massive gap in interoperability - This exists both at the switch level, where existing products from different vendors have limited compatibility, making it difficult to connect an arbitrary switch and controller and at an API level, where its difficult to write a portable application across multiple controller platforms.
- Work closely with network operators on deployable use-cases, so that they could download near production quality code from one location, and get started with functioning software defined networks on real hardware.

Architecture

The key components of Atrium BGP Peering Router Application are as follows:

- Data Plane Switch - Data plane switch is the entity that uses flow table entries installed by BGP Routing Application through SDN controller. In the simplest form data plane switch with the installed flows act like a BGP Router.
- OpenDaylight Controller - OpenDaylight SDN controller has many utility applications or plugins which are leveraged by the BGP Router application to manage the control plane information.
- BGP Routing Application - An application running within the OpenDaylight runtime environment to handle I-BGP updates.
- *DIDM* - DIDM manages the drivers specific to each data plane switch connected to the controller. The drivers are created primarily to hide the underlying complexity of the devices and to expose a uniform API to applications.
- Flow Objectives API - The driver implementation provides a pipeline abstraction and exposes Flow Objectives API. This means applications need to be aware of only the Flow Objectives API without worrying about the Table IDs or the pipelines.
- Control Plane Switch - This component is primarily used to connect the OpenDaylight SDN controller with the Quagga Soft-Router and establish a path for forwarding E-BGP packets to and from Quagga.
- Quagga soft router - An open source routing software that handles E-BGP updates.

Key APIs and Interfaces

BGP Routing Configuration

The BGP Routing Configuration maintains information about its BGP Speakers & BGP Peers.

- Configuration data about BGP speakers can be accessed from the below URL:

```
GET http://<controller_ip>:8181/restconf/config/bgpconfig:bgpSpeakers/
```

- Configuration data about BGP peers can be accessed from the below URL:

```
GET http://<controller_ip>:8181/restconf/config/bgpconfig:bgpPeers/
```

Host Service

Host Service API contains the host specific details that can be used during address resolution

- Host specific data can be accessed by using the below REST request:

```
GET http://<controller_ip>:8181/restconf/config/hostservice-api:addresses/
```

BGP Routing Information Base

The BGP RIB module stores all the route information that it has learnt from its peers.

- Routing Information Base entries can be accessed from the URL below:

```
GET http://<controller_ip>:8181/restconf/operational/bgp-rib:bgp-rib/
```

Forwarding Information Base

The Forwarding Information Base is used to keep track of active FIB entries.

- FIB entries can be accessed from the URL below:

```
GET http://<controller_ip>:8181/restconf/config/routingservice-api:fibEntries/
```

BGP Developer Guide

Overview

This section provides an overview of the `odl-bgpcep-bgp-all` Karaf feature. This feature will install everything needed for BGP (Border Gateway Protocol) from establishing the connection, storing the data in RIBs (Route Information Base) and displaying data in network-topology overview.

BGP Architecture

Each feature represents a module in the BGPCEP codebase. The following diagram illustrates how the features are related.

Key APIs and Interfaces

BGP concepts

This module contains the base BGP concepts contained in RFC 4271, RFC 4760, RFC 4456, RFC 1997 and RFC 4360.

All the concepts are described in one yang model: bgp-types.yang.

Outside generated classes, there is just one class NextHopUtil that contains methods for serializing and parsing NextHop.

BGP parser

Base BGP parser includes messages and attributes from RFC 4271, RFC 4760, RFC 1997 and RFC 4360.

API module defines BGP messages in YANG.

IMPL module contains actual parsers and serializers for BGP messages and Activator class

SPI module contains helper classes needed for registering parsers into activators

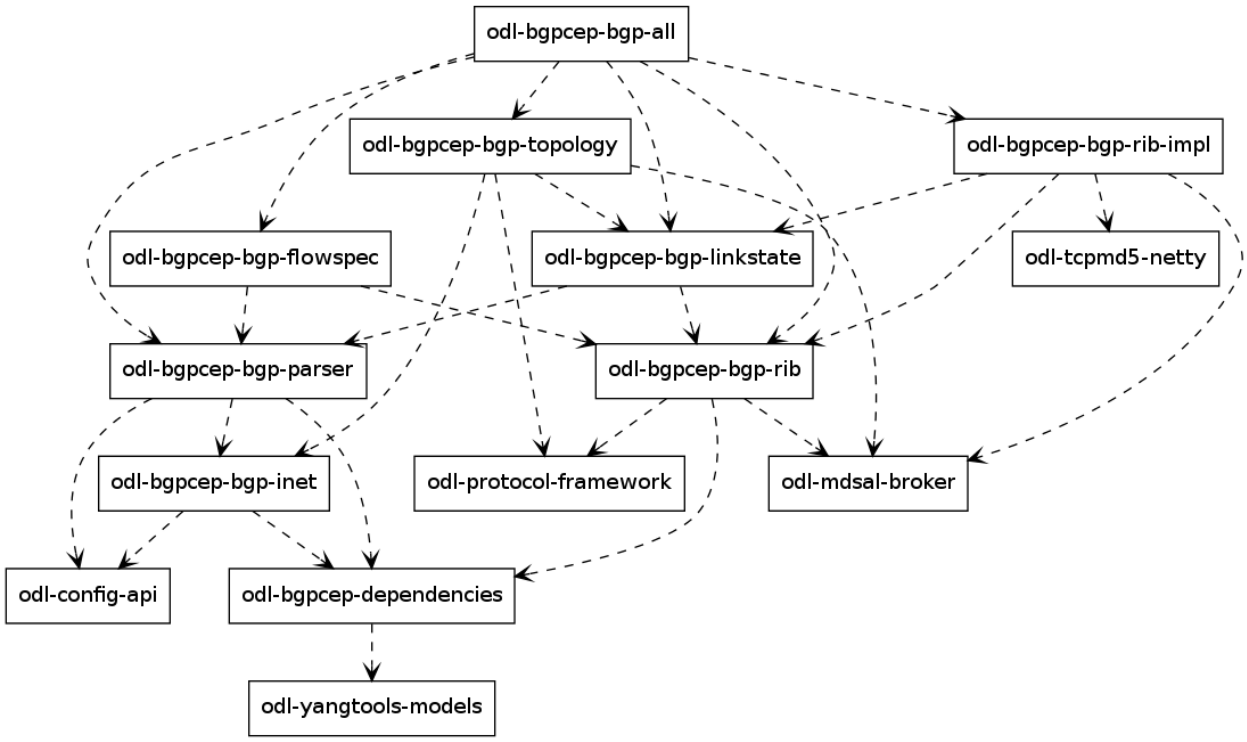

Fig. 12: BGP Dependency Tree

Registration

All parsers and serializers need to be registered into the *Extension provider*. This *Extension provider* is configured in initial configuration of the parser-spi module (`31-bgp.xml`).

```
<module>
 <type xmlns:prefix="urn:opendaylight:params:xml:ns:yang:controller:bgp:parser:spi">
↪prefix:bgp-extensions-impl</type>
 <name>global-bgp-extensions</name>
 <extension>
  <type xmlns:bgpspi="urn:opendaylight:params:xml:ns:yang:controller:bgp:parser:spi">
↪bgpspi:extension</type>
  <name>base-bgp-parser</name>
 </extension>
 <extension>
  <type xmlns:bgpspi="urn:opendaylight:params:xml:ns:yang:controller:bgp:parser:spi">
↪bgpspi:extension</type>
  <name>bgp-linkstate</name>
 </extension>
</module>
```

- *base-bgp-parser* - will register parsers and serializers implemented in the bgp-parser-impl module
- *bgp-linkstate* - will register parsers and serializers implemented in the bgp-linkstate module

The bgp-linkstate module is a good example of a BGP parser extension.

The configuration of bgp-parser-spi specifies one implementation of *Extension provider* that will take care of registering mentioned parser extensions: SimpleBGPExtensionProviderContext. All registries are implemented in package bgp-parser-spi.

Serializing

The serializing of BGP elements is mostly done in the same way as in *PCEP*, the only exception is the serialization of path attributes, which is described here. Path attributes are different from any other BGP element, as path attributes don't implement one common interface, but this interface contains getters for individual path attributes (this structure is because update message can contain exactly one instance of each path attribute). This means, that a given *PathAttributes* object, you can only get to the specific type of the path attribute through checking its presence. Therefore method *serialize()* in *AttributeRegistry*, won't look up the registered class, instead it will go through the registrations and offer this object to the each registered parser. This way the object will be passed also to serializers unknown to module bgp-parser, for example to LinkstateAttributeParser. RFC 4271 recommends ordering path attributes, hence the serializers are ordered in a list as they are registered in the *Activator*. In other words, this is the only case, where registration ordering matters.

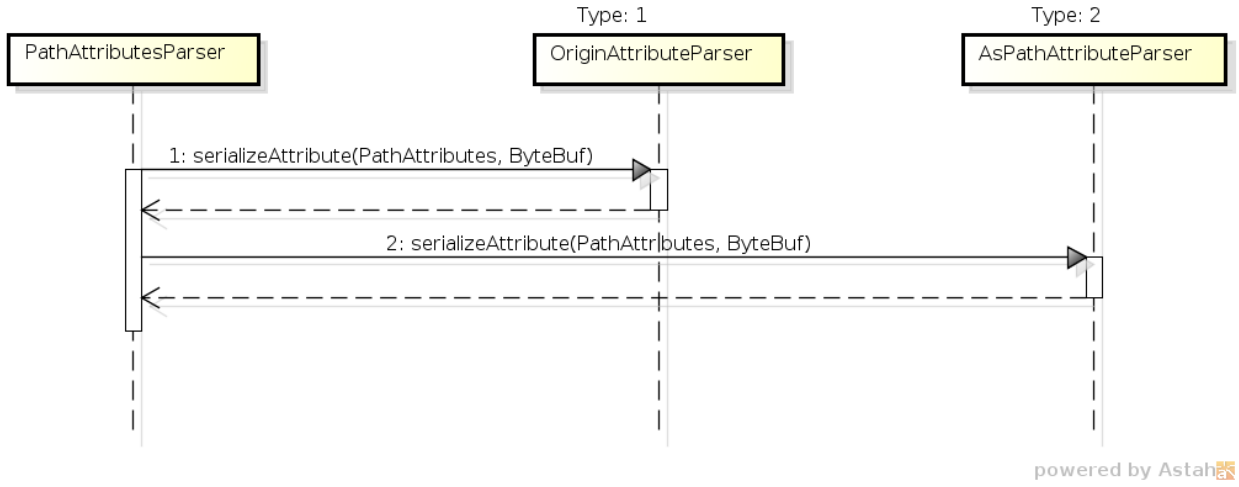

Fig. 13: PathAttributesSerialization

serialize() method in each Path Attribute parser contains check for presence of its attribute in the PathAttributes object, which simply returns, if the attribute is not there:

```
if (pathAttributes.getAtomicAggregate() == null) {
    return;
}
//continue with serialization of Atomic Aggregate
```

BGP RIB

The BGP RIB module can be divided into two parts:

- BGP listener and speaker session handling
- RIB handling.

Session handling

`31-bgp.xml` defines only bgp-dispatcher and the parser it should be using (global-bgp-extensions).

```
<module>
 <type>prefix:bgp-dispatcher-impl</type>
 <name>global-bgp-dispatcher</name>
 <bgp-extensions>
  <type>bgpspi:extensions</type>
  <name>global-bgp-extensions</name>
 </bgp-extensions>
 <boss-group>
  <type>netty:netty-threadgroup</type>
  <name>global-boss-group</name>
 </boss-group>
 <worker-group>
  <type>netty:netty-threadgroup</type>
  <name>global-worker-group</name>
 </worker-group>
</module>
```

For user configuration of BGP, check User Guide.

Synchronization

Synchronization is a phase, where upon connection, a BGP speaker sends all available data about topology to its new client. After the whole topology has been advertised, the synchronization is over. For the listener, the synchronization is over when the RIB receives End-of-RIB (EOR) messages. There is a special EOR message for each AFI (Address Family Identifier).

- IPv4 EOR is an empty Update message.
- Ipv6 EOR is an Update message with empty MP_UNREACH attribute where AFI and SAFI (Subsequent Address Family Identifier) are set to Ipv6. OpenDaylight also supports EOR for IPv4 in this format.
- Linkstate EOR is an Update message with empty MP_UNREACH attribute where AFI and SAFI are set to Linkstate.

For BGP connections, where both peers support graceful restart, the EORs are sent by the BGP speaker and are redirected to RIB, where the specific AFI/SAFI table is set to *true*. Without graceful restart, the messages are generated by OpenDaylight itself and sent after second keepalive for each AFI/SAFI. This is done in BGPSynchronization.

Peers

BGPPeer has various meanings. If you configure BGP listener, *BGPPeer* represents the BGP listener itself. If you are configuring BGP speaker, you need to provide a list of peers, that are allowed to connect to this speaker. Unknown peer represents, in this case, a peer that is allowed to be refused. *BGPPeer* represents in this case peer, that is supposed to connect to your speaker. *BGPPeer* is stored in BGPPeerRegistry. This registry controls the number of sessions. Our strict implementation limits sessions to one per peer.

ApplicationPeer is a special case of peer, that has it's own RIB. This RIB is populated from RESTCONF. The RIB is synchronized with default BGP RIB. Incoming routes to the default RIB are treated in the same way as they were from a BGP peer (speaker or listener) in the network.

RIB handling

RIB (Route Information Base) is defined as a concept in RFC 4271. RFC does not define how it should be implemented. In our implementation, the routes are stored in the MD-SAL datastore. There are four supported routes - *Ipv4Routes*, *Ipv6Routes*, *LinkstateRoutes* and *FlowspecRoutes*.

Each route type needs to provide a RIBSupport.java implementation. *RIBSupport* tells RIB how to parse binding-aware data (BGP Update message) to binding-independent (datastore format).

Following picture describes the data flow from BGP message that is sent to *BGPPeer* to datastore and various types of RIB.

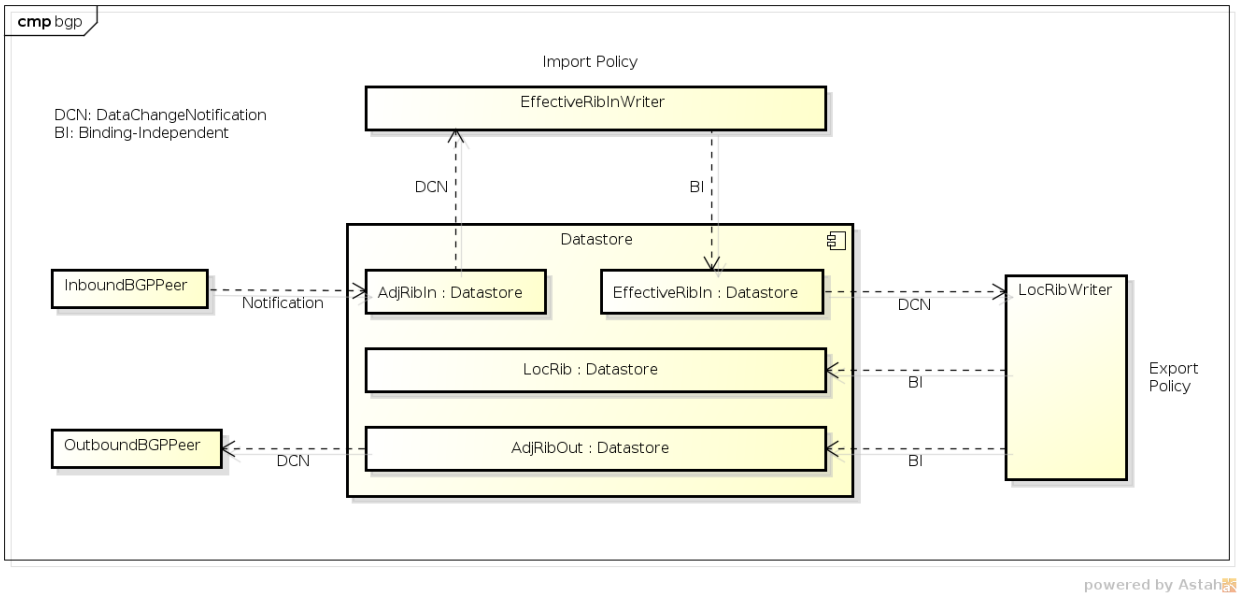

Fig. 14: RIB

AdjRibInWriter - represents the first step in putting data to datastore. This writer is notified whenever a peer receives an Update message. The message is transformed into binding-independent format and pushed into datastore to *adj-rib-in*. This RIB is associated with a peer.

EffectiveRibInWriter - this writer is notified whenever *adj-rib-in* is updated. It applies all configured import policies to the routes and stores them in *effective-rib-in*. This RIB is also associated with a peer.

LocRibWriter - this writer is notified whenever **any** *effective-rib-in* is updated (in any peer). Performs best path selection filtering and stores the routes in *loc-rib*. It also determines which routes need to be advertised and fills in *adj-rib-out* that is per peer as well.

AdjRibOutListener - listens for changes in *adj-rib-out*, transforms the routes into BGPUpdate messages and sends them to its associated peer.

BGP inet

This module contains only one YANG model bgp-inet.yang that summarizes the ipv4 and ipv6 extensions to RIB routes and BGP messages.

BGP flowspec

BGP flowspec is a module that implements RFC 5575 for IPv4 AFI and draft-ietf-idr-flow-spec-v6-06 for IPv6 AFI. The RFC defines an extension to BGP in form of a new subsequent address family, NLRI and extended communities. All of those are defined in the bgp-flowspec.yang model. In addition to generated sources, the module contains parsers for newly defined elements and RIBSupport for flowspec-routes. The route key of flowspec routes is a string representing human-readable flowspec request.

BGP linkstate

BGP linkstate is a module that implements draft-ietf-idr-ls-distribution version 04. The draft defines an extension to BGP in form of a new address family, subsequent address family, NLRI and path attribute. All of those are defined in the bgp-linkstate.yang model. In addition to generated sources, the module contains LinkstateAttributeParser, LinkstateNlriParser, activators for both, parser and RIB, and RIBSupport handler for linkstate address family. As each route needs a key, in case of linkstate, the route key is defined as a binary string, containing all the NLRI serialized to byte format. The BGP linkstate extension also supports distribution of MPLS TE state as defined in draft-ietf-idr-te-lsp-distribution-03, extension for Segment Routing draft-gredler-idr-bgp-ls-segment-routing-ext-00 and Segment Routing Egress Peer Engineering draft-ietf-idr-bgpls-segment-routing-epe-02.

BGP labeled-unicast

BGP labeled unicast is a module that implements RFC 3107. The RFC defines an extension to the BGP MP to carry Label Mapping Information as a part of the NLRI. The AFI indicates, as usual, the address family of the associated route. The fact that the NLRI contains a label is indicated by using SAFI value 4. All of those are defined in bgp-labeled-unicast.yang model. In addition to the generated sources, the module contains new NLRI codec and RIBSupport. The route key is defined as a binary, where whole NLRI information is encoded.

BGP topology provider

BGP data besides RIB, is stored in network-topology view. The format of how the data is displayed there conforms to draft-clemm-netmod-yang-network-topo.

API Reference Documentation

Javadocs are generated while creating mvn:site and they are located in target/ directory in each module.

BGP Monitoring Protocol Developer Guide

Overview

This section provides an overview of **feature odl-bgpcep-bmp**. This feature will install everything needed for BMP (BGP Monitoring Protocol) including establishing the connection, processing messages, storing information about monitored routers, peers and their Adj-RIB-In (unprocessed routing information) and Post-Policy Adj-RIB-In and displaying data in BGP RIBs overview. The OpenDaylight BMP plugin plays the role of a monitoring station.

Key APIs and Interfaces

Session handling

32-bmp.xml defines only bmp-dispatcher the parser should be using (global-bmp-extensions).

```
<module>
 <type xmlns:prefix="urn:opendaylight:params:xml:ns:yang:controller:bmp:impl">
↪prefix:bmp-dispatcher-impl</type>
 <name>global-bmp-dispatcher</name>
  <bmp-extensions>
```

(continues on next page)

(continued from previous page)

```
   <type xmlns:bmp-spi="urn:opendaylight:params:xml:ns:yang:controller:bmp:spi">bmp-
↪spi:extensions</type>
   <name>global-bmp-extensions</name>
  </bmp-extensions>
  <boss-group>
   <type xmlns:netty="urn:opendaylight:params:xml:ns:yang:controller:netty">
↪netty:netty-threadgroup</type>
   <name>global-boss-group</name>
  </boss-group>
  <worker-group>
   <type xmlns:netty="urn:opendaylight:params:xml:ns:yang:controller:netty">
↪netty:netty-threadgroup</type>
   <name>global-worker-group</name>
 </worker-group>
</module>
```

For user configuration of BMP, check User Guide.

Parser

The base BMP parser includes messages and attributes from https://tools.ietf.org/html/draft-ietf-grow-bmp-15

Registration

All parsers and serializers need to be registered into *Extension provider*. This *Extension provider* is configured in initial configuration of the parser (*32-bmp.xml*).

```
<module>
 <type xmlns:prefix="urn:opendaylight:params:xml:ns:yang:controller:bmp:spi">
↪prefix:bmp-extensions-impl</type>
 <name>global-bmp-extensions</name>
 <extension>
  <type xmlns:bmp-spi="urn:opendaylight:params:xml:ns:yang:controller:bmp:spi">bmp-
↪spi:extension</type>
  <name>bmp-parser-base</name>
 </extension>
</module>
```

- *bmp-parser-base* - will register parsers and serializers implemented in bmp-impl module

Parsing

Parsing of BMP elements is mostly done equally to BGP. Some of the BMP messages includes wrapped BGP messages.

BMP Monitoring Station

The BMP application (Monitoring Station) serves as message processor incoming from monitored routers. The processed message is transformed and relevant information is stored. Route information is stored in a BGP RIB data structure.

BMP data is displayed only through one URL that is accessible from the base BMP URL:

'http://<controllerIP>:8181/restconf/operational/bmp-monitor:bmp-monitor <http://<controllerIP>:8181/restconf/operational/bmp-monitor:bmp-monitor>'__

Each Monitor station will be displayed and it may contains multiple monitored routers and peers within:

```
<bmp-monitor xmlns="urn:opendaylight:params:xml:ns:yang:bmp-monitor">
 <monitor>
 <monitor-id>example-bmp-monitor</monitor-id>
  <router>
  <router-id>127.0.0.11</router-id>
   <status>up</status>
   <peer>
    <peer-id>20.20.20.20</peer-id>
    <as>72</as>
    <type>global</type>
    <peer-session>
     <remote-port>5000</remote-port>
     <timestamp-sec>5</timestamp-sec>
     <status>up</status>
     <local-address>10.10.10.10</local-address>
     <local-port>220</local-port>
    </peer-session>
    <pre-policy-rib>
     <tables>
      <afi xmlns:x="urn:opendaylight:params:xml:ns:yang:bgp-types">x:ipv4-address-
↪family</afi>
      <safi xmlns:x="urn:opendaylight:params:xml:ns:yang:bgp-types">x:unicast-
↪subsequent-address-family</safi>
      <ipv4-routes xmlns="urn:opendaylight:params:xml:ns:yang:bgp-inet">
       <ipv4-route>
        <prefix>10.10.10.0/24</prefix>
         <attributes>
          ...
         </attributes>
       </ipv4-route>
      </ipv4-routes>
      <attributes>
       <uptodate>true</uptodate>
      </attributes>
     </tables>
    </pre-policy-rib>
    <address>10.10.10.10</address>
    <post-policy-rib>
     ...
    </post-policy-rib>
    <bgp-id>20.20.20.20</bgp-id>
    <stats>
     <timestamp-sec>5</timestamp-sec>
     <invalidated-cluster-list-loop>53</invalidated-cluster-list-loop>
     <duplicate-prefix-advertisements>16</duplicate-prefix-advertisements>
     <loc-rib-routes>100</loc-rib-routes>
     <duplicate-withdraws>11</duplicate-withdraws>
     <invalidated-as-confed-loop>55</invalidated-as-confed-loop>
     <adj-ribs-in-routes>10</adj-ribs-in-routes>
     <invalidated-as-path-loop>66</invalidated-as-path-loop>
     <invalidated-originator-id>70</invalidated-originator-id>
     <rejected-prefixes>8</rejected-prefixes>
    </stats>
```

(continues on next page)

(continued from previous page)

```
    </peer>
    <name>name</name>
    <description>description</description>
    <info>some info;</info>
   </router>
  </monitor>
</bmp-monitor>
</source>
```

API Reference Documentation

Javadocs are generated while creating mvn:site and they are located in target/ directory in each module.

CAPWAP Developer Guide

Overview

The Control And Provisioning of Wireless Access Points (CAPWAP) plugin project aims to provide new southbound interface for controller to be able to monitor and manage CAPWAP compliant wireless termination point (WTP) network devices. The CAPWAP feature will provide REST based northbound APIs.

CAPWAP Architecture

The CAPWAP feature is implemented as an MD-SAL based provider module, which helps discover WTP devices and update their states in the MD-SAL operational datastore.

CAPWAP APIs and Interfaces

This section describes the APIs for interacting with the CAPWAP plugin.

Discovered WTPs

The CAPWAP project maintains list of discovered CAPWAP WTPs that is YANG-based in MD-SAL. These models are available via RESTCONF.

- Name: Discovered-WTPs
- URL: http://\protect\T1\textdollar\protect\T1\textbraceleftipaddress\protect\T1\textbraceright:8181/restconf/operational/capwap-impl:capwap-ac-root/
- Description: Displays list of discovered WTPs and their basic attributes

API Reference Documentation

Go to http://\protect\T1\textdollar\protect\T1\textbraceleftipaddress\protect\T1\textbraceright:8181/apidoc/explorer/index.html, sign in, and expand the capwap-impl panel. From there, users can execute various API calls to test their CAPWAP deployment.

Cardinal: OpenDaylight Monitoring as a Service

Overview

Cardinal (OpenDaylight Monitoring as a Service) enables OpenDaylight and the underlying software defined network to be remotely monitored by deployed Network Management Systems (NMS) or Analytics suite. In the Boron release, Cardinal adds:

1. OpenDaylight MIB.
2. Enable ODL diagnostics/monitoring to be exposed across SNMP (v2c, v3) and REST north-bound.
3. Extend ODL System health, Karaf parameter and feature info, ODL plugin scalability and network parameters.
4. Support autonomous notifications (SNMP Traps).

Cardinal Architecture

The Cardinal architecture can be found at the below link:

https://wiki.opendaylight.org/images/8/89/Cardinal-ODL_Monitoring_as_a_Service_V2.pdf

Key APIs and Interfaces

There are 2 main APIs for requesting snmpget request of the Karaf info and System info. To expose these APIs, it assumes that you already have the `odl-cardinal` and `odl-restconf` features installed. You can do that by entering the following at the Karaf console:

```
feature:install odl-cardinal
feature:install odl-restconf-all
```

System Info APIs

Open the REST interface and using the basic authentication, execute REST APIs for system info as:

```
http://localhost:8181/restconf/operational/cardinal:CardinalSystemInfo/
```

You should get the response code of the same as 200 OK with the following output as:

```
{
  "CardinalSystemInfo": {
    "odlSystemMemUsage": " 9",
    "odlSystemSysInfo": " OpenDaylight Node Information",
    "odlSystemOdlUptime": " 00:29",
    "odlSystemCpuUsage": " 271",
    "odlSystemHostAddress": " Address of the Host should come up"
  }
}
```

Karaf Info APIs

Open the REST interface and using the basic authentication, execute REST APIs for system info as:

```
http://localhost:8181/restconf/operational/cardinal-karaf:CardinalKarafInfo/
```

You should get the response code of the same as 200 OK with the following output as:


```
  {
  "CardinalKarafInfo": {
    "odlKarafBundleListActive1": " org.ops4j.pax.url.mvn_2.4.5 [1]",
    "odlKarafBundleListActive2": " org.ops4j.pax.url.wrap_2.4.5 [2]",
    "odlKarafBundleListActive3": " org.ops4j.pax.logging.pax-logging-api_1.8.4 [3]",
    "odlKarafBundleListActive4": " org.ops4j.pax.logging.pax-logging-service_1.8.4 [4]
↪",
    "odlKarafBundleListActive5": " org.apache.karaf.service.guard_3.0.6 [5]",
    "odlKarafBundleListActive6": " org.apache.felix.configadmin_1.8.4 [6]",
    "odlKarafBundleListActive7": " org.apache.felix.fileinstall_3.5.2 [7]",
    "odlKarafBundleListActive8": " org.objectweb.asm.all_5.0.3 [8]",
    "odlKarafBundleListActive9": " org.apache.aries.util_1.1.1 [9]",
    "odlKarafBundleListActive10": " org.apache.aries.proxy.api_1.0.1 [10]",
    "odlKarafBundleListInstalled1": " org.ops4j.pax.url.mvn_2.4.5 [1]",
    "odlKarafBundleListInstalled2": " org.ops4j.pax.url.wrap_2.4.5 [2]",
    "odlKarafBundleListInstalled3": " org.ops4j.pax.logging.pax-logging-api_1.8.4 [3]
↪",
    "odlKarafBundleListInstalled4": " org.ops4j.pax.logging.pax-logging-service_1.8.4␣
↪[4]",
    "odlKarafBundleListInstalled5": " org.apache.karaf.service.guard_3.0.6 [5]",
    "odlKarafFeatureListInstalled1": " config",
    "odlKarafFeatureListInstalled2": " region",
    "odlKarafFeatureListInstalled3": " package",
    "odlKarafFeatureListInstalled4": " http",
    "odlKarafFeatureListInstalled5": " war",
    "odlKarafFeatureListInstalled6": " kar",
    "odlKarafFeatureListInstalled7": " ssh",
    "odlKarafFeatureListInstalled8": " management",
    "odlKarafFeatureListInstalled9": " odl-netty",
    "odlKarafFeatureListInstalled10": " odl-lmax",
    "odlKarafBundleListResolved1": " org.ops4j.pax.url.mvn_2.4.5 [1]",
    "odlKarafBundleListResolved2": " org.ops4j.pax.url.wrap_2.4.5 [2]",
    "odlKarafBundleListResolved3": " org.ops4j.pax.logging.pax-logging-api_1.8.4 [3]",
    "odlKarafBundleListResolved4": " org.ops4j.pax.logging.pax-logging-service_1.8.4␣
↪[4]",
    "odlKarafBundleListResolved5": " org.apache.karaf.service.guard_3.0.6 [5]",
    "odlKarafFeatureListUnInstalled1": " aries-annotation",
    "odlKarafFeatureListUnInstalled2": " wrapper",
    "odlKarafFeatureListUnInstalled3": " service-wrapper",
    "odlKarafFeatureListUnInstalled4": " obr",
    "odlKarafFeatureListUnInstalled5": " http-whiteboard",
    "odlKarafFeatureListUnInstalled6": " jetty",
    "odlKarafFeatureListUnInstalled7": " webconsole",
    "odlKarafFeatureListUnInstalled8": " scheduler",
    "odlKarafFeatureListUnInstalled9": " eventadmin",
    "odlKarafFeatureListUnInstalled10": " jasypt-encryption"
  }
}
```

Controller

Overview

OpenDaylight Controller is Java-based, model-driven controller using YANG as its modeling language for various aspects of the system and applications and with its components serves as a base platform for other OpenDaylight applications.

The OpenDaylight Controller relies on the following technologies:

- **OSGI** - This framework is the back-end of OpenDaylight as it allows dynamically loading of bundles and packages JAR files, and binding bundles together for exchanging information.
- **Karaf** - Application container built on top of OSGI, which simplifies operational aspects of packaging and installing applications.
- **YANG** - a data modeling language used to model configuration and state data manipulated by the applications, remote procedure calls, and notifications.

The OpenDaylight Controller provides following model-driven subsystems as a foundation for Java applications:

- *Config Subsystem* - an activation, dependency-injection and configuration framework, which allows two-phase commits of configuration and dependency-injection, and allows for run-time rewiring.
- MD-SAL - messaging and data storage functionality for data, notifications and RPCs modeled by application developers. MD-SAL uses YANG as the modeling for both interface and data definitions, and provides a messaging and data-centric runtime for such services based on YANG modeling.
- **MD-SAL Clustering** - enables cluster support for core MD-SAL functionality and provides location-transparent accesss to YANG-modeled data.

The OpenDaylight Controller supports external access to applications and data using following model-driven protocols:

- **NETCONF** - XML-based RPC protocol, which provides abilities for client to invoke YANG-modeled RPCs, receive notifications and to read, modify and manipulate YANG modeled data.
- **RESTCONF** - HTTP-based protocol, which provides REST-like APIs to manipulate YANG modeled data and invoke YANG modeled RPCs, using XML or JSON as payload format.

MD-SAL Overview

The Model-Driven Service Adaptation Layer (MD-SAL) is message-bus inspired extensible middleware component that provides messaging and data storage functionality based on data and interface models defined by application developers (i.e. user-defined models).

The MD-SAL:

- Defines a **common-layer, concepts, data model building blocks and messaging patterns** and provides infrastructure / framework for applications and inter-application communication.
- Provide common support for user-defined transport and payload formats, including payload serialization and adaptation (e.g. binary, XML or JSON).

The MD-SAL uses **YANG** as the modeling language for both interface and data definitions, and provides a messaging and data-centric runtime for such services based on YANG modeling.

The MD-SAL provides two different API types (flavours):

- **MD-SAL Binding:** MD-SAL APIs which extensively uses APIs and classes generated from YANG models, which provides compile-time safety.
- **MD-SAL DOM:** (Document Object Model) APIs which uses DOM-like representation of data, which makes them more powerful, but provides less compile-time safety.

Note: Model-driven nature of the MD-SAL and **DOM**-based APIs allows for behind-the-scene API and payload type mediation and transformation to facilitate seamless communication between applications - this enables for other components and applications to provide connectors / expose different set of APIs and derive most of its functionality purely from models, which all existing code can benefit from without modification. For example **RESTCONF Connector** is an application built on top of MD-SAL and exposes YANG-modeled application APIs transparently via HTTP and adds support for XML and JSON payload type.

Basic concepts

Basic concepts are building blocks which are used by applications, and from which MD-SAL uses to define messaging patterns and to provide services and behavior based on developer-supplied YANG models.

Data Tree All state-related data are modeled and represented as data tree, with possibility to address any element / subtree

- **Operational Data Tree** - Reported state of the system, published by the providers using MD-SAL. Represents a feedback loop for applications to observe state of the network / system.
- **Configuration Data Tree** - Intended state of the system or network, populated by consumers, which expresses their intention.

Instance Identifier Unique identifier of node / subtree in data tree, which provides unambiguous information, how to reference and retrieve node / subtree from conceptual data trees.

Notification Asynchronous transient event which may be consumed by subscribers and they may act upon it

RPC asynchronous request-reply message pair, when request is triggered by consumer, send to the provider, which in future replies with reply message.

Note: In MD-SAL terminology, the term *RPC* is used to define the input and output for a procedure (function) that is to be provided by a provider, and mediated by the MD-SAL, that means it may not result in remote call.

Messaging Patterns

MD-SAL provides several messaging patterns using broker derived from basic concepts, which are intended to transfer YANG modeled data between applications to provide data-centric integration between applications instead of API-centric integration.

- **Unicast communication**
  - **Remote Procedure Calls** - unicast between consumer and provider, where consumer sends **request** message to provider, which asynchronously responds with **reply** message
- **Publish / Subscribe**
  - **Notifications** - multicast transient message which is published by provider and is delivered to subscribers
  - **Data Change Events** - multicast asynchronous event, which is sent by data broker if there is change in conceptual data tree, and is delivered to subscribers

- **Transactional access to Data Tree**
  - Transactional **reads** from conceptual **data tree** - read-only transactions with isolation from other running transactions.
  - Transactional **modification** to conceptual **data tree** - write transactions with isolation from other running transactions.
  - **Transaction chaining**

MD-SAL Data Transactions

MD-SAL **Data Broker** provides transactional access to conceptual **data trees** representing configuration and operational state.

Note: **Data tree** usually represents state of the modeled data, usually this is state of controller, applications and also external systems (network devices).

Transactions provide *stable and isolated view* from other currently running transactions. The state of running transaction and underlying data tree is not affected by other concurrently running transactions.

Write-Only Transaction provides only modification capabilities, but does not provide read capabilities. Write-only transaction is allocated using `newWriteOnlyTransaction()`.

> **Note:** This allows less state tracking for write-only transactions and allows MD-SAL Clustering to optimize internal representation of transaction in cluster.

Read-Write Transaction provides both read and write capabilities. It is allocated using `newReadWriteTransaction()`.

Read-Only Transaction provides stable read-only view based on current data tree. Read-only view is not affected by any subsequent write transactions. Read-only transaction is allocated using `newReadOnlyTransaction()`.

> **Note:** If an application needs to observe changes itself in data tree, it should use **data tree listeners** instead of read-only transactions and polling data tree.

Transactions may be allocated using the **data broker** itself or using **transaction chain**. In the case of **transaction chain**, the new allocated transaction is not based on current state of data tree, but rather on state introduced by previous transaction from the same chain, even if the commit for previous transaction has not yet occurred (but transaction was submitted).

Write-Only & Read-Write Transaction

Write-Only and Read-Write transactions provide modification capabilities for the conceptual data trees.

1. application allocates new transactions using `newWriteOnlyTransaction()` or `newReadWriteTransaction()`.
2. application *modifies data tree* using `put`, `merge` and/or `delete`.
3. application finishes transaction using `submit()`, which *seals transaction and submits* it to be processed.
4. application observes the result of the transaction commit using either blocking or asynchronous calls.

The **initial state** of the write transaction is a **stable snapshot** of the current data tree state captured when transaction was created and it's state and underlying data tree are not affected by other concurrently running transactions.

Write transactions are **isolated** from other concurrent write transactions. All *writes are local* to the transaction and represents only a **proposal of state change** for data tree and **are not visible** to any other concurrently running transactions (including read-only transactions).

The transaction *commit may fail* due to failing verification of data or concurrent transaction modifying and affected data in an incompatible way.

Modification of Data Tree

Write-only and read-write transaction provides following methods to modify data tree:

put

```
<T> void put(LogicalDatastoreType store, InstanceIdentifier<T> path, T data);
```

Stores a piece of data at a specified path. This acts as an **add / replace** operation, which is to say that whole subtree will be replaced by the specified data.

merge

```
<T> void merge(LogicalDatastoreType store, InstanceIdentifier<T> path, T data);
```

Merges a piece of data with the existing data at a specified path. Any **pre-existing data** which are not explicitly overwritten **will be preserved**. This means that if you store a container, its child subtrees will be merged.

delete

```
void delete(LogicalDatastoreType store, InstanceIdentifier<?> path);
```

Removes a whole subtree from a specified path.

Submitting transaction

Transaction is submitted to be processed and committed using following method:

```
CheckedFuture<Void,TransactionCommitFailedException> submit();
```

Applications publish the changes proposed in the transaction by calling `submit()` on the transaction. This **seals the transaction** (preventing any further writes using this transaction) and submits it to be processed and applied to global conceptual data tree. The `submit()` method does not block, but rather returns `ListenableFuture`, which will complete successfully once processing of transaction is finished and changes are applied to data tree. If **commit** of data failed, the future will fail with `TransactionFailedException`.

Application may listen on commit state asynchronously using `ListenableFuture`.

```
Futures.addCallback( writeTx.submit(), new FutureCallback<Void>() {
        public void onSuccess( Void result ) {
            LOG.debug("Transaction committed successfully.");
        }

        public void onFailure( Throwable t ) {
            LOG.error("Commit failed.",e);
        }
    });
```

- Submits `writeTx` and registers application provided `FutureCallback` on returned future.
- Invoked when future completed successfully - transaction `writeTx` was successfully committed to data tree.
- Invoked when future failed - commit of transaction `writeTx` failed. Supplied exception provides additional details and cause of failure.

If application need to block till commit is finished it may use `checkedGet()` to wait till commit is finished.

```
try {
    writeTx.submit().checkedGet();
} catch (TransactionCommitFailedException e) {
    LOG.error("Commit failed.",e);
}
```

- Submits `writeTx` and blocks till commit of `writeTx` is finished. If commit fails `TransactionCommitFailedException` will be thrown.
- Catches `TransactionCommitFailedException` and logs it.

Transaction local state

Read-Write transactions maintain transaction-local state, which renders all modifications as if they happened, but this is only local to transaction.

Reads from the transaction returns data as if the previous modifications in transaction already happened.

Let assume initial state of data tree for `PATH` is `A`.

```
ReadWriteTransaction rwTx = broker.newReadWriteTransaction();

rwRx.read(OPERATIONAL,PATH).get();
rwRx.put(OPERATIONAL,PATH,B);
rwRx.read(OPERATIONAL,PATH).get();
rwRx.put(OPERATIONAL,PATH,C);
rwRx.read(OPERATIONAL,PATH).get();
```

- Allocates new `ReadWriteTransaction`.
- Read from `rwTx` will return value `A` for `PATH`.
- Writes value `B` to `PATH` using `rwTx`.
- Read will return value `B` for `PATH`, since previous write occurred in same transaction.
- Writes value `C` to `PATH` using `rwTx`.
- Read will return value `C` for `PATH`, since previous write occurred in same transaction.

Transaction isolation

Running (not submitted) transactions are isolated from each other and changes done in one transaction are not observable in other currently running transaction.

Lets assume initial state of data tree for `PATH` is `A`.

```
ReadOnlyTransaction txRead = broker.newReadOnlyTransaction();
ReadWriteTransaction txWrite = broker.newReadWriteTransaction();
```

(continues on next page)

(continued from previous page)

```
txRead.read(OPERATIONAL,PATH).get();
txWrite.put(OPERATIONAL,PATH,B);
txWrite.read(OPERATIONAL,PATH).get();
txWrite.submit().get();
txRead.read(OPERATIONAL,PATH).get();
txAfterCommit = broker.newReadOnlyTransaction();
txAfterCommit.read(OPERATIONAL,PATH).get();
```

- Allocates read only transaction, which is based on data tree which contains value `A` for `PATH`.
- Allocates read write transaction, which is based on data tree which contains value `A` for `PATH`.
- Read from read-only transaction returns value `A` for `PATH`.
- Data tree is updated using read-write transaction, `PATH` contains `B`. Change is not public and only local to transaction.
- Read from read-write transaction returns value `B` for `PATH`.
- Submits changes in read-write transaction to be committed to data tree. Once commit will finish, changes will be published and `PATH` will be updated for value `B`. Previously allocated transactions are not affected by this change.
- Read from previously allocated read-only transaction still returns value `A` for `PATH`, since it provides stable and isolated view.
- Allocates new read-only transaction, which is based on data tree, which contains value `B` for `PATH`.
- Read from new read-only transaction return value `B` for `PATH` since read-write transaction was committed.

Note: Examples contain blocking calls on future only to illustrate that action happened after other asynchronous action. The use of the blocking call `ListenableFuture#get()` is discouraged for most use-cases and you should use `Futures#addCallback(ListenableFuture, FutureCallback)` to listen asynchronously for result.

Commit failure scenarios

A transaction commit may fail because of following reasons:

Optimistic Lock Failure Another transaction finished earlier and **modified the same node in a non-compatible way**. The commit (and the returned future) will fail with an `OptimisticLockFailedException`.

It is the responsibility of the caller to create a new transaction and submit the same modification again in order to update data tree.

> **Warning:** `OptimisticLockFailedException` usually exposes **multiple writers** to the same data subtree, which may conflict on same resources.
>
> In most cases, retrying may result in a probability of success.
>
> There are scenarios, albeit unusual, where any number of retries will not succeed. Therefore it is strongly recommended to limit the number of retries (2 or 3) to avoid an endless loop.

Data Validation The data change introduced by this transaction **did not pass validation** by commit handlers or data was incorrectly structured. The returned future will fail with a `DataValidationFailedException`. User

should not retry to create new transaction with same data, since it probably will fail again.

Example conflict of two transactions

This example illustrates two concurrent transactions, which derived from same initial state of data tree and proposes conflicting modifications.

```
WriteTransaction txA = broker.newWriteTransaction();
WriteTransaction txB = broker.newWriteTransaction();

txA.put(CONFIGURATION, PATH, A);
txB.put(CONFIGURATION, PATH, B);

CheckedFuture<?,?> futureA = txA.submit();
CheckedFuture<?,?> futureB = txB.submit();
```

- Updates `PATH` to value `A` using `txA`
- Updates `PATH` to value `B` using `txB`
- Seals & submits `txA`. The commit will be processed asynchronously and data tree will be updated to contain value `A` for `PATH`. The returned 'ListenableFuture' will complete successfully once state is applied to data tree.
- Seals & submits `txB`. Commit of `txB` will fail, because previous transaction also modified path in a concurrent way. The state introduced by `txB` will not be applied. The returned `ListenableFuture` will fail with `OptimisticLockFailedException` exception, which indicates that concurrent transaction prevented the submitted transaction from being applied.

Example asynchronous retry-loop

```
private void doWrite( final int tries ) {
    WriteTransaction writeTx = dataBroker.newWriteOnlyTransaction();

    MyDataObject data = ...;
    InstanceIdentifier<MyDataObject> path = ...;
    writeTx.put( LogicalDatastoreType.OPERATIONAL, path, data );

    Futures.addCallback( writeTx.submit(), new FutureCallback<Void>() {
        public void onSuccess( Void result ) {
            // succeeded
        }

        public void onFailure( Throwable t ) {
            if( t instanceof OptimisticLockFailedException && (( tries - 1 ) > 0)) {
                doWrite( tries - 1 );
            }
        }
      });
}
...
doWrite( 2 );
```

Concurrent change compatibility

There are several sets of changes which could be considered incompatible between two transactions which are derived from same initial state. Rules for conflict detection applies recursively for each subtree level.

Following table shows state changes and failures between two concurrent transactions, which are based on same initial state, `tx1` is submitted before `tx2`.

INFO: Following tables stores numeric values and shows data using `toString()` to simplify examples.

Initial state	tx1	tx2	Observable Result
Empty	`put(A,1)`	`put(A,2)`	`tx2` will fail, value of `A` is `1`
Empty	`put(A,1)`	`merge(A,2)`	value of `A` is `2`
Empty	`merge(A,1)`	`put(A,2)`	`tx2` will fail, value of `A` is `1`
Empty	`merge(A,1)`	`merge(A,2)`	`A` is `2`
A=0	`put(A,1)`	`put(A,2)`	`tx2` will fail, `A` is `1`
A=0	`put(A,1)`	`merge(A,2)`	`A` is `2`
A=0	`merge(A,1)`	`put(A,2)`	`tx2` will fail, value of `A` is `1`
A=0	`merge(A,1)`	`merge(A,2)`	`A` is `2`
A=0	`delete(A)`	`put(A,2)`	`tx2` will fail, `A` does not exists
A=0	`delete(A)`	`merge(A,2)`	`A` is `2`

Table: Concurrent change resolution for leaves and leaf-list items

Initial state	`tx1`	`tx2`	Result
Empty	put(TOP,[])	put(TOP,[])	`tx2` will fail, state is TOP=[]
Empty	put(TOP,[])	merge(TOP,[])	TOP=[]
Empty	put(TOP,[FOO=1])	put(TOP,[BAR=1])	`tx2` will fail, state is TOP=[FOO=1]
Empty	put(TOP,[FOO=1])	merge(TOP,[BAR=1])	TOP=[FOO=1,BAR=1]
Empty	merge(TOP,[FOO=1])	put(TOP,[BAR=1])	`tx2` will fail, state is TOP=[FOO=1]
Empty	merge(TOP,[FOO=1])	merge(TOP,[BAR=1])	TOP=[FOO=1,BAR=1]
TOP=[]	put(TOP,[FOO=1])	put(TOP,[BAR=1])	`tx2` will fail, state is TOP=[FOO=1]
TOP=[]	put(TOP,[FOO=1])	merge(TOP,[BAR=1])	state is TOP=[FOO=1,BAR=1]
TOP=[]	merge(TOP,[FOO=1])	put(TOP,[BAR=1])	`tx2` will fail, state is TOP=[FOO=1]
TOP=[]	merge(TOP,[FOO=1])	merge(TOP,[BAR=1])	state is TOP=[FOO=1,BAR=1]
TOP=[]	delete(TOP)	put(TOP,[BAR=1])	`tx2` will fail, state is empty store
TOP=[]	delete(TOP)	merge(TOP,[BAR=1])	state is TOP=[BAR=1]
TOP=[]	put(TOP/FOO,1)	put(TOP/BAR,1])	state is TOP=[FOO=1,BAR=1]
TOP=[]	put(TOP/FOO,1)	merge(TOP/BAR,1)	state is TOP=[FOO=1,BAR=1]
TOP=[]	merge(TOP/FOO,1)	put(TOP/BAR,1)	state is TOP=[FOO=1,BAR=1]
TOP=[]	merge(TOP/FOO,1)	merge(TOP/BAR,1)	state is TOP=[FOO=1,BAR=1]
TOP=[]	delete(TOP)	put(TOP/BAR,1)	`tx2` will fail, state is empty store
TOP=[]	delete(TOP)	merge(TOP/BAR,1]	`tx2` will fail, state is empty store
TOP=[FOO=1]	put(TOP/FOO,2)	put(TOP/BAR,1)	state is TOP=[FOO=2,BAR=1]
TOP=[FOO=1]	put(TOP/FOO,2)	merge(TOP/BAR,1)	state is TOP=[FOO=2,BAR=1]
TOP=[FOO=1]	merge(TOP/FOO,2)	put(TOP/BAR,1)	state is TOP=[FOO=2,BAR=1]
TOP=[FOO=1]	merge(TOP/FOO,2)	merge(TOP/BAR,1)	state is TOP=[FOO=2,BAR=1]
TOP=[FOO=1]	delete(TOP/FOO)	put(TOP/BAR,1)	state is TOP=[BAR=1]
TOP=[FOO=1]	delete(TOP/FOO)	merge(TOP/BAR,1]	state is TOP=[BAR=1]

Table: Concurrent change resolution for containers, lists, list items

MD-SAL RPC routing

The MD-SAL provides a way to deliver Remote Procedure Calls (RPCs) to a particular implementation based on content in the input as it is modeled in YANG. This part of the the RPC input is referred to as a **context reference**.

The MD-SAL does not dictate the name of the leaf which is used for this RPC routing, but provides necessary functionality for YANG model author to define their **context reference** in their model of RPCs.

MD-SAL routing behavior is modeled using following terminology and its application to YANG models:

Context Type Logical type of RPC routing. Context type is modeled as YANG `identity` and is referenced in model to provide scoping information.

Context Instance Conceptual location in data tree, which represents context in which RPC could be executed. Context instance usually represent logical point to which RPC execution is attached.

Context Reference Field of RPC input payload which contains Instance Identifier referencing **context instance** in which the RPC should be executed.

Modeling a routed RPC

In order to define routed RPCs, the YANG model author needs to declare (or reuse) a **context type**, set of possible **context instances** and finally RPCs which will contain **context reference** on which they will be routed.

Declaring a routing context type

This declares an identity named `node-context`, which is used as marker for node-based routing and is used in other places to reference that routing type.

Declaring possible context instances

In order to define possible values of **context instances** for routed RPCs, we need to model that set accordingly using `context-instance` extension from the `yang-ext` model.

The statement `ext:context-instance "node-context";` marks any element of the `list node` as a possible valid **context instance** in `node-context` based routing.

Note: The existence of a **context instance** node in operational or config data tree is not strongly tied to existence of RPC implementation.

For most routed RPC models, there is relationship between the data present in operational data tree and RPC implementation availability, but this is not enforced by MD-SAL. This provides some flexibility for YANG model writers to better specify their routing model and requirements for implementations. Details when RPC implementations are available should be documented in YANG model.

If user invokes RPC with a **context instance** that has no registered implementation, the RPC invocation will fail with the exception `DOMRpcImplementationNotAvailableException`.

Declaring a routed RPC

To declare RPC to be routed based on `node-context` we need to add leaf of `instance-identifier` type (or type derived from `instance-identifier`) to the RPC and mark it as **context reference**.

This is achieved using YANG extension `context-reference` from `yang-ext` model on leaf, which will be used for RPC routing.

The statement `ext:context-reference "node-context"` marks `leaf node` as **context reference** of type `node-context`. The value of this leaf, will be used by the MD-SAL to select the particular RPC implementation that registered itself as the implementation of the RPC for particular **context instance**.

Using routed RPCs

From a user perspective (e.g. invoking RPCs) there is no difference between routed and non-routed RPCs. Routing information is just an additional leaf in RPC which must be populated.

Implementing a routed RPC

Implementation

Registering implementations

Implementations of a routed RPC (e.g., southbound plugins) will specify an instance-identifier for the **context reference** (in this case a node) for which they want to provide an implementation during registration. Consumers, e.g., those calling the RPC are required to specify that instance-identifier (in this case the identifier of a node) when invoking RPC.

Simple code which showcases that for add-flow via Binding-Aware APIs (RoutedServiceTest.java):

```
61  @Override
62  public void onSessionInitiated(ProviderContext session) {
63      assertNotNull(session);
64      firstReg = session.addRoutedRpcImplementation(SalFlowService.class,␣
↪salFlowService1);
65  }
```

Line 64: We are registering salFlowService1 as implementation of SalFlowService RPC

```
107  NodeRef nodeOne = createNodeRef("foo:node:1");
109  /**
110   * Provider 1 registers path of node 1
111   */
112  firstReg.registerPath(NodeContext.class, nodeOne);
```

Line 107: We are creating NodeRef (encapsulation of InstanceIdentifier) for “foo:node:1”.

Line 112: We register salFlowService1 as implementation for nodeOne.

The salFlowService1 will be executed only for RPCs which contains Instance Identifier for foo:node:1.

OpenDaylight Controller MD-SAL: RESTCONF

RESTCONF operations overview

RESTCONF allows access to datastores in the controller.
There are two datastores:

- Config: Contains data inserted via controller
- Operational: Contains other data

Note:

Each request must start with the URI /restconf.
RESTCONF listens on port 8080 for HTTP requests.

RESTCONF supports **OPTIONS**, **GET**, **PUT**, **POST**, and **DELETE** operations. Request and response data can either be in the XML or JSON format. XML structures according to yang are defined at: XML-YANG. JSON structures are defined at: JSON-YANG. Data in the request must have a correctly set **Content-Type** field in the http header with the allowed value of the media type. The media type of the requested data has to be set in the **Accept** field. Get the media types for each resource by calling the OPTIONS operation. Most of the paths of the pathsRestconf endpoints use Instance Identifier. `<identifier>` is used in the explanation of the operations.

<identifier>

- It must start with <moduleName>:<nodeName> where <moduleName> is a name of the module and <nodeName> is the name of a node in the module. It is sufficient to just use <nodeName> after <moduleName>:<nodeName>. Each <nodeName> has to be separated by /.
- <nodeName> can represent a data node which is a list or container yang built-in type. If the data node is a list, there must be defined keys of the list behind the data node name for example, <nodeName>/<valueOfKey1>/<valueOfKey2>.
- The format <moduleName>:<nodeName> has to be used in this case as well:
  Module A has node A1. Module B augments node A1 by adding node X. Module C augments node A1 by adding node X. For clarity, it has to be known which node is X (for example: C:X). For more details about encoding, see: RESTCONF 02 - Encoding YANG Instance Identifiers in the Request URI.

Mount point

A Node can be behind a mount point. In this case, the URI has to be in format <identifier>**/yang-ext:mount/**<identifier>. The first <identifier> is the path to a mount point and the second <identifier> is the path to a node behind the mount point. A URI can end in a mount point itself by using <identifier>**/yang-ext:mount**.
More information on how to actually use mountpoints is available at: OpenDaylight Controller:Config:Examples:Netconf.

HTTP methods

OPTIONS /restconf

- Returns the XML description of the resources with the required request and response media types in Web Application Description Language (WADL)

GET /restconf/config/<identifier>

- Returns a data node from the Config datastore.
- <identifier> points to a data node which must be retrieved.

GET /restconf/operational/<identifier>

- Returns the value of the data node from the Operational datastore.
- <identifier> points to a data node which must be retrieved.

PUT /restconf/config/<identifier>

- Updates or creates data in the Config datastore and returns the state about success.
- <identifier> points to a data node which must be stored.

Example:

```
PUT http://<controllerIP>:8080/restconf/config/module1:foo/bar
Content-Type: applicaton/xml
<bar>
  ...
</bar>
```

Example with mount point:

```
PUT http://<controllerIP>:8080/restconf/config/module1:foo1/foo2/yang-ext:mount/
↪module2:foo/bar
Content-Type: applicaton/xml
<bar>
  ...
</bar>
```

POST /restconf/config

- Creates the data if it does not exist

For example:

```
POST URL: http://localhost:8080/restconf/config/
content-type: application/yang.data+json
JSON payload:
```

(continues on next page)

(continued from previous page)

```
  {
    "toaster:toaster" :
    {
      "toaster:toasterManufacturer" : "General Electric",
      "toaster:toasterModelNumber" : "123",
      "toaster:toasterStatus" : "up"
    }
  }
```

POST /restconf/config/<identifier>

- Creates the data if it does not exist in the Config datastore, and returns the state about success.
- <identifier> points to a data node where data must be stored.
- The root element of data must have the namespace (data are in XML) or module name (data are in JSON.)

Example:

```
POST http://<controllerIP>:8080/restconf/config/module1:foo
Content-Type: applicaton/xml/
<bar xmlns="module1namespace">
  ...
</bar>
```

Example with mount point:

```
http://<controllerIP>:8080/restconf/config/module1:foo1/foo2/yang-ext:mount/
↪module2:foo
Content-Type: applicaton/xml
<bar xmlns="module2namespace">
  ...
</bar>
```

POST /restconf/operations/<moduleName>:<rpcName>

- Invokes RPC.
- <moduleName>:<rpcName> - <moduleName> is the name of the module and <rpcName> is the name of the RPC in this module.
- The Root element of the data sent to RPC must have the name "input".
- The result can be the status code or the retrieved data having the root element "output".

Example:

```
POST http://<controllerIP>:8080/restconf/operations/module1:fooRpc
Content-Type: applicaton/xml
Accept: applicaton/xml
<input>
  ...
</input>

The answer from the server could be:
<output>
  ...
</output>
```

An example using a JSON payload:

```
POST http://localhost:8080/restconf/operations/toaster:make-toast
Content-Type: application/yang.data+json
{
  "input" :
  {
     "toaster:toasterDoneness" : "10",
     "toaster:toasterToastType":"wheat-bread"
  }
}
```

Note: Even though this is a default for the toasterToastType value in the yang, you still need to define it.

DELETE /restconf/config/<identifier>

- Removes the data node in the Config datastore and returns the state about success.
- <identifier> points to a data node which must be removed.

More information is available in the RESTCONF RFC.

How RESTCONF works

RESTCONF uses these base classes:

InstanceIdentifier Represents the path in the data tree

ConsumerSession Used for invoking RPCs

DataBrokerService Offers manipulation with transactions and reading data from the datastores

SchemaContext Holds information about yang modules

MountService Returns MountInstance based on the InstanceIdentifier pointing to a mount point

MountInstace Contains the SchemaContext behind the mount point

DataSchemaNode Provides information about the schema node

SimpleNode Possesses the same name as the schema node, and contains the value representing the data node value

CompositeNode Can contain CompositeNode-s and SimpleNode-s

GET in action

Figure 1 shows the GET operation with URI restconf/config/M:N where M is the module name, and N is the node name.

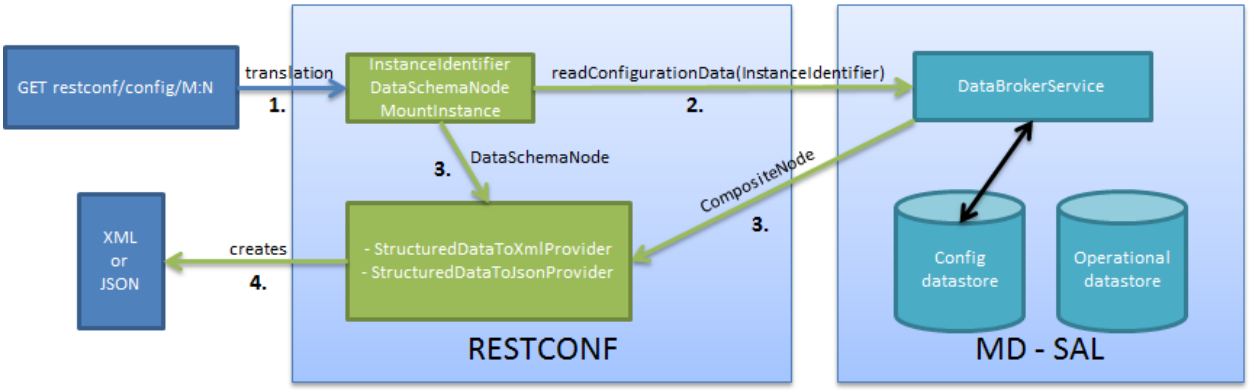

Fig. 15: Get

1. The requested URI is translated into the InstanceIdentifier which points to the data node. During this translation, the DataSchemaNode that conforms to the data node is obtained. If the data node is behind the mount point, the MountInstance is obtained as well.
2. RESTCONF asks for the value of the data node from DataBrokerService based on InstanceIdentifier.
3. DataBrokerService returns CompositeNode as data.
4. StructuredDataToXmlProvider or StructuredDataToJsonProvider is called based on the **Accept** field from the http request. These two providers can transform CompositeNode regarding DataSchemaNode to an XML or JSON document.
5. XML or JSON is returned as the answer on the request from the client.

PUT in action

Figure 2 shows the PUT operation with the URI restconf/config/M:N where M is the module name, and N is the node name. Data is sent in the request either in the XML or JSON format.

1. Input data is sent to JsonToCompositeNodeProvider or XmlToCompositeNodeProvider. The correct provider is selected based on the Content-Type field from the http request. These two providers can transform input data to CompositeNode. However, this CompositeNode does not contain enough information for transactions.
2. The requested URI is translated into InstanceIdentifier which points to the data node. DataSchemaNode conforming to the data node is obtained during this translation. If the data node is behind the mount point, the MountInstance is obtained as well.
3. CompositeNode can be normalized by adding additional information from DataSchemaNode.
4. RESTCONF begins the transaction, and puts CompositeNode with InstanceIdentifier into it. The response on the request from the client is the status code which depends on the result from the transaction.

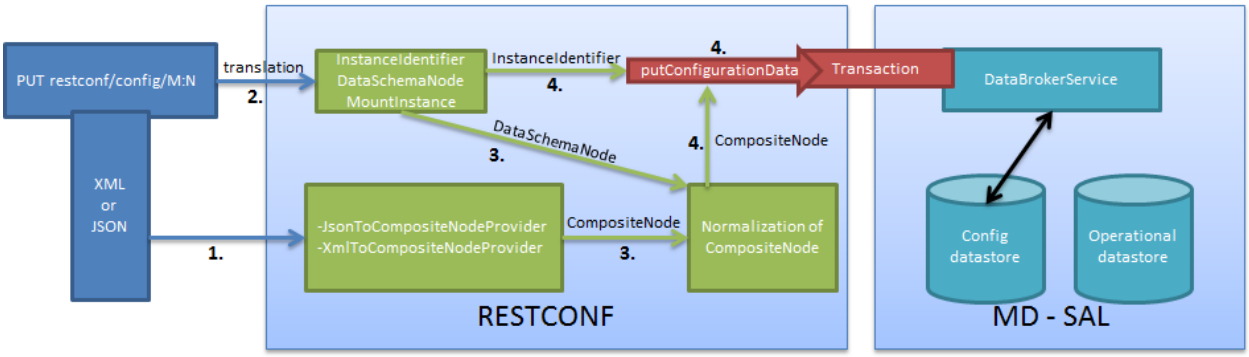

Fig. 16: Put

Something practical

1. Create a new flow on the switch openflow:1 in table 2.

HTTP request

```
Operation: POST
URI: http://192.168.11.1:8080/restconf/config/opendaylight-inventory:nodes/node/
↪openflow:1/table/2
Content-Type: application/xml
```

```
<?xml version="1.0" encoding="UTF-8" standalone="no"?>
<flow
    xmlns="urn:opendaylight:flow:inventory">
    <strict>false</strict>
    <instructions>
        <instruction>
            <order>1</order>
            <apply-actions>
                <action>
                  <order>1</order>
                    <flood-all-action/>
                </action>
            </apply-actions>
        </instruction>
    </instructions>
    <table_id>2</table_id>
    <id>111</id>
    <cookie_mask>10</cookie_mask>
    <out_port>10</out_port>
    <installHw>false</installHw>
    <out_group>2</out_group>
    <match>
        <ethernet-match>
            <ethernet-type>
                <type>2048</type>
            </ethernet-type>
        </ethernet-match>
```

(continues on next page)

(continued from previous page)

```
        <ipv4-destination>10.0.0.1/24</ipv4-destination>
    </match>
    <hard-timeout>0</hard-timeout>
    <cookie>10</cookie>
    <idle-timeout>0</idle-timeout>
    <flow-name>FooXf22</flow-name>
    <priority>2</priority>
    <barrier>false</barrier>
</flow>
```

HTTP response

```
Status: 204 No Content
```

1. Change *strict* to *true* in the previous flow.

HTTP request

```
Operation: PUT
URI: http://192.168.11.1:8080/restconf/config/opendaylight-inventory:nodes/node/
↪openflow:1/table/2/flow/111
Content-Type: application/xml
```

```
<?xml version="1.0" encoding="UTF-8" standalone="no"?>
<flow
    xmlns="urn:opendaylight:flow:inventory">
    <strict>true</strict>
    <instructions>
        <instruction>
            <order>1</order>
            <apply-actions>
                <action>
                  <order>1</order>
                    <flood-all-action/>
                </action>
            </apply-actions>
        </instruction>
    </instructions>
    <table_id>2</table_id>
    <id>111</id>
    <cookie_mask>10</cookie_mask>
    <out_port>10</out_port>
    <installHw>false</installHw>
    <out_group>2</out_group>
    <match>
        <ethernet-match>
            <ethernet-type>
                <type>2048</type>
            </ethernet-type>
        </ethernet-match>
```

(continues on next page)

(continued from previous page)

```
        <ipv4-destination>10.0.0.1/24</ipv4-destination>
    </match>
    <hard-timeout>0</hard-timeout>
    <cookie>10</cookie>
    <idle-timeout>0</idle-timeout>
    <flow-name>FooXf22</flow-name>
    <priority>2</priority>
    <barrier>false</barrier>
</flow>
```

HTTP response

```
Status: 200 OK
```

1. Show flow: check that *strict* is *true*.

HTTP request

```
Operation: GET
URI: http://192.168.11.1:8080/restconf/config/opendaylight-inventory:nodes/node/
↪openflow:1/table/2/flow/111
Accept: application/xml
```

HTTP response

```
Status: 200 OK
```

```
<?xml version="1.0" encoding="UTF-8" standalone="no"?>
<flow
    xmlns="urn:opendaylight:flow:inventory">
    <strict>true</strict>
    <instructions>
        <instruction>
            <order>1</order>
            <apply-actions>
                <action>
                  <order>1</order>
                    <flood-all-action/>
                </action>
            </apply-actions>
        </instruction>
    </instructions>
    <table_id>2</table_id>
    <id>111</id>
    <cookie_mask>10</cookie_mask>
    <out_port>10</out_port>
    <installHw>false</installHw>
```

(continues on next page)

(continued from previous page)

```
    <out_group>2</out_group>
    <match>
        <ethernet-match>
            <ethernet-type>
                <type>2048</type>
            </ethernet-type>
        </ethernet-match>
        <ipv4-destination>10.0.0.1/24</ipv4-destination>
    </match>
    <hard-timeout>0</hard-timeout>
    <cookie>10</cookie>
    <idle-timeout>0</idle-timeout>
    <flow-name>FooXf22</flow-name>
    <priority>2</priority>
    <barrier>false</barrier>
</flow>
```

1. Delete the flow created.

HTTP request

```
Operation: DELETE
URI: http://192.168.11.1:8080/restconf/config/opendaylight-inventory:nodes/node/
↪openflow:1/table/2/flow/111
```

HTTP response

```
Status: 200 OK
```

Websocket change event notification subscription tutorial

Subscribing to data change notifications makes it possible to obtain notifications about data manipulation (insert, change, delete) which are done on any specified **path** of any specified **datastore** with specific **scope**. In following examples *{odlAddress}* is address of server where ODL is running and *{odlPort}* is port on which OpenDaylight is running.

Websocket notifications subscription process

In this section we will learn what steps need to be taken in order to successfully subscribe to data change event notifications.

Create stream

In order to use event notifications you first need to call RPC that creates notification stream that you can later listen to. You need to provide three parameters to this RPC:

- **path**: data store path that you plan to listen to. You can register listener on containers, lists and leaves.
- **datastore**: data store type. *OPERATIONAL* or *CONFIGURATION*.
- **scope**: Represents scope of data change. Possible options are:
  - BASE: only changes directly to the data tree node specified in the path will be reported
  - ONE: changes to the node and to direct child nodes will be reported
  - SUBTREE: changes anywhere in the subtree starting at the node will be reported

The RPC to create the stream can be invoked via RESTCONF like this:

- URI: http://{odlAddress}:{odlPort}/restconf/operations/sal-remote:create-data-change-event-subscription
- HEADER: Content-Type=application/json
- OPERATION: POST
- DATA:

```
{
    "input": {
        "path": "/toaster:toaster/toaster:toasterStatus",
        "sal-remote-augment:datastore": "OPERATIONAL",
        "sal-remote-augment:scope": "ONE"
    }
}
```

The response should look something like this:

```
{
    "output": {
        "stream-name": "toaster:toaster/toaster:toasterStatus/datastore=CONFIGURATION/
↪scope=SUBTREE"
    }
}
```

stream-name is important because you will need to use it when you subscribe to the stream in the next step.

Note: Internally, this will create a new listener for *stream-name* if it did not already exist.

Subscribe to stream

In order to subscribe to stream and obtain WebSocket location you need to call *GET* on your stream path. The URI should generally be http://{odlAddress}:{odlPort}/restconf/streams/stream/{streamName}, where *{streamName}* is the *stream-name* parameter contained in response from *create-data-change-event-subscription* RPC from the previous step.

- URI: http://{odlAddress}:{odlPort}/restconf/streams/stream/toaster:toaster/datastore=CONFIGURATION/scope=SUBTREE
- OPERATION: GET

The expected response status is 200 OK and response body should be empty. You will get your WebSocket location from **Location** header of response. For example in our particular toaster example location header would have this value: *ws://{odlAddress}:8185/toaster:toaster/datastore=CONFIGURATION/scope=SUBTREE*

Note: During this phase there is an internal check for to see if a listener for the *stream-name* from the URI exists. If not, new a new listener is registered with the DOM data broker.

Receive notifications

You should now have a data change notification stream created and have location of a WebSocket. You can use this WebSocket to listen to data change notifications. To listen to notifications you can use a JavaScript client or if you are using chrome browser you can use the Simple WebSocket Client.

Also, for testing purposes, there is simple Java application named WebSocketClient. The application is placed in the *-sal-rest-connector-classes.class* project. It accepts a WebSocket URI as and input parameter. After starting the utility (WebSocketClient class directly in Eclipse/InteliJ Idea) received notifications should be displayed in console.

Notifications are always in XML format and look like this:

```
<notification xmlns="urn:ietf:params:xml:ns:netconf:notification:1.0">
    <eventTime>2014-09-11T09:58:23+02:00</eventTime>
    <data-changed-notification xmlns=
↪"urn:opendaylight:params:xml:ns:yang:controller:md:sal:remote">
        <data-change-event>
            <path xmlns:meae="http://netconfcentral.org/ns/toaster">/meae:toaster</
↪path>
            <operation>updated</operation>
            <data>
               <!-- updated data -->
            </data>
        </data-change-event>
    </data-changed-notification>
</notification>
```

Example use case

The typical use case is listening to data change events to update web page data in real-time. In this tutorial we will be using toaster as the base.

When you call *make-toast* RPC, it sets *toasterStatus* to "down" to reflect that the toaster is busy making toast. When it finishes, *toasterStatus* is set to "up" again. We will listen to this toaster status changes in data store and will reflect it on our web page in real-time thanks to WebSocket data change notification.

Simple javascript client implementation

We will create simple JavaScript web application that will listen updates on *toasterStatus* leaf and update some element of our web page according to new toaster status state.

Create stream

First you need to create stream that you are planing to subscribe to. This can be achieved by invoking "create-data-change-event-subscription" RPC on RESTCONF via AJAX request. You need to provide data store **path** that you plan to listen on, **data store type** and **scope**. If the request is successful you can extract the **stream-name** from the

response and use that to subscribe to the newly created stream. The *{username}* and *{password}* fields represent your credentials that you use to connect to OpenDaylight via RESTCONF:

Note: The default user name and password are "admin".

```
function createStream() {
    $.ajax(
        {
            url: 'http://{odlAddress}:{odlPort}/restconf/operations/sal-remote:create-
↪data-change-event-subscription',
            type: 'POST',
            headers: {
              'Authorization': 'Basic ' + btoa('{username}:{password}'),
              'Content-Type': 'application/json'
            },
            data: JSON.stringify(
                {
                    'input': {
                        'path': '/toaster:toaster/toaster:toasterStatus',
                        'sal-remote-augment:datastore': 'OPERATIONAL',
                        'sal-remote-augment:scope': 'ONE'
                    }
                }
            )
        }).done(function (data) {
            // this function will be called when ajax call is executed successfully
            subscribeToStream(data.output['stream-name']);
        }).fail(function (data) {
            // this function will be called when ajax call fails
            console.log("Create stream call unsuccessful");
        })
}
```

Subscribe to stream

The Next step is to subscribe to the stream. To subscribe to the stream you need to call *GET* on *http://{odlAddress}:{odlPort}/restconf/streams/stream/{stream-name}*. If the call is successful, you get WebSocket address for this stream in **Location** parameter inside response header. You can get response header by calling *getResponseHeader(*Location)** on HttpRequest object inside *done()* function call:

```
function subscribeToStream(streamName) {
    $.ajax(
        {
            url: 'http://{odlAddress}:{odlPort}/restconf/streams/stream/' +␣
↪streamName;
            type: 'GET',
            headers: {
              'Authorization': 'Basic ' + btoa('{username}:{password}'),
            }
        }
    ).done(function (data, textStatus, httpReq) {
        // we need function that has http request object parameter in order to access␣
↪response headers.
        listenToNotifications(httpReq.getResponseHeader('Location'));
```

(continues on next page)

(continued from previous page)

```
    }).fail(function (data) {
        console.log("Subscribe to stream call unsuccessful");
    });
}
```

Receive notifications

Once you got WebSocket server location you can now connect to it and start receiving data change events. You need to define functions that will handle events on WebSocket. In order to process incoming events from OpenDaylight you need to provide a function that will handle *onmessage* events. The function must have one parameter that represents the received event object. The event data will be stored in *event.data*. The data will be in an XML format that you can then easily parse using jQuery.

```
function listenToNotifications(socketLocation) {
    try {
        var notificatinSocket = new WebSocket(socketLocation);

        notificatinSocket.onmessage = function (event) {
            // we process our received event here
            console.log('Received toaster data change event.');
            $($.parseXML(event.data)).find('data-change-event').each(
                function (index) {
                    var operation = $(this).find('operation').text();
                    if (operation == 'updated') {
                        // toaster status was updated so we call function that gets␣
↪the value of toasterStatus leaf
                        updateToasterStatus();
                        return false;
                    }
                }
            );
        }
        notificatinSocket.onerror = function (error) {
            console.log("Socket error: " + error);
        }
        notificatinSocket.onopen = function (event) {
            console.log("Socket connection opened.");
        }
        notificatinSocket.onclose = function (event) {
            console.log("Socket connection closed.");
        }
        // if there is a problem on socket creation we get exception (i.e. when␣
↪socket address is incorrect)
    } catch(e) {
        alert("Error when creating WebSocket" + e );
    }
}
```

The *updateToasterStatus()* function represents function that calls *GET* on the path that was modified and sets toaster status in some web page element according to received data. After the WebSocket connection has been established you can test events by calling make-toast RPC via RESTCONF.

Note: for more information about WebSockets in JavaScript visit Writing WebSocket client applications

Config Subsystem

Overview

The Controller configuration operation has three stages:

- First, a Proposed configuration is created. Its target is to replace the old configuration.
- Second, the Proposed configuration is validated, and then committed. If it passes validation successfully, the Proposed configuration state will be changed to Validated.
- Finally, a Validated configuration can be Committed, and the affected modules can be reconfigured.

In fact, each configuration operation is wrapped in a transaction. Once a transaction is created, it can be configured, that is to say, a user can abort the transaction during this stage. After the transaction configuration is done, it is committed to the validation stage. In this stage, the validation procedures are invoked. If one or more validations fail, the transaction can be reconfigured. Upon success, the second phase commit is invoked. If this commit is successful, the transaction enters the last stage, committed. After that, the desired modules are reconfigured. If the second phase commit fails, it means that the transaction is unhealthy - basically, a new configuration instance creation failed, and the application can be in an inconsistent state.

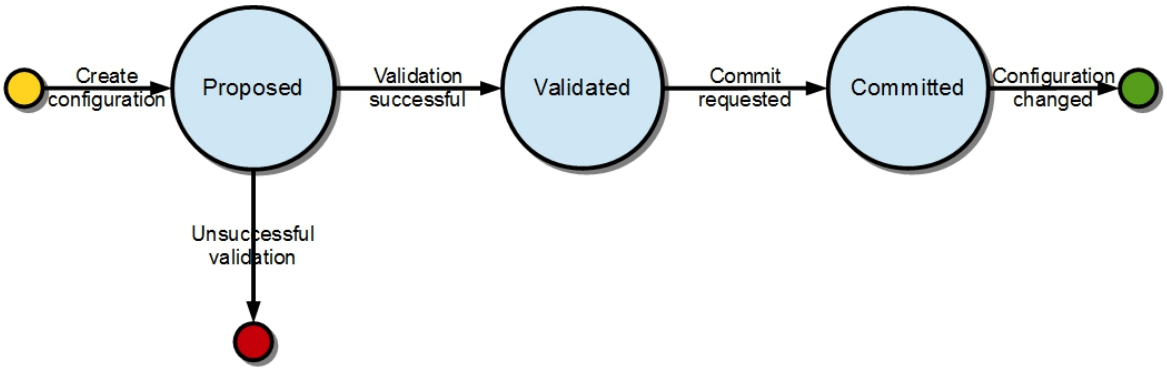

Fig. 17: Configuration states

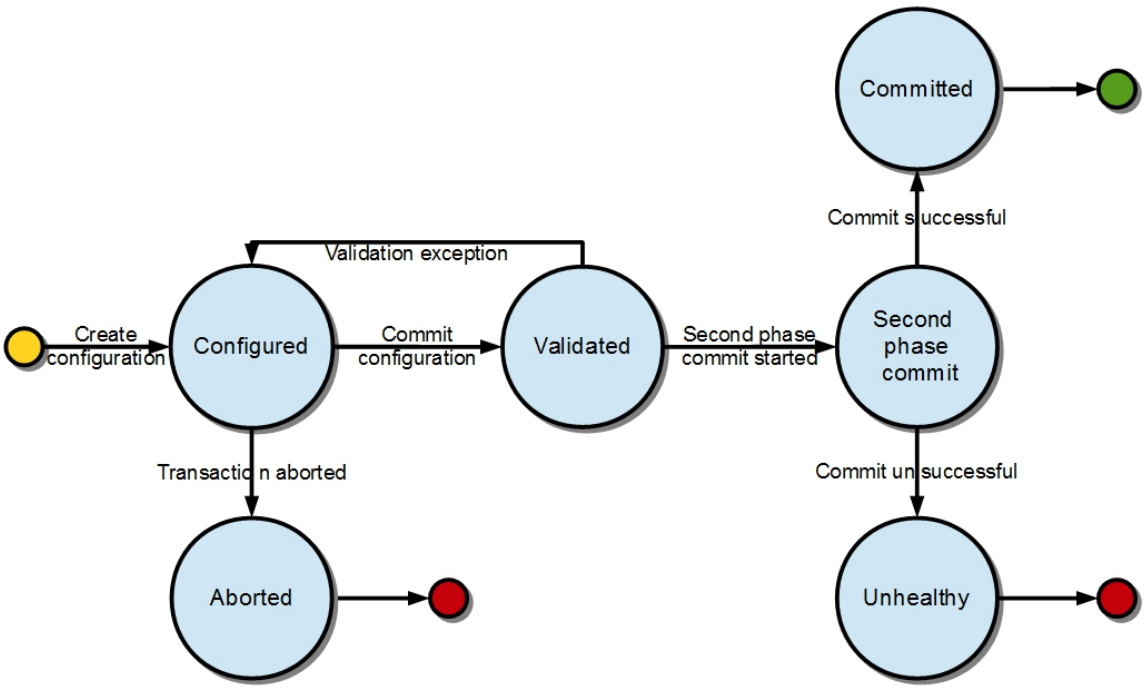

Fig. 18: Transaction states

Validation

To secure the consistency and safety of the new configuration and to avoid conflicts, the configuration validation process is necessary. Usually, validation checks the input parameters of a new configuration, and mostly verifies module-specific relationships. The validation procedure results in a decision on whether the proposed configuration is healthy.

Dependency resolver

Since there can be dependencies between modules, a change in a module configuration can affect the state of other modules. Therefore, we need to verify whether dependencies on other modules can be resolved. The Dependency Resolver acts in a manner similar to dependency injectors. Basically, a dependency tree is built.

APIs and SPIs

This section describes configuration system APIs and SPIs.

SPIs

Module org.opendaylight.controller.config.spi. Module is the common interface for all modules: every module must implement it. The module is designated to hold configuration attributes, validate them, and create instances of service based on the attributes. This instance must implement the AutoCloseable interface, owing to resources clean up. If the module was created from an already running instance, it contains an old instance of the module. A module can implement multiple services. If the module depends on other modules, setters need to be annotated with @RequireInterface.

Module creation

1. The module needs to be configured, set with all required attributes.
2. The module is then moved to the commit stage for validation. If the validation fails, the module attributes can be reconfigured. Otherwise, a new instance is either created, or an old instance is reconfigured. A module instance is identified by ModuleIdentifier, consisting of the factory name and instance name.

ModuleFactory org.opendaylight.controller.config.spi. The ModuleFactory interface must be implemented by each module factory.
A module factory can create a new module instance in two ways:

- From an existing module instance
- An entirely new instance
  ModuleFactory can also return default modules, useful for populating registry with already existing configurations. A module factory implementation must have a globally unique name.

APIs

ConfigRegistry	Represents functionality provided by a configuration transaction (create, destroy module, validate, or abort transaction).
ConfigTransactionController	Represents functionality for manipulating with configuration transactions (begin, commit config).
RuntimeBeanRegistratorAwareConfiBean	The module implementing this interface will receive RuntimeBeanRegistrator before getInstance is invoked.

Runtime APIs

RuntimeBean	Common interface for all runtime beans
RootRuntimeBeanRegistrator	Represents functionality for root runtime bean registration, which subsequently allows hierarchical registrations
HierarchicalRuntimeBeanRegistration	Represents functionality for runtime bean registration and unreregistration from hierarchy

JMX APIs

JMX API is purposed as a transition between the Client API and the JMX platform.

ConfigTransactionControllerMXBean	Extends ConfigTransactionController, executed by Jolokia clients on configuration transaction.
ConfigRegistryMXBean	Represents entry point of configuration management for MXBeans.
Object names	Object Name is the pattern used in JMX to locate JMX beans. It consists of domain and key properties (at least one key-value pair). Domain is defined as "org.opendaylight.controller". The only mandatory property is "type".

Use case scenarios

A few samples of successful and unsuccessful transaction scenarios follow:

Successful commit scenario

1. The user creates a transaction calling creteTransaction() method on ConfigRegistry.
2. ConfigRegisty creates a transaction controller, and registers the transaction as a new bean.
3. Runtime configurations are copied to the transaction. The user can create modules and set their attributes.
4. The configuration transaction is to be committed.
5. The validation process is performed.
6. After successful validation, the second phase commit begins.
7. Modules proposed to be destroyed are destroyed, and their service instances are closed.

8. Runtime beans are set to registrator.
9. The transaction controller invokes the method getInstance on each module.
10. The transaction is committed, and resources are either closed or released.

Validation failure scenario
The transaction is the same as the previous case until the validation process.

1. If validation fails, (that is to day, illegal input attributes values or dependency resolver failure), the validationException is thrown and exposed to the user.
2. The user can decide to reconfigure the transaction and commit again, or abort the current transaction.
3. On aborted transactions, TransactionController and JMXRegistrator are properly closed.
4. Unregistration event is sent to ConfigRegistry.

Default module instances

The configuration subsystem provides a way for modules to create default instances. A default instance is an instance of a module, that is created at the module bundle start-up (module becomes visible for configuration subsystem, for example, its bundle is activated in the OSGi environment). By default, no default instances are produced.

The default instance does not differ from instances created later in the module life-cycle. The only difference is that the configuration for the default instance cannot be provided by the configuration subsystem. The module has to acquire the configuration for these instances on its own. It can be acquired from, for example, environment variables. After the creation of a default instance, it acts as a regular instance and fully participates in the configuration subsystem (It can be reconfigured or deleted in following transactions.).

DIDM Developer Guide

Overview

The Device Identification and Driver Management (DIDM) project addresses the need to provide device-specific functionality. Device-specific functionality is code that performs a feature, and the code is knowledgeable of the capability and limitations of the device. For example, configuring VLANs and adjusting FlowMods are features, and there may be different implementations for different device types. Device-specific functionality is implemented as Device Drivers. Device Drivers need to be associated with the devices they can be used with. To determine this association requires the ability to identify the device type.

DIDM Architecture

The DIDM project creates the infrastructure to support the following functions:

- **Discovery** - Determination that a device exists in the controller management domain and connectivity to the device can be established. For devices that support the OpenFlow protocol, the existing discovery mechanism in OpenDaylight suffices. Devices that do not support OpenFlow will be discovered through manual means such as the operator entering device information via GUI or REST API.
- **Identification** – Determination of the device type.
- **Driver Registration** – Registration of Device Drivers as routed RPCs.

- **Synchronization** – Collection of device information, device configuration, and link (connection) information.
- **Data Models for Common Features** – Data models will be defined to perform common features such as VLAN configuration. For example, applications can configure a VLAN by writing the VLAN data to the data store as specified by the common data model.
- **RPCs for Common Features** – Configuring VLANs and adjusting FlowMods are example of features. RPCs will be defined that specify the APIs for these features. Drivers implement features for specific devices and support the APIs defined by the RPCs. There may be different Driver implementations for different device types.

Key APIs and Interfaces

FlowObjective API

Following are the list of the APIs to create the flow objectives to install the flow rule in OpenFlow switch in pipeline agnostic way. Currently these APIs are getting consumed by Atrium project.

Install the Forwarding Objective:

`http://<CONTROLLER-IP>:8181/restconf/operations/atrium-flow-objective:forward`

Install the Filter Objective

`http://<CONTROLLER-IP>:8181/restconf/operations/atrium-flow-objective:filter`

Install the Next Objective:

`http://<CONTROLLER-IP>:8181/restconf/operations/atrium-flow-objective:next`

Flow mod driver API

This release includes a flow mod driver for the HP 3800. This driver adjusts the flows and push the same to the device. This API takes the flow to be adjusted as input and displays the adjusted flow as output in the REST output container. Here is the REST API to adjust and push flows to HP 3800 device:

```
http://<CONTROLLER-IP:8181>/restconf/operations/openflow-feature:adjust-flow
```

Here is an example of an ARP flow and how it gets adjusted and pushed to device HP3800:

adjust-flow input.

```
<?xml version="1.0" encoding="UTF-8" standalone="no"?>
<input xmlns="urn:opendaylight:params:xml:ns:yang:didm:drivers:openflow"␣
↪xmlns:opendaylight-inventory="urn:opendaylight:inventory">
  <node>/opendaylight-inventory:nodes/opendaylight-inventory:node[opendaylight-
↪inventory:id='openflow:673249119553088']</node>
    <flow>
      <match>
        <ethernet-match>
            <ethernet-type>
                <type>2054</type>
            </ethernet-type>
        </ethernet-match>
      </match>
      <flags>SEND_FLOW_REM</flags>
      <priority>0</priority>
```

(continues on next page)

(continued from previous page)

```
      <flow-name>ARP_FLOW</flow-name>
      <instructions>
        <instruction>
            <order>0</order>
            <apply-actions>
                <action>
                    <order>0</order>
                    <output-action>
                        <output-node-connector>CONTROLLER</output-node-connector>
                        <max-length>65535</max-length>
                    </output-action>
                </action>
                <action>
                    <order>1</order>
                    <output-action>
                        <output-node-connector>NORMAL</output-node-connector>
                        <max-length>65535</max-length>
                    </output-action>
                </action>
            </apply-actions>
        </instruction>
      </instructions>
      <idle-timeout>180</idle-timeout>
      <hard-timeout>1800</hard-timeout>
      <cookie>10</cookie>
    </flow>
</input>
```

In the output, you can see that the table ID has been identified for the given flow and two flow mods are created as a result of adjustment. The first one is to catch ARP packets in Hardware table 100 with an action to goto table 200. The second flow mod is in table 200 with actions: output normal and output controller.

adjust-flow output.

```
{
  "output": {
    "flow": [
      {
        "idle-timeout": 180,
        "instructions": {
          "instruction": [
            {
              "order": 0,
              "apply-actions": {
                "action": [
                  {
                    "order": 1,
                    "output-action": {
                      "output-node-connector": "NORMAL",
                      "max-length": 65535
                    }
                  },
                  {
                    "order": 0,
                    "output-action": {
                      "output-node-connector": "CONTROLLER",
                      "max-length": 65535
```

(continues on next page)

(continued from previous page)

```
                    }
                  }
                ]
              }
            }
          ]
        },
        "strict": false,
        "table_id": 200,
        "flags": "SEND_FLOW_REM",
        "cookie": 10,
        "hard-timeout": 1800,
        "match": {
          "ethernet-match": {
            "ethernet-type": {
              "type": 2054
            }
          }
        },
        "flow-name": "ARP_FLOW",
        "priority": 0
      },
      {
        "idle-timeout": 180,
        "instructions": {
          "instruction": [
            {
              "order": 0,
              "go-to-table": {
                "table_id": 200
              }
            }
          ]
        },
        "strict": false,
        "table_id": 100,
        "flags": "SEND_FLOW_REM",
        "cookie": 10,
        "hard-timeout": 1800,
        "match": {},
        "flow-name": "ARP_FLOW",
        "priority": 0
      }
    ]
  }
}
```

API Reference Documentation

Go to http://${controller-ip}:8181/apidoc/explorer/index.html, and look under DIDM section to see all the available REST calls and tables

Distribution Version reporting

Overview

This section provides an overview of **odl-distribution-version** feature.

A remote user of OpenDaylight usually has access to RESTCONF and NETCONF northbound interfaces, but does not have access to the system OpenDaylight is running on. OpenDaylight has released multiple versions including Service Releases, and there are incompatible changes between them. In order to know which YANG modules to use, which bugs to expect and which workarounds to apply, such user would need to know the exact version of at least one OpenDaylight component.

There are indirect ways to deduce such version, but the direct way is enabled by odl-distribution-version feature. Administrator can specify version strings, which would be available to users via NETCONF, or via RESTCONF if OpenDaylight is configured to initiate NETCONF connection to its config subsystem northbound interface.

By default, users have write access to config subsystem, so they can add, modify or delete any version strings present there. Admins can only influence whether the feature is installed, and initial values.

Config subsystem is local only, not cluster aware, so each member reports versions independently. This is suitable for heterogeneous clusters. On homogeneous clusters, make sure you set and check every member.

Key APIs and Interfaces

Current implementation relies heavily on `config-parent` parent POM file from Controller project.

YANG model for config subsystem

Throughout this chapter, *model* denotes YANG module, and *module* denotes item in config subsystem module list.

Version functionality relies on config subsystem and its `config` YANG model. The YANG model `odl-distribution-version` adds an identity `odl-version` and augments `/config:modules/module/configuration` adding new case for `odl-version` type. This case contains single leaf `version`, which would hold the version string.

Config subsystem can hold multiple modules, the version string should contain version of OpenDaylight component corresponding to the module name. As this is pure metadata with no consequence on OpenDaylight behavior, there is no prescribed scheme for chosing config module names. But see the default configuration file for examples.

Java API

Each config module needs to come with java classes which override `customValidation()` and `createInstance()`. Version related modules have no impact on OpenDaylight internal behavior, so the methods return void and dummy closeable respectively, without any side effect.

Default config file

Initial version values are set via config file `odl-version.xml` which is created in `$KARAF_HOME/etc/opendaylight/karaf/` upon installation of `odl-distribution-version` feature. If admin wants to use different content, the file with desired content has to be created there before feature installation happens.

By default, the config file defines two config modules, named `odl-distribution-version` and `odl-odlparent-version`.

Currently the default version values are set to Maven property strings (as opposed to valid values), as the needed new functionality did not make it into Controller project in Boron. See Bug number 6003.

Karaf Feature

The `odl-distribution-version` feature is currently the only feature defined in feature repository of artifactId `features-distribution`, which is available (transitively) in OpenDaylight Karaf distribution.

RESTCONF usage

Opendaylight config subsystem NETCONF northbound is not made available just by installing `odl-distribution-version`, but most other feature installations would enable it. RESTCONF interfaces are enabled by installing `odl-restconf` feature, but that do not allow access to config subsystem by itself.

On single node deployments, installation of `odl-netconf-connector-ssh` is recommended, which would configure `controller-config` device and its MD-SAL mount point. See documentation for clustering on how to create similar devices for member modes, as `controller-config` name is not unique in that context.

Assuming single node deployment and user located on the same system, here is an example `curl` command accessing `odl-odlparent-version` config module:

```
curl 127.0.0.1:8181/restconf/config/network-topology:network-topology/topology/
↪topology-netconf/node/controller-config/yang-ext:mount/config:modules/module/odl-
↪distribution-version:odl-version/odl-odlparent-version
```

DLUX

Setup and Run

Required Technology Stack

- AngularJS (JavaScript client-side framework, http://www.angularjs.org)

Run DLUX

To turn on the DLUX UI, install DLUX core feature via running following command on the Karaf console -

```
feature:install odl-dlux-core
```

The above command will install odl-restconf and DLUX topology application internally, along with core DLUX components. Once this feature is successfully installed, access the UI at http://localhost:8181/index.html. The default credentials for login are admin/admin.

All the applications in DLUX are Karaf features. A user can install other dlux applications such as node and yang-ui from Karaf console using commands such as -

```
$ feature:install odl-dlux-node

$ feature:install odl-dlux-yangui
```

DLUX Modules

DLUX modules are the individual features such as nodes and topology. Each module has a defined structure and you can find all existing modules at https://github.com/opendaylight/dlux/tree/stable/boron/modules.

Module Structure

- module_folder
  - <module_name>.module.js
  - <module_name>.controller.js
  - <module_name>.services.js
  - <module_name>.directives.js
  - <module_name>.filter.js
  - index.tpl.html
  - <a_stylesheet>.css

Create New Module

Define the module

1. Create an empty maven project and create your module folder under src/main/resources.
2. Create an empty file with pattern <module_name>.module.js.
3. Next, you need to surround the angular module with a define function. This allows RequireJs to see our module.js files. The first argument is an array which contains all the module's dependencies. The second argument is a callback function, whose body contain the AngularJS code base. The function parameters correspond with the order of dependencies. Each dependency is injected into a parameter, if it is provided.
4. Finally, you will return the angular module to be able to inject it as a parameter in others modules.

For each new module, you must have at least these two dependencies :

- angularAMD : It's a wrapper around AngularJS to provide an AMD (Asynchronous Module Definition) support, which is used by RequireJs. For more information see the AMD documentation.
- app/core/core.services : This one is mandatory, if you want to add content in the navigation menu, the left bar or the top bar.

The following are not mandatory, but very often used.

- angular-ui-router : A library to provide URL routing.
- routingConfig : To set the level access to a page.

Your module.js file might look like this:

```
define(['angularAMD','app/routingConfig', 'angular-ui-router','app/core/core.services
↪'], function(ng) {
   var module = angular.module('app.a_module', ['ui.router.state', 'app.core']);
   // module configuration
   module.config(function() {
       [...]
```

(continues on next page)

(continued from previous page)

```
  });
  return module;
});
```

Set the register function

AngularJS allows lazy registration of a module's components such as controller, factory etc. Once you will install your application, DLUX will load your module javascript, but not your angular component during bootstrap phase. You have to register your angular components to make sure they are available at the runtime.

Here is how to register your module's component for lazy initialization -

```
module.config(function($compileProvider, $controllerProvider, $provide) {
   module.register = {
     controller : $controllerProvider.register,
     directive : $compileProvider.directive,
     factory : $provide.factory,
     service : $provide.service
   };
});
```

Set the route

The next step is to set up the route for your module. This part is also done in the configuration method of the module. We have to add **$stateProvider** as a parameter.

```
module.config(function($stateProvider) {
   var access = routingConfig.accessLevels;
   $stateProvider.state('main.module', {
     url: 'module',
     views : {
       'content' : {
         templateUrl: 'src/app/module/module.tpl.html',
         controller: 'ModuleCtrl'
       }
     }
   });
});
```

Adding element to the navigation menu

To be able to add item to the navigation menu, the module requires the **NavHelperProvider** parameter in the configuration method. **addToMenu** method in **NavMenuHelper** helper allows an item addition to the menu.

```
var module = angular.module('app.a_module', ['app.core']);
module.config(function(NavMenuHelper) {
    NavMenuHelper.addToMenu('myFirstModule', {
        "link" : "#/module/index",
        "active" : "module",
        "title" : "My First Module",
        "icon" : "icon-sitemap",
```

(continues on next page)

(continued from previous page)

```
        "page" : {
            "title" : "My First Module",
            "description" : "My first module"
        }
    });
});
```

The first parameter is an ID that refers to the level of your menu and the second is a object. For now, The ID parameter supports two levels of depth. If your ID looks like *rootNode.childNode*, the helper will look for a node named *rootNode* and it will append the *childNode* to it. If the root node doesn't exist, it will create it.

Link the AngularJS module's controller file

To include the module's controller file, you can use the NavHelperProvider. It contains a method that will load the given file.

```
[...]
   NavHelperProvider.addControllerUrl('<path_to_module_folder>/<module_name>.
↪controller');
```

This completes your module.js file.

Create the controller, factory, directive, etc

Creating the controller and other components is similar to the module.

- First, add the define method.
- Second, add the relative path to the module definition.
- Last, create your methods as you usually do it with AngularJS.

For example -

```
define(['<relative_path_to_module>/<module_name>.module'], function(module) {
   module.register.controller('ModuleCtrl', function($rootScope, $scope) {
   });
});
```

Add new application using DLUX modularity

DLUX works as a Karaf based UI platform, where you can create a new Karaf feature of your UI component and install that UI applications in DLUX using blueprint. This page will help you to create and load a new application for DLUX. You don't have to add new module in DLUX repository.

Add a new OSGi blueprint bundle

The OSGi Blueprint Container specification allows us to use dependency injection in our OSGi environment. Each DLUX application module registers itself via blueprint configuration. Each application will have its own blueprint.xml to place its configuration.

1. Create a maven project to place blueprint configuration. For reference, take a look at topology bundle, present at https://github.com/opendaylight/dlux/tree/stable/boron/bundles/topology. All the existing DLUX modules' configurations are available under bundles directory of DLUX code.
2. In pom.xml, you have to add a maven plugin to unpack your module code under generated-resources of this project. For reference, you can check pom.xml of dlux/bundles/topology at https://github.com/opendaylight/dlux/tree/stable/boron/bundles/topology. Your bundle will eventually get deployed in Karaf as feature, so your bundle should contain all your module code. If you want to combine module and bundle project, that should not be an issue either.
3. Create a blueprint.xml configuration file under src/main/resources/OSGI-INF/blueprint. Below is the content of the blueprint.xml taken from topology bundles's blueprint.xml. Any new application should create a blueprint.xml in following format -

```
<blueprint xmlns="http://www.osgi.org/xmlns/blueprint/v1.0.0">
    <reference id="httpService" availability="mandatory" activation="eager" interface=
↪"org.osgi.service.http.HttpService"/>
    <reference id="loader" availability="mandatory" activation="eager" interface="org.
↪opendaylight.dlux.loader.DluxModuleLoader"/>

    <bean id="bundle" init-method="initialize" destroy-method="clean" class="org.
↪opendaylight.dlux.loader.DluxModule">
      <property name="httpService" ref="httpService"/>
      <property name="loader" ref="loader"/>
      <property name="moduleName" value="topology "/>
      <property name="url" value="/src/app/topology"/>
      <property name="directory" value="/topology"/>
      <property name="requireJs" value="app/topology/topology.module"/>
      <property name="angularJs" value="app.topology"/>
      <property name="cssDependencies">
          <list>
              <value>http://yui.yahooapis.com/3.18.1/build/cssreset/cssreset-min.css</
↪value>
              <value>src/app/topology/topology-custom.css</value>
          </list>
      </property>
    </bean>
</blueprint>
```

In above configuration, there are two references with id httpService and loader. These two beans will already be initialized by dlux-core, so any new application can use them. Without these two bean references, a new application will not be able to register.

Next is the initialization of your application bean, which will be an instance of class org.opendaylight.dlux.loader.DluxModule. There are 5 properties that you should provide in this bean besides the references of httpService and loader. Lets talk about those bean properties in little more detail.

moduleName : Name of your module. This name should be unique in DLUX.

url: This is the url via which RequireJS in DLUX will try to load your module JS/HTML files. Also, this is the url that browser will use to load the static HTML, JS or CSS files. RequireJS in DLUX has a base path of **src**, so all the url should start with /src so RequireJS and the browser can correctly find the files.

directory: In your bundle's pom.xml, you unpack your module code. This is the directory where your actual static files will reside. The above mentioned url is registered with httpService, so when browser makes a call to that url, it will be redirected to the directory mentioned here. In the above example, all the topology files are present under /topology directory and the browser/RequireJS can access those files with uri /src/app/topology.

requireJS: This is the path to your RequireJS module. If you notice closely, you will see the initial path of RequireJS app/topology in the above example matches with the last part of url. This path will be be used by RequireJS. As

mentioned above, we have kept **src** as base path in RequireJS, that is the exact reason that url start with /src.

angularJS: name of your AngularJS module.

cssDependencies: If the application has any external/internal css dependencies, then those can be added here. If you create your own css files, just point to those css files here. Use the url path that you mentioned above, so the browser can find your css file.

OSGi understands blueprint.xml, once you will deploy your bundle in karaf (or you can create a new feature for your application), karaf will read your blueprint.xml and it will try to register your application with dlux. Once successful, if you refresh your dlux UI, you will see your application in left hand navigation bar of dlux.

Yang Utils

Yang Utils are used by UI to perform all CRUD operations. All of these utilities are present in yangutils.services.js file. It has following AngularJS factories -

- **arrayUtils** – defines functions for working with arrays.
- **pathUtils** – defines functions for working with xpath (paths to APIs and subAPIs). It divides xpath string to array of elements, so this array can be later used for search functions.
- **syncFact** – provides synchronization between requests to and from OpenDaylight when it's needed.
- **custFunct** – it is linked with apiConnector.createCustomFunctionalityApis in yangui controller in yangui.controller.js. That function makes it possible to create some custom function called by the click on button in index.tpl.html. All custom functions are stored in array and linked to specific subAPI. When particular subAPI is expanded and clicked, its inputs (linked root node with its child nodes) are displayed in the bottom part of the page and its buttons with custom functionality are displayed also.
- **reqBuilder** – Builds object in JSON format from input fields of the UI page. **Show Preview** button on Yang UI use this builder. This request is sent to OpenDaylight when button PUT or POST is clicked.
- **yinParser** – factory for reading .xml files of yang models and creating object hierarchy. Every statement from yang is represented by a node.
- **nodeWrapper** – adds functions to objects in tree hierarchy created with yinParser. These functions provide functionality for every type of node.
- **apiConnector** – the main functionality is filling the main structures and linking them. Structure of APIs and subAPIs which is two level array - first level is filled by main APIs, second level is filled by others sub APIs. Second main structure is array of root nodes, which are objects including root node and its children nodes. Linking these two structures is creating links between every subAPI (second level of APIs array) and its root node, which must be displayed like inputs when subAPI is expanded.
- **yangUtils** – some top level functions which are used by yangui controller for creating the main structures.

Fabric As A Service

FaaS (Fabric As A service) has two layers of APIs. We describe the top level API in the user guide. This document focuses on the Fabric level API and describes each API's semantics and example implementation. The second layer defines an abstraction layer called *''Fabric*" API. The idea is to abstract network into a topology formed by a collections of fabric objects other than varies of physical devices.Each Fabric object provides a collection of unified services.The top level API enables application developers or users to write applications to map high level model such as GBP, Intent etc... into a logical network model, while the lower level gives the application more control to individual fabric object level. More importantly the Fabric API is more like SP (Service Provider API) a fabric provider or vendor can implement the SPI based on its own Fabric technique such as TRILL, SPB etc ...

For how to use first level API operation, please refer to user guide for more details.

FaaS Architecture

FaaS Architecture is an 3 layered architecture, on the top is the FaaS Application layer, in the middle is the Fabric manager and at the bottom are different types of fabric objects. From bottom up, it is

Fabric and its controller (Fabric Controller) The Fabric object provides an abstraction of a homogeneous network or portion of the network and also has a built in Fabric controller which provides management plane and control plane for the fabric. The fabric controller implements the services required in Fabric Service and monitor and control the fabric operation.

Fabric Manager Fabric Manager manages all the fabric objects. also Fabric manager acts as a Unified Fabric Controller which provides inter-connect fabric control and configuration Also Fabric Manager is FaaS API service via Which FaaS user level logical network API (the top level API as mentioned previously) exposed and implemented.

FaaS renderer for GBP (Group Based Policy) FaaS renderer for GBP is an application of FaaS and provides the rendering service between GBP model and logical network model provided by Fabric Manager.

Fabric APIs and Interfaces

FaaS APIs have 4 groups as defined below

Fabric Provisioning API This set of APIs is used to create and remove Fabric Abstractions, in other words, those APIs is to provision the underlay networks and prepare to create overlay network(the logical network) on top of it.

Fabric Service API This set of APIs is used to create logical network over the Fabrics.

EndPoint API EndPoint API is used to bind a physical port which is the location of the attachment of an EndPoint happens or will happen.

OAM API Those APIs are for Operations, Administration and Maintenance purpose and In current release, OAM API is not implemented yet.

Fabric Provisioning API

- http://\protect\T1\textdollar\protect\T1\textbraceleftipaddress\protect\T1\textbraceright:8181/restconf/operations/fabric:compose-fabric
- http://\protect\T1\textdollar\protect\T1\textbraceleftipaddress\protect\T1\textbraceright:8181/restconf/operations/fabric:decompose-fabric
- http://\protect\T1\textdollar\protect\T1\textbraceleftipaddress\protect\T1\textbraceright:8181/restconf/operations/fabric:get-all-fabrics

Fabric Service API

- RESTCONF for creating Logical port, switch, router, routing entries and link. Among them, both switches and routers have ports. links connect ports.these 5 logical elements are basic building blocks of a logical network.
  - http://\protect\T1\textdollar\protect\T1\textbraceleftipaddress\protect\T1\textbraceright:8181/restconf/operations/fabric-service:create-logical-switch
  - http://\protect\T1\textdollar\protect\T1\textbraceleftipaddress\protect\T1\textbraceright:8181/restconf/operations/fabric-service:rm-logical-switch

  - http://\protect\T1\textdollar\protect\T1\textbraceleftipaddress\protect\T1\textbraceright:8181/restconf/operations/fabric-service:create-logical-router
  - http://\protect\T1\textdollar\protect\T1\textbraceleftipaddress\protect\T1\textbraceright:8181/restconf/operations/fabric-service:rm-logical-router
  - http://\protect\T1\textdollar\protect\T1\textbraceleftipaddress\protect\T1\textbraceright:8181/restconf/operations/fabric-service:add-static-route
  - http://\protect\T1\textdollar\protect\T1\textbraceleftipaddress\protect\T1\textbraceright:8181/restconf/operations/fabric-service:create-logic-port
  - http://\protect\T1\textdollar\protect\T1\textbraceleftipaddress\protect\T1\textbraceright:8181/restconf/operations/fabric-service:rm-logic-port
  - http://\protect\T1\textdollar\protect\T1\textbraceleftipaddress\protect\T1\textbraceright:8181/restconf/operations/fabric-service:create-gateway
  - http://\protect\T1\textdollar\protect\T1\textbraceleftipaddress\protect\T1\textbraceright:8181/restconf/operations/fabric-service:rm-gateway
  - http://\protect\T1\textdollar\protect\T1\textbraceleftipaddress\protect\T1\textbraceright:8181/restconf/operations/fabric-service:port-binding-logical-to-fabric
  - http://\protect\T1\textdollar\protect\T1\textbraceleftipaddress\protect\T1\textbraceright:8181/restconf/operations/fabric-service:port-binding-logical-to-device
  - http://\protect\T1\textdollar\protect\T1\textbraceleftipaddress\protect\T1\textbraceright:8181/restconf/operations/fabric-service:add-port-function
  - http://\protect\T1\textdollar\protect\T1\textbraceleftipaddress\protect\T1\textbraceright:8181/restconf/operations/fabric-service:add-acl
  - http://\protect\T1\textdollar\protect\T1\textbraceleftipaddress\protect\T1\textbraceright:8181/restconf/operations/fabric-service:del-acl

EndPoint API

The following APIs is to bind the physical ports to the logical ports on the logical switches:

- http://\protect\T1\textdollar\protect\T1\textbraceleftipaddress\protect\T1\textbraceright:8181/restconf/operations/fabric-endpoint:register-endpoint
- http://\protect\T1\textdollar\protect\T1\textbraceleftipaddress\protect\T1\textbraceright:8181/restconf/operations/fabric-endpoint:unregister-endpoint
- http://\protect\T1\textdollar\protect\T1\textbraceleftipaddress\protect\T1\textbraceright:8181/restconf/operations/fabric-endpoint:locate-endpoint

Others API

- http://\protect\T1\textdollar\protect\T1\textbraceleftipaddress\protect\T1\textbraceright:8181/restconf/operations/fabric-resource:create-fabric-port

API Reference Documentation

Go to http://\protect\T1\textdollar\protect\T1\textbraceleftipaddress\protect\T1\textbraceright:8181/restconf/apidoc/index.html and expand on *''FaaS*" related panel for more APIs.

Infrautils

Overview

Infrautils offer various utilities and infrastructures for other projects to use:

Counters Infrastructure

Create, update and output counters is a basic tool for debugging and generating statistics in any system. We have developed a counter infrastructure integrated into ODL which has already been successfully used with multiple products, and more recently in debugging and fixing the OpenFlow plugin/Java and LACP modules. Getting started with Counters

Async Infrastructure

The decision to split a service into one or more threads with asynchronous interactions between them is frequently dependent on constraints learned late in the development and even the deployment cycle. In order to allow flexibility in making these decisions we have developed an infrastructure which is configuration driven allowing agnostic code to be written under generic constrains which can then later be customized according to the required constraints. Getting started with Async

IoTDM Developer Guide

Overview

The Internet of Things Data Management (IoTDM) on OpenDaylight project is about developing a data-centric middleware that will act as a oneM2M compliant IoT Data Broker and enable authorized applications to retrieve IoT data uploaded by any device. The OpenDaylight platform is used to implement the oneM2M data store which models a hierarchical containment tree, where each node in the tree represents an oneM2M resource. Typically, IoT devices and applications interact with the resource tree over standard protocols such as CoAP, MQTT, and HTTP. Initially, the oneM2M resource tree is used by applications to retrieve data. Possible applications are inventory or device management systems or big data analytic systems designed to make sense of the collected data. But, at some point, applications will need to configure the devices. Features and tools will have to be provided to enable configuration of the devices based on applications responding to user input, network conditions, or some set of programmable rules or policies possibly triggered by the receipt of data collected from the devices. The OpenDaylight platform, with its rich unique cross-section of SDN capabilities, NFV, and now IoT device and application management, can be bundled with a targeted set of features and deployed anywhere in the network to give the network service provider ultimate control. Depending on the use case, the OpenDaylight IoT platform can be configured with only IoT data collection capabilities where it is deployed near the IoT devices and its footprint needs to be small, or it can be configured to run as a highly scaled up and out distributed cluster with IoT, SDN and NFV functions enabled and deployed in a high traffic data center.

oneM2M Architecture

The architecture provides a framework that enables the support of the oneM2M resource containment tree. The onem2m-core implements the MDSAL RPCs defined in the onem2m-api YANG files. These RPCs enable oneM2M resources to be created, read, updated, and deleted (CRUD), and also enables the management of subscriptions. When resources are CRUDed, the onem2m-notifier issues oneM2M notification events to interested subscribers. TS0001: oneM2M Functional Architecture and TS0004: oneM2M Service Layer Protocol are great reference documents to

learn details of oneM2M resource types, message flow, formats, and CRUD/N semantics. Both of these specifications can be found at http://onem2m.org/technical/published-documents

The oneM2M resource tree is modeled in YANG and essentially is a meta-model for the tree. The oneM2M wire protocols allow the resource tree to be constructed via HTTP or CoAP messages that populate nodes in the tree with resource specific attributes. Each oneM2M resource type has semantic behaviour associated with it. For example: a container resource has attributes which control quotas on how many and how big the collection of data or content instance objects that can exist below it in the tree. Depending on the resource type, the oneM2M core software implements and enforces the resource type specific rules to ensure a well-behaved resource tree.

The resource tree can be simultaneously accessed by many concurrent applications wishing to manage or access the tree, and also many devices can be reporting in new data or sensor readings into their appropriate place in the tree.

Key APIs and Interfaces

The API's to access the oneM2M datastore are well documented in TS0004 (referred above) found on onem2m.org

RESTCONF is available too but generally HTTP and CoAP are used to access the oneM2M data tree.

L2Switch Developer Guide

Overview

The L2Switch project provides Layer2 switch functionality.

L2Switch Architecture

- Packet Handler
  - Decodes the packets coming to the controller and dispatches them appropriately
- Loop Remover
  - Removes loops in the network
- Arp Handler
  - Handles the decoded ARP packets
- Address Tracker
  - Learns the Addresses (MAC and IP) of entities in the network
- Host Tracker
  - Tracks the locations of hosts in the network
- L2Switch Main
  - Installs flows on each switch based on network traffic

Key APIs and Interfaces

- Packet Handler
- Loop Remover
- Arp Handler

- Address Tracker
- Host Tracker
- L2Switch Main

Packet Dispatcher

Classes

- AbstractPacketDecoder
  - Defines the methods that all decoders must implement
- EthernetDecoder
  - The base decoder which decodes the packet into an Ethernet packet
- ArpDecoder, Ipv4Decoder, Ipv6Decoder
  - Decodes Ethernet packets into the either an ARP or IPv4 or IPv6 packet

Further development

There is a need for more decoders. A developer can write

- A decoder for another EtherType, i.e. LLDP.
- A higher layer decoder for the body of the IPv4 packet or IPv6 packet, i.e. TCP and UDP.

How to write a new decoder

- extends AbstractDecoder<A, B>
  - A refers to the notification that the new decoder consumes
  - B refers to the notification that the new decoder produces
- implements xPacketListener
  - The new decoder must specify which notification it is listening to
- canDecode method
  - This method should examine the consumed notification to see whether the new decoder can decode the contents of the packet
- decode method
  - This method does the actual decoding of the packet

Loop Remover

Classes

- **LoopRemoverModule**
  - Reads config subsystem value for *is-install-lldp-flow*
    - If *is-install-lldp-flow* is true, then an **InitialFlowWriter** is created

  - Creates and initializes the other LoopRemover classes
- **InitialFlowWriter**
  - Only created when *is-install-lldp-flow* is true
  - Installs a flow, which forwards all LLDP packets to the controller, on each switch
- **TopologyLinkDataChangeHandler**
  - Listens to data change events on the Topology tree
  - When these changes occur, it waits *graph-refresh-delay* seconds and then tells **NetworkGraphImpl** to update
  - Writes an STP (Spanning Tree Protocol) status of "forwarding" or "discarding" to each link in the Topology data tree
    - Forwarding links can forward packets.
    - Discarding links cannot forward packets.
- **NetworkGraphImpl**
  - Creates a loop-free graph of the network

Configuration

- graph-refresh-delay
  - Used in TopologyLinkDataChangeHandler
  - A higher value has the advantage of doing less graph updates, at the potential cost of losing some packets because the graph didn't update immediately.
  - A lower value has the advantage of handling network topology changes quicker, at the cost of doing more computation.
- is-install-lldp-flow
  - Used in LoopRemoverModule
  - "true" means a flow that sends all LLDP packets to the controller will be installed on each switch
  - "false" means this flow will not be installed
- lldp-flow-table-id
  - The LLDP flow will be installed on the specified flow table of each switch
- lldp-flow-priority
  - The LLDP flow will be installed with the specified priority
- lldp-flow-idle-timeout
  - The LLDP flow will timeout (removed from the switch) if the flow doesn't forward a packet for *x* seconds
- lldp-flow-hard-timeout
  - The LLDP flow will timeout (removed from the switch) after *x* seconds, regardless of how many packets it is forwarding

Further development

No suggestions at the moment.

Validating changes to Loop Remover

STP Status information is added to the Inventory data tree.

- A status of "forwarding" means the link is active and packets are flowing on it.
- A status of "discarding" means the link is inactive and packets are not sent over it.

The STP status of a link can be checked through a browser or a REST Client.

```
http://10.194.126.91:8080/restconf/operational/opendaylight-inventory:nodes/node/
↪openflow:1/node-connector/openflow:1:2
```

The STP status should still be there after changes are made.

Arp Handler

Classes

- **ArpHandlerModule**
  - Reads config subsystem value for *is-proactive-flood-mode*
    - If *is-proactive-flood-mode* is true, then a ProactiveFloodFlowWriter is created
    - If *is-proactive-flood-mode* is false, then an InitialFlowWriter is created
- **ProactiveFloodFlowWriter**
  - Only created when *is-proactive-flood-mode* is true
  - Installs a flood flow on each switch. With this flood flow, a packet that doesn't match any other flows will be flooded/broadcast from that switch.
- **InitialFlowWriter**
  - Only created when *is-proactive-flood-mode* is false
  - Installs a flow, which sends all ARP packets to the controller, on each switch
- **ArpPacketHandler**
  - Only created when *is-proactive-flood-mode* is false
  - Handles and processes the controller's incoming ARP packets
  - Uses **PacketDispatcher** to send the ARP packet back into the network
- **PacketDispatcher**
  - Only created when *is-proactive-flood-mode* is false
  - Sends packets out to the network
  - Uses **InventoryReader** to determine which node-connector to a send a packet on
- **InventoryReader**
  - Only created when *is-proactive-flood-mode* is false

  - Maintains a list of each switch's node-connectors

Configuration

- is-proactive-flood-mode
  - "true" means that flood flows will be installed on each switch. With this flood flow, each switch will flood a packet that doesn't match any other flows.
    * Advantage: Fewer packets are sent to the controller because those packets are flooded to the network.
    * Disadvantage: A lot of network traffic is generated.
  - "false" means the previously mentioned flood flows will not be installed. Instead an ARP flow will be installed on each switch that sends all ARP packets to the controller.
    * Advantage: Less network traffic is generated.
    * Disadvantage: The controller handles more packets (ARP requests & replies) and the ARP process takes longer than if there were flood flows.
- flood-flow-table-id
  - The flood flow will be installed on the specified flow table of each switch
- flood-flow-priority
  - The flood flow will be installed with the specified priority
- flood-flow-idle-timeout
  - The flood flow will timeout (removed from the switch) if the flow doesn't forward a packet for *x* seconds
- flood-flow-hard-timeout
  - The flood flow will timeout (removed from the switch) after *x* seconds, regardless of how many packets it is forwarding
- arp-flow-table-id
  - The ARP flow will be installed on the specified flow table of each switch
- arp-flow-priority
  - The ARP flow will be installed with the specified priority
- arp-flow-idle-timeout
  - The ARP flow will timeout (removed from the switch) if the flow doesn't forward a packet for *x* seconds
- arp-flow-hard-timeout
  - The ARP flow will timeout (removed from the switch) after *arp-flow-hard-timeout* seconds, regardless of how many packets it is forwarding

Further development

The **ProactiveFloodFlowWriter** needs to be improved. It does have the advantage of having less traffic come to the controller; however, it generates too much network traffic.

Address Tracker

Classes

- AddressTrackerModule
  - Reads config subsystem value for *observe-addresses-from*
  - If *observe-addresses-from* contains "arp", then an AddressObserverUsingArp is created
  - If *observe-addresses-from* contains "ipv4", then an AddressObserverUsingIpv4 is created
  - If *observe-addresses-from* contains "ipv6", then an AddressObserverUsingIpv6 is created
- AddressObserverUsingArp
  - Registers for ARP packet notifications
  - Uses **AddressObservationWriter** to write address observations from ARP packets
- AddressObserverUsingIpv4
  - Registers for IPv4 packet notifications
  - Uses **AddressObservationWriter** to write address observations from IPv4 packets
- AddressObserverUsingIpv6
  - Registers for IPv6 packet notifications
  - Uses **AddressObservationWriter** to write address observations from IPv6 packets
- AddressObservationWriter
  - Writes new Address Observations to the Inventory data tree
  - Updates existing Address Observations with updated "last seen" timestamps
    - Uses the *timestamp-update-intervval* configuration variable to determine whether or not to update

Configuration

- timestamp-update-interval
  - A last-seen timestamp is associated with each address. This last-seen timestamp will only be updated after *timestamp-update-interval* milliseconds.
  - A higher value has the advantage of performing less writes to the database.
  - A lower value has the advantage of knowing how fresh an address is.
- observe-addresses-from
  - IP and MAC addresses can be observed/learned from ARP, IPv4, and IPv6 packets. Set which packets to make these observations from.

Further development

Further improvements can be made to the **AddressObservationWriter** so that it (1) doesn't make any unnecessary writes to the DB and (2) is optimized for multi-threaded environments.

Validating changes to Address Tracker

Address Observations are added to the Inventory data tree.

The Address Observations on a Node Connector can be checked through a browser or a REST Client.

```
http://10.194.126.91:8080/restconf/operational/opendaylight-inventory:nodes/node/
↪openflow:1/node-connector/openflow:1:1
```

The Address Observations should still be there after changes.

Developer's Guide for Host Tracker

Validationg changes to Host Tracker

Host information is added to the Topology data tree.

- Host address
- Attachment point (link) to a node/switch

This host information and attachment point information can be checked through a browser or a REST Client.

```
http://10.194.126.91:8080/restconf/operational/network-topology:network-topology/
↪topology/flow:1/
```

Host information should still be there after changes.

L2Switch Main

Classes

- L2SwitchMainModule
  - Reads config subsystem value for *is-install-dropall-flow*
    - If *is-install-dropall-flow* is true, then an **InitialFlowWriter** is created
  - Reads config subsystem value for *is-learning-only-mode*
    - If *is-learning-only-mode* is false, then a **ReactiveFlowWriter** is created
- InitialFlowWriter
  - Only created when *is-install-dropall-flow* is true
  - Installs a flow, which drops all packets, on each switch. This flow has low priority and means that packets that don't match any higher-priority flows will simply be dropped.
- ReactiveFlowWriter
  - Reacts to network traffic and installs MAC-to-MAC flows on switches. These flows have matches based on MAC source and MAC destination.
  - Uses **FlowWriterServiceImpl** to write these flows to the switches
- FlowWriterService / FlowWriterServiceImpl
  - Writes flows to switches

Configuration

- is-install-dropall-flow
  - "true" means a drop-all flow will be installed on each switch, so the default action will be to drop a packet instead of sending it to the controller
  - "false" means this flow will not be installed
- dropall-flow-table-id
  - The dropall flow will be installed on the specified flow table of each switch
  - This field is only relevant when "is-install-dropall-flow" is set to "true"
- dropall-flow-priority
  - The dropall flow will be installed with the specified priority
  - This field is only relevant when "is-install-dropall-flow" is set to "true"
- dropall-flow-idle-timeout
  - The dropall flow will timeout (removed from the switch) if the flow doesn't forward a packet for *x* seconds
  - This field is only relevant when "is-install-dropall-flow" is set to "true"
- dropall-flow-hard-timeout
  - The dropall flow will timeout (removed from the switch) after *x* seconds, regardless of how many packets it is forwarding
  - This field is only relevant when "is-install-dropall-flow" is set to "true"
- is-learning-only-mode
  - "true" means that the L2Switch will only be learning addresses. No additional flows to optimize network traffic will be installed.
  - "false" means that the L2Switch will react to network traffic and install flows on the switches to optimize traffic. Currently, MAC-to-MAC flows are installed.
- reactive-flow-table-id
  - The reactive flow will be installed on the specified flow table of each switch
  - This field is only relevant when "is-learning-only-mode" is set to "false"
- reactive-flow-priority
  - The reactive flow will be installed with the specified priority
  - This field is only relevant when "is-learning-only-mode" is set to "false"
- reactive-flow-idle-timeout
  - The reactive flow will timeout (removed from the switch) if the flow doesn't forward a packet for *x* seconds
  - This field is only relevant when "is-learning-only-mode" is set to "false"
- reactive-flow-hard-timeout
  - The reactive flow will timeout (removed from the switch) after *x* seconds, regardless of how many packets it is forwarding
  - This field is only relevant when "is-learning-only-mode" is set to "false"

Further development

The **ReactiveFlowWriter** needs to be improved to install the MAC-to-MAC flows faster. For the first ping, the ARP request and reply are successful. However, then the ping packets are sent out. The first ping packet is dropped sometimes because the MAC-to-MAC flow isn't installed quickly enough. The second, third, and following ping packets are successful though.

API Reference Documentation

Further documentation can be found by checking out the L2Switch project.

Checking out the L2Switch project

```
git clone https://git.opendaylight.org/gerrit/p/l2switch.git
```

The above command will create a directory called "l2switch" with the project.

Testing your changes to the L2Switch project

Running the L2Switch project

To run the base distribution, you can use the following command

```
./distribution/base/target/distributions-l2switch-base-0.1.0-SNAPSHOT-osgipackage/
↪opendaylight/run.sh
```

If you need additional resources, you can use these command line arguments:

```
-Xms1024m -Xmx2048m -XX:PermSize=512m -XX:MaxPermSize=1024m'
```

To run the karaf distribution, you can use the following command:

```
./distribution/karaf/target/assembly/bin/karaf
```

Create a network using mininet

```
sudo mn --controller=remote,ip=<Controller IP> --topo=linear,3 --switch ovsk,
↪protocols=OpenFlow13
sudo mn --controller=remote,ip=127.0.0.1 --topo=linear,3 --switch ovsk,
↪protocols=OpenFlow13
```

The above command will create a virtual network consisting of 3 switches. Each switch will connect to the controller located at the specified IP, i.e. 127.0.0.1

```
sudo mn --controller=remote,ip=127.0.0.1 --mac --topo=linear,3 --switch ovsk,
↪protocols=OpenFlow13
```

The above command has the "mac" option, which makes it easier to distinguish between Host MAC addresses and Switch MAC addresses.

Generating network traffic using mininet

```
h1 ping h2
```

The above command will cause host1 (h1) to ping host2 (h2)

```
pingall
```

pingall will cause each host to ping every other host.

Miscellaneous mininet commands

```
link s1 s2 down
```

This will bring the link between switch1 (s1) and switch2 (s2) down

```
link s1 s2 up
```

This will bring the link between switch1 (s1) and switch2 (s2) up

```
link s1 h1 down
```

This will bring the link between switch1 (s1) and host1 (h1) down

LACP Developer Guide

LACP Overview

The OpenDaylight LACP (Link Aggregation Control Protocol) project can be used to aggregate multiple links between OpenDaylight controlled network switches and LACP enabled legacy switches or hosts operating in active LACP mode.

OpenDaylight LACP passively negotiates automatic bundling of multiple links to form a single LAG (Link Aggregation Group). LAGs are realised in the OpenDaylight controlled switches using OpenFlow 1.3+ group table functionality.

LACP Architecture

- **inventory**
  - Maintains list of OpenDaylight controlled switches and port information
  - List of LAGs created and physical ports that are part of the LAG
  - Interacts with MD-SAL to update LACP related information
- **inventorylistener**
  - This module interacts with MD-SAL for receiving node/node-connector notifications
- **flow**
  - Programs the switch to punt LACP PDU (Protocol Data Unit) to controller
- **packethandler**

  - Receives and transmits LACP PDUs to the LACP enabled endpoint
  - Provides infrastructure services for group table programming

- **core**
  - Performs LACP state machine processing

How LAG programming is implemented

The LAG representing the aggregated multiple physical ports are realized in the OpenDaylight controlled switches by creating a group table entry (Group table supported from OpenFlow 1.3 onwards). The group table entry has a group type **Select** and action referring to the aggregated physical ports. Any data traffic to be sent out through the LAG can be sent through the **group entry** available for the LAG.

Suppose there are ports P1-P8 in a node. When LACP project is installed, a group table entry for handling broadcast traffic is automatically created on all the switches that have registered to the controller.

GroupID	GroupType	EgressPorts
<B'castgID>	ALL	P1,P2,...P8

Now, assume P1 & P2 are now part of LAG1. The group table would be programmed as follows:

GroupID	GroupType	EgressPorts
<B'castgID>	ALL	P3,P4,...P8
<LAG1>	SELECT	P1,P2

When a second LAG, LAG2, is formed with ports P3 and P4,

GroupID	GroupType	EgressPorts
<B'castgID>	ALL	P5,P6,...P8
<LAG1>	SELECT	P1,P2
<LAG2>	SELECT	P3,P4

How applications can program OpenFlow flows using LACP-created LAG groups

OpenDaylight controller modules can get the information of LAG by listening/querying the LACP Aggregator datastore.

When any application receives packets, it can check, if the ingress port is part of a LAG by verifying the LAG Aggregator reference (lacp-agg-ref) for the source nodeConnector that OpenFlow plugin provides.

When applications want to add flows to egress out of the LAG, they must use the group entry corresponding to the LAG.

From the above example, for a flow to egress out of LAG1,

add-flow eth_type=<xxxx>,ip_dst=<x.x.x.x>,actions=output:<LAG1>

Similarly, when applications want traffic to be broadcasted, they should use the group table entries **<B'castgID>,<LAG1>,<LAG2>** in output action.

For all applications, the group table information is accessible from LACP Aggregator datastore.

NATApp Developer Guide

Overview

NATApp acts as a basic framework for providing NAT functionality to the SDN controller. One can use REST or Java APIs to enter global IP address into YANG Data Store which will be used by the odl-natapp-feature to map local IP to global IP addresses.

NATApp Architecture

NATApp listens on OpenFlow southbound interface for Packet_In messages. The application parses the message for header information. If the received message has a local IP address the application installs rules on the OpenFlow switch for network address translation from local to global IP addresses. NATApp has NATPacketHandler class that implements the PacketProcessing interface to override the OnPacketReceived notification by which the application is notified of Packet_In messages.

NATApp is implemented with the help of a few java classes.

1. NATPacketHandler
   - Receives Packet_In messages coming to the controller and process them appropriately
2. NATPacketParsing
   - Decodes Packet_In messages for packet header information (L2, L3 & L4 information)
3. NATInventoryUtility
   - Decodes Packet_In messages for OpenFlow Switch and Port information
4. NATFlowBuilder
   - Creates NAT flow rules at the OpenFlow Switch
5. NATYangStore
   - Reads Global IP entered by user and maps local IP to Global IP information
6. NATFlowHandler
   - Manages expired flows in the switch and frees up used global IP address for future natting.

Key APIs and Interfaces

1. RPC APIs
   - Static - Configure Static Natting Functionality
   - Dynamic - Configure Static Dynamic Functionality
   - PAT - Configure PAT Functionality
2. DataStore APIs
   - StaticNatIp - Configure floating IP addresses for Static Natting
   - StaticIpMapInfo - Mapped Information between floating and private IP addresses in Static Natting
   - DynamicNatIp - Configure floating IP addresses for Dynamic Natting
   - DynamicIpMapInfo - Mapped Information between floating and private IP addresses in Dynamic Natting
   - PatIp - Configure floating IP addresses for Port Address Translation

- PatIpMapInfo - Mapped Information between TCP Port numbers of floating IP and private IP addresses

3. Notification APIs
   - DynamicIPExhaustion - Exhaustion of Dynamic Global IP Addresses
   - PatOverConnection - More than 10 TCP or UDP connections from one private IP address

NEtwork MOdeling (NEMO)

Overview

The NEMO engine provides REST APIs to express intent, and manage it. With this northbound API, user could query what intents have been handled successfully, and what types have been predefined.

NEMO Architecture

In NEMO project, it provides three features facing developer.

- `odl-nemo-engine`: it is a whole model to handle intent.
- `odl-nemo-openflow-renderer`: it is a southbound render to translate intent to flow table in devices supporting for OpenFlow protocol.
- `odl-nemo-cli-render`: it is also a southbound render to translate intent into forwarding table in devices supporting for traditional protocol.

Key APIs and Interfaces

NEMO projects provide four basic REST methods for user to use.

- PUT: store the information expressed in NEMO model directly without handled by NEMO engine.
- POST: the information expressed in NEMO model will be handled by NEMO engine, and will be translated into southbound configuration.
- GET: obtain the data stored in data store.
- DELETE: delete the data in data store.

NEMO Intent API

NEMO provides several RPCs to handle user's intent. All RPCs use POST method.

- `http://{controller-ip}:8181/restconf/operations/nemo-intent:register-user`: a REST API to register a new user. It is the first and necessary step to express intent.
- `http://{controller-ip}:8181/restconf/operations/nemo-intent:transaction-begin`: a REST type to start a transaction. The intent exist in the transaction will be handled together.
- `http://{controller-ip}:8181/restconf/operations/nemo-intent:transaction-end`: a REST API to end a transaction. The intent exist in the transaction will be handled together.
- `http://{controller-ip}:8181/restconf/operations/nemo-intent:structure-style-nemo-updat` a REST API to create, import or update intent in a structure style, that is, user could express the structure of intent in json body.

- `http://{controller-ip}:8181/restconf/operations/nemo-intent:structure-style-nemo-delet` a REST API to delete intent in a structure style.
- `http://{controller-ip}:8181/restconf/operations/nemo-intent:language-style-nemo-reques` a REST API to create, import, update and delete intent in a language style, that is, user could express intent with NEMO script. On the other hand, with this interface, user could query which intent have been handled successfully.

API Reference Documentation

Go to `http://${IPADDRESS}:8181/apidoc/explorer/index.html`. User could see many useful APIs to deploy or query intent.

NETCONF Developer Guide

Note: Reading the NETCONF section in the User Guide is likely useful as it contains an overview of NETCONF in OpenDaylight and a how-to for spawning and configuring NETCONF connectors.

This chapter is recommended for application developers who want to interact with mounted NETCONF devices from their application code. It tries to demonstrate all the use cases from user guide with RESTCONF but now from the code level. One important difference would be the demonstration of NETCONF notifications and notification listeners. The notifications were not shown using RESTCONF because **RESTCONF does not support notifications from mounted NETCONF devices.**

Note: It may also be useful to read the generic OpenDaylight MD-SAL app development tutorial before diving into this chapter. This guide assumes awareness of basic OpenDaylight application development.

Sample app overview

All the examples presented here are implemented by a sample OpenDaylight application called **ncmount** in the `coretutorials` OpenDaylight project. It can be found on the github mirror of OpenDaylight's repositories:

- https://github.com/opendaylight/coretutorials/tree/stable/boron/ncmount

or checked out from the official OpenDaylight repository:

- https://git.opendaylight.org/gerrit/#/admin/projects/coretutorials

The application was built using the project startup maven archetype **and demonstrates how to:**

- preconfigure connectors to NETCONF devices
- retrieve MountPointService (registry of available mount points)
- listen and react to changing connection state of netconf-connector
- add custom device YANG models to the app and work with them
- read data from device in binding aware format (generated java APIs from provided YANG models)
- write data into device in binding aware format
- trigger and listen to NETCONF notifications in binding aware format

Detailed information about the structure of the application can be found at: https://wiki.opendaylight.org/view/Controller_Core_Functionality_Tutorials:Tutorials:Netconf_Mount

Note: The code in ncmount is fully **binding aware** (works with generated java APIs from provided YANG models). However it is also possible to perform the same operations in **binding independent** manner.

NcmountProvider

The NcmountProvider class (found in NcmountProvider.java) is the central point of the ncmount application and all the application logic is contained there. The following sections will detail its most interesting pieces.

Retrieve MountPointService

The MountPointService is a central registry of all available mount points in OpenDaylight. It is just another MD-SAL service and is available from the `session` attribute passed by `onSessionInitiated` callback:

```
@Override
public void onSessionInitiated(ProviderContext session) {
    LOG.info("NcmountProvider Session Initiated");

    // Get references to the data broker and mount service
    this.mountService = session.getSALService(MountPointService.class);

    ...

    }
}
```

Listen for connection state changes

It is important to know when a mount point appears, when it is fully connected and when it is disconnected or removed. The exact states of a mount point are:

- Connected
- Connecting
- Unable to connect

To receive this kind of information, an application has to register itself as a notification listener for the preconfigured netconf-topology subtree in MD-SAL's datastore. This can be performed in the `onSessionInitiated` callback as well:

```
@Override
public void onSessionInitiated(ProviderContext session) {

    ...

    this.dataBroker = session.getSALService(DataBroker.class);

    // Register ourselves as the REST API RPC implementation
    this.rpcReg = session.addRpcImplementation(NcmountService.class, this);
```

(continues on next page)

(continued from previous page)

```
    // Register ourselves as data change listener for changes on Netconf
    // nodes. Netconf nodes are accessed via "Netconf Topology" - a special
    // topology that is created by the system infrastructure. It contains
    // all Netconf nodes the Netconf connector knows about. NETCONF_TOPO_IID
    // is equivalent to the following URL:
    // .../restconf/operational/network-topology:network-topology/topology/topology-
↪netconf
    if (dataBroker != null) {
        this.dclReg = dataBroker.registerDataChangeListener(LogicalDatastoreType.
↪OPERATIONAL,
                NETCONF_TOPO_IID.child(Node.class),
                this,
                DataChangeScope.SUBTREE);
    }
}
```

The implementation of the callback from MD-SAL when the data change can be found in the `onDataChanged(AsyncDataChangeEvent<InstanceIdentifier<?>, DataObject> change)` callback of NcmountProvider class.

Reading data from the device

The first step when trying to interact with the device is to get the exact mount point instance (identified by an instance identifier) from the MountPointService:

```
@Override
public Future<RpcResult<ShowNodeOutput>> showNode(ShowNodeInput input) {
    LOG.info("showNode called, input {}", input);

    // Get the mount point for the specified node
    // Equivalent to '.../restconf/<config | operational>/opendaylight-
↪inventory:nodes/node/<node-name>/yang-ext:mount/'
    // Note that we can read both config and operational data from the same
    // mount point
    final Optional<MountPoint> xrNodeOptional = mountService.getMountPoint(NETCONF_
↪TOPO_IID
            .child(Node.class, new NodeKey(new NodeId(input.getNodeName()))));

    Preconditions.checkArgument(xrNodeOptional.isPresent(),
            "Unable to locate mountpoint: %s, not mounted yet or not configured",
            input.getNodeName());
    final MountPoint xrNode = xrNodeOptional.get();

    ....
}
```

Note: The triggering method in this case is called `showNode`. It is a YANG-defined RPC and NcmountProvider serves as an MD-SAL RPC implementation among other things. This means that `showNode` an be triggered using RESTCONF.

The next step is to retrieve an instance of the `DataBroker` API from the mount point and start a read transaction:

```
@Override
public Future<RpcResult<ShowNodeOutput>> showNode(ShowNodeInput input) {

    ...

    // Get the DataBroker for the mounted node
    final DataBroker xrNodeBroker = xrNode.getService(DataBroker.class).get();
    // Start a new read only transaction that we will use to read data
    // from the device
    final ReadOnlyTransaction xrNodeReadTx = xrNodeBroker.newReadOnlyTransaction();

    ...
}
```

Finally, it is possible to perform the read operation:

```
@Override
public Future<RpcResult<ShowNodeOutput>> showNode(ShowNodeInput input) {

    ...

    InstanceIdentifier<InterfaceConfigurations> iid =
            InstanceIdentifier.create(InterfaceConfigurations.class);

    Optional<InterfaceConfigurations> ifConfig;
    try {
        // Read from a transaction is asynchronous, but a simple
        // get/checkedGet makes the call synchronous
        ifConfig = xrNodeReadTx.read(LogicalDatastoreType.CONFIGURATION, iid).
↪checkedGet();
    } catch (ReadFailedException e) {
        throw new IllegalStateException("Unexpected error reading data from " + input.
↪getNodeName(), e);
    }

    ...
}
```

The instance identifier is used here again to specify a subtree to read from the device. At this point application can process the data as it sees fit. The ncmount app transforms the data into its own format and returns it from `showNode`.

Note: More information can be found in the source code of ncmount sample app + on wiki: https://wiki.opendaylight.org/view/Controller_Core_Functionality_Tutorials:Tutorials:Netconf_Mount

---

Network Intent Composition (NIC) Developer Guide

Overview

The Network Intent Composition (NIC) provides four features:

- odl-nic-core-hazelcast: Provides a distributed intent mapping service, implemented using hazelcast, that stores metadata needed by odl-nic-core feature.
- odl-nic-core-mdsal: Provides an intent rest API to external applications for CRUD operations on intents, conflict resolution and event handling. Uses MD-SAL as backend.

- odl-nic-console: Provides a karaf CLI extension for intent CRUD operations and mapping service operations.
- odl-nic-renderer-of - Generic OpenFlow Renderer.
- odl-nic-renderer-vtn - a feature that transforms an intent to a network modification using the VTN project
- odl-nic-renderer-gbp - a feature that transforms an intent to a network modification using the Group Policy project
- odl-nic-renderer-nemo - a feature that transforms an intent to a network modification using the NEMO project
- odl-nic-listeners - adds support for event listening. (depends on: odl-nic-renderer-of)
- odl-nic-neutron-integration - allow integration with openstack neutron to allow coexistence between existing neutron security rules and intents pushed by ODL applications.

Only a single renderer feature should be installed at a time for the Boron release.

odl-nic-core-mdsal XOR odl-nic-core-hazelcast

This feature supplies the base models for the Network Intent Composition (NIC) capability. This includes the definition of intent as well as the configuration and operational data trees.

This feature only provides an information model. The interface for NIC is to modify the information model via the configuraiton data tree, which will trigger the renderer to make the appropriate changes in the controlled network.

Installation

First you need to install one of the core installations:

```
feature:install odl-nic-core-service-mdsal odl-nic-console
```

OR

```
feature:install odl-nic-core-service-hazelcast odl-nic-console
```

Then pick a renderer:

```
feature:install odl-nic-listeners (will install odl-nic-renderer-of)
```

OR

```
feature:install odl-nic-renderer-vtn
```

OR

```
feature:install odl-nic-renderer-gbp
```

OR

```
feature:install odl-nic-renderer-nemo
```

REST Supported operations

POST / PUT (configuration)

This operations create instances of an intent in the configuration data tree and trigger the creation or modification of an intent.

GET (configuration / operational)

This operation lists all or fetches a single intent from the data tree.

DELETE (configuration)

This operation will cause an intent to be removed from the system and trigger any configuration changes on the network rendered from this intent to be removed.

odl-nic-cli user guide

This feature provides karaf console CLI command to manipulate the intent data model. The CLI essentailly invokes the equivalent data operations.

intent:add

Creates a new intent in the configuration data tree

```
DESCRIPTION
        intent:add

    Adds an intent to the controller.

Examples: --actions [ALLOW] --from <subject> --to <subject>
          --actions [BLOCK] --from <subject>

SYNTAX
        intent:add [options]

OPTIONS
        -a, --actions
                Action to be performed.
                -a / --actions BLOCK/ALLOW
                (defaults to [BLOCK])
        --help
                Display this help message
        -t, --to
                Second Subject.
                -t / --to <subject>
                (defaults to any)
        -f, --from
                First subject.
                -f / --from <subject>
                (defaults to any)
```

intent:delete

Removes an existing intent from the system

```
DESCRIPTION
        intent:remove

    Removes an intent from the controller.

SYNTAX
        intent:remove id

ARGUMENTS
        id  Intent Id
```

intent:list

Lists all the intents in the system

```
DESCRIPTION
        intent:list

    Lists all intents in the controller.

SYNTAX
        intent:list [options]

OPTIONS
        -c, --config
                List Configuration Data (optional).
                -c / --config <ENTER>
        --help
                Display this help message
```

intent:show

Displays the details of a single intent

```
DESCRIPTION
        intent:show

    Shows detailed information about an intent.

SYNTAX
        intent:show id

ARGUMENTS
        id  Intent Id
```

intent:map

List/Add/Delete current state from/to the mapping service.

```
DESCRIPTION
        intent:map

        List/Add/Delete current state from/to the mapping service.

SYNTAX
        intent:map [options]

         Examples: --list, -l [ENTER], to retrieve all keys.
                   --add-key <key> [ENTER], to add a new key with empty contents.
                   --del-key <key> [ENTER], to remove a key with it's values."
                   --add-key <key> --value [<value 1>, <value 2>, ...] [ENTER],
                     to add a new key with some values (json format).
OPTIONS
       --help
           Display this help message
       -l, --list
           List values associated with a particular key.
       -l / --filter <regular expression> [ENTER]
       --add-key
           Adds a new key to the mapping service.
       --add-key <key name> [ENTER]
       --value
           Specifies which value should be added/delete from the mapping service.
       --value "key=>value"... --value "key=>value" [ENTER]
           (defaults to [])
       --del-key
           Deletes a key from the mapping service.
       --del-key <key name> [ENTER]
```

Sample Use case: MPLS

Description

The scope of this use-case is to add MPLS intents between two MPLS endpoints. The use-case tries to address the real-world scenario illustrated in the diagram below:

where PE (Provider Edge) and P (Provider) switches are managed by OpenDaylight. In NIC's terminology the endpoints are the PE switches. There could be many P switches between the PEs.

In order for NIC to recognize endpoints as MPLS endpoints, the user is expected to add mapping information about the PE switches to NIC's mapping service to include the below properties:

1. MPLS Label to identify a PE
2. IPv4 Prefix for the customer site that are connected to a PE
3. Switch-Port: Ingress (or Egress) for source (or Destination) endpoint of the source (or Destination) PE

An intent:add between two MPLS endpoints renders OpenFlow rules for: 1. push/pop labels to the MPLS endpoint nodes after an IPv4 Prefix match. 2. forward to port rule after MPLS label match to all the switches that form the shortest path between the endpoints (calculated using Dijkstra algorithm).

Additionally, we have also added constraints to Intent model for protection and failover mechanism to ensure end-to-end connectivity between endpoints. By specifying these constraints to intent:add the use-case aims to reduces the risk of connectivity failure due to a single link or port down event on a forwarding device.

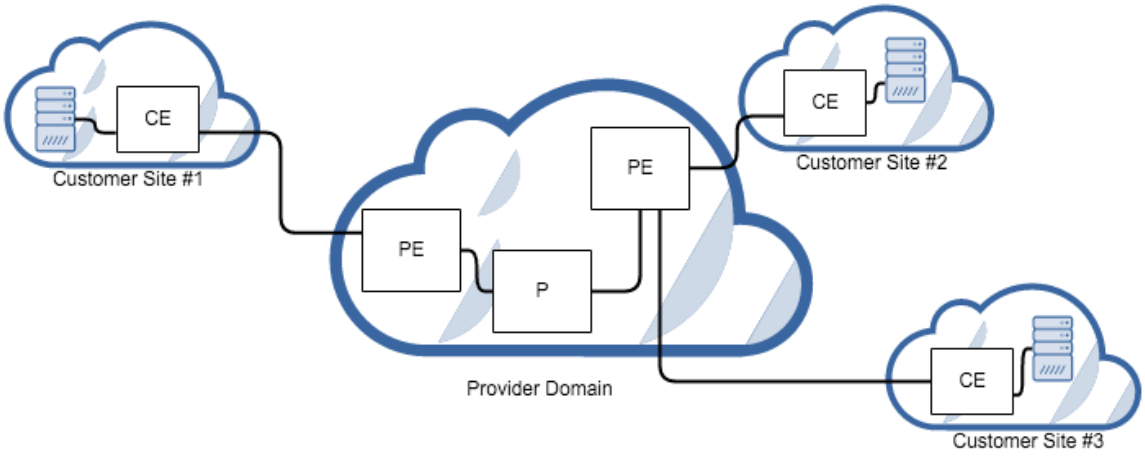

Fig. 19: MPLS VPN Service Diagram

- Protection constraint: Constraint that requires an end-to-end connectivity to be protected by providing redundant paths.
- Failover constraint: Constraint that specifies the type of failover implementation. slow-reroute: Uses disjoint path calculation algorithms like Suurballe to provide alternate end-to-end routes. fast-reroute: Uses failure detection feature in hardware forwarding device through OF group table features (Future plans) When no constraint is requested by the user we default to offering a since end-to-end route using Dijkstra shortest path.

How to use it?

1. Start Karaf and install related features:

```
feature:install odl-nic-core-service-mdsal odl-nic-core odl-nic-console odl-nic-
↪listeners
feature:install odl-dlux-all odl-dlux-core odl-dlux-yangui odl-dlux-yangvisualizer
```

2. Start mininet topology and verify in DLUX Topology page for the nodes and link.

```
mn --controller=remote,ip=$CONTROLLER_IP --custom ~/shortest_path.py --topo␣
↪shortest_path --switch ovsk,protocols=OpenFlow13
```

```
cat shortest.py -->
from mininet.topo import Topo
from mininet.cli import CLI
from mininet.net import Mininet
from mininet.link import TCLink
from mininet.util import irange,dumpNodeConnections
from mininet.log import setLogLevel
```

```
class Fast_Failover_Demo_Topo(Topo):
```

```
def __init__(self):
    # Initialize topology and default options
    Topo.__init__(self)
```

```
s1 = self.addSwitch('s1',dpid='0000000000000001')
s2a = self.addSwitch('s2a',dpid='000000000000002a')
s2b = self.addSwitch('s2b',dpid='000000000000002b')
s2c = self.addSwitch('s2c',dpid='000000000000002c')
s3 = self.addSwitch('s3',dpid='0000000000000003')
self.addLink(s1, s2a)
self.addLink(s1, s2b)
self.addLink(s2b, s2c)
self.addLink(s3, s2a)
self.addLink(s3, s2c)
host_1 = self.addHost('h1',ip='10.0.0.1',mac='10:00:00:00:00:01')
host_2 = self.addHost('h2',ip='10.0.0.2',mac='10:00:00:00:00:02')
self.addLink(host_1, s1)
self.addLink(host_2, s3)
```

```
topos = { 'shortest_path': ( lambda: Demo_Topo() ) }
```

3. Update the mapping service with required information

   Sample payload:

```
{
  "mappings": {
    "outer-map": [
      {
        "id": "uva",
        "inner-map": [
          {
            "inner-key": "ip_prefix",
            "value": "10.0.0.1/32"
          },
          {
            "inner-key": "mpls_label",
            "value": "15"
          },
          {
            "inner-key": "switch_port",
            "value": "openflow:1:1"
          }
        ]
      },
      {
        "id": "eur",
        "inner-map": [
          {
            "inner-key": "ip_prefix",
            "value": "10.0.0.2/32"
          },
          {
            "inner-key": "mpls_label",
            "value": "16"
          },
          {
```

(continues on next page)

(continued from previous page)

```
            "inner-key": "switch_port",
            "value": "openflow:3:1"
          }
        ]
      }
    ]
  }
}
```

4. Create bidirectional Intents using Karaf command line or RestCONF:

   Example:

```
intent:add -f uva -t eur -a ALLOW
intent:add -f eur -t uva -a ALLOW
```

5. Verify by running ovs-ofctl command on mininet if the flows were pushed correctly to the nodes that form the shortest path.

   Example:

```
ovs-ofctl -O OpenFlow13 dump-flows s1
```

NetIDE Developer Guide

Overview

The NetIDE Network Engine enables portability and cooperation inside a single network by using a client/server multi-controller SDN architecture. Separate "Client SDN Controllers" host the various SDN Applications with their access to the actual physical network abstracted and coordinated through a single "Server SDN Controller", in this instance OpenDaylight. This allows applications written for Ryu/Floodlight/Pyretic to execute on OpenDaylight managed infrastructure.

The "Network Engine" is modular by design:

- An OpenDaylight plugin, "shim", sends/receives messages to/from subscribed SDN Client Controllers. This consumes the ODL OpenFlow Plugin
- An initial suite of SDN Client Controller "Backends": Floodlight, Ryu, Pyretic. Further controllers may be added over time as the engine is extensible.

The Network Engine provides a compatibility layer capable of translating calls of the network applications running on top of the client controllers, into calls for the server controller framework. The communication between the client and the server layers is achieved through the NetIDE intermediate protocol, which is an application-layer protocol on top of TCP that transmits the network control/management messages from the client to the server controller and vice-versa. Between client and server controller sits the Core Layer which also "speaks" the intermediate protocol. The core layer implements three main functions:

1. interfacing with the client backends and server shim, controlling the lifecycle of controllers as well as modules in them,
2. orchestrating the execution of individual modules (in one client controller) or complete applications (possibly spread across multiple client controllers),
3. interfacing with the tools.

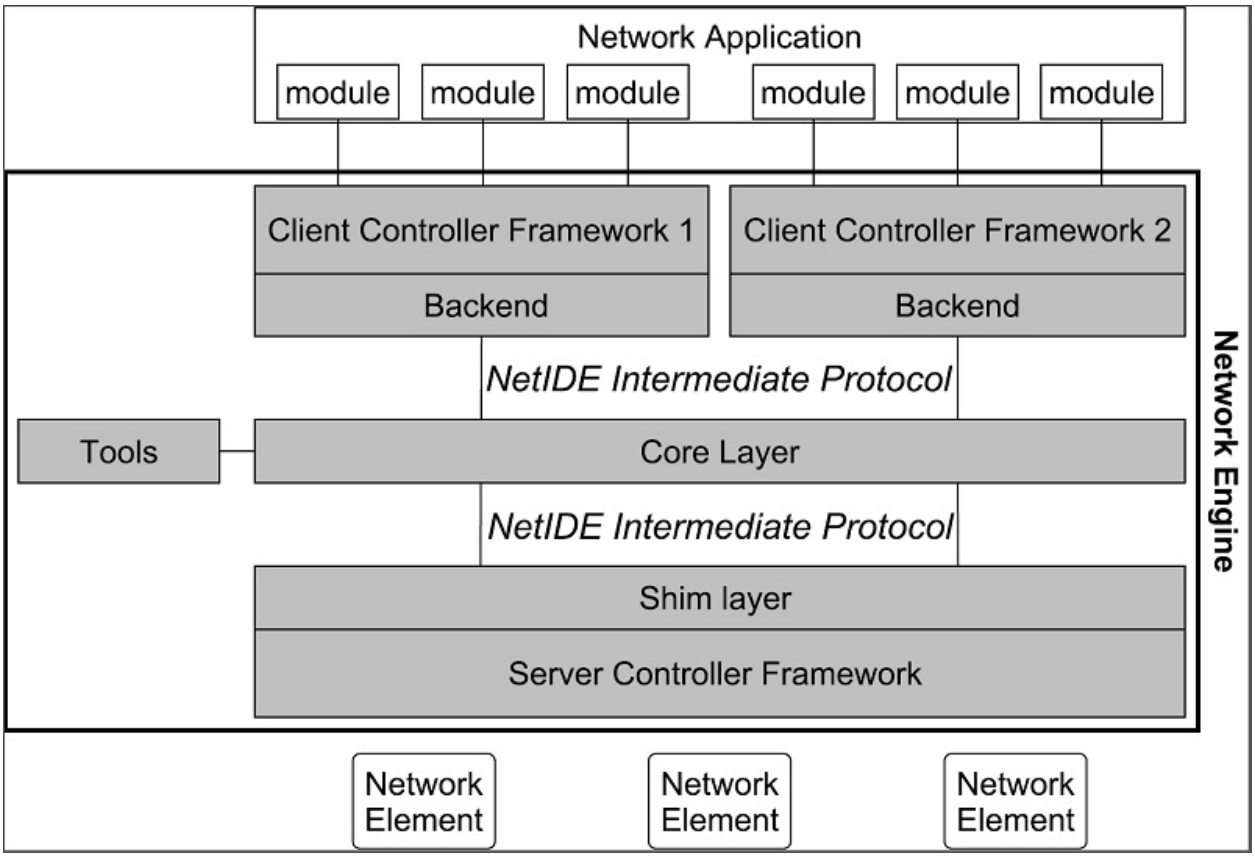

Fig. 20: NetIDE Network Engine Architecture

NetIDE Intermediate Protocol

The Intermediate Protocol serves several needs, it has to:

1. carry control messages between core and shim/backend, e.g., to start up/take down a particular module, providing unique identifiers for modules,
2. carry event and action messages between shim, core, and backend, properly demultiplexing such messages to the right module based on identifiers,
3. encapsulate messages specific to a particular SBI protocol version (e.g., OpenFlow 1.X, NETCONF, etc.) towards the client controllers with proper information to recognize these messages as such.

The NetIDE packages can be added as dependencies in Maven projects by putting the following code in the *pom.xml* file.

```
<dependency>
    <groupId>org.opendaylight.netide</groupId>
    <artifactId>api</artifactId>
    <version>${NETIDE_VERSION}</version>
</dependency>
```

The current stable version for NetIDE is `0.2.0-Boron`.

Protocol specification

Messages of the NetIDE protocol contain two basic elements: the NetIDE header and the data (or payload). The NetIDE header, described below, is placed before the payload and serves as the communication and control link between the different components of the Network Engine. The payload can contain management messages, used by the components of the Network Engine to exchange relevant information, or control/configuration messages (such as OpenFlow, NETCONF, etc.) crossing the Network Engine generated by either network application modules or by the network elements.

The NetIDE header is defined as follows:

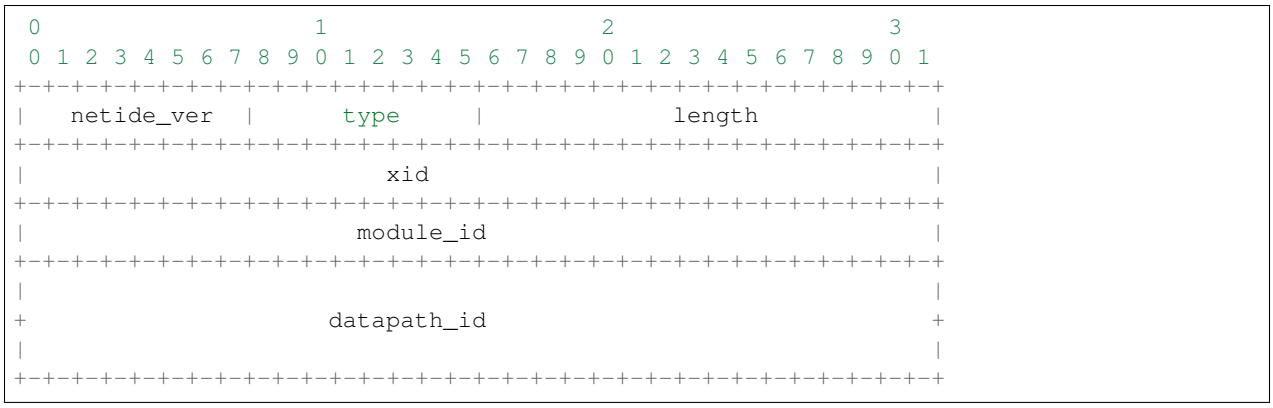

where each tick mark represents one bit position. Alternatively, in a C-style coding format, the NetIDE header can be represented with the following structure:

```
struct netide_header {
    uint8_t netide_ver ;
    uint8_t type ;
    uint16_t length ;
    uint32_t xid
```

(continues on next page)

(continued from previous page)

```
    uint32_t module_id
    uint64_t datapath_id
};
```

- `netide_ver` is the version of the NetIDE protocol (the current version is v1.2, which is identified with value 0x03).
- `length` is the total length of the payload in bytes.
- `type` contains a code that indicates the type of the message according with the following values:

```
enum type {
    NETIDE_HELLO = 0x01 ,
    NETIDE_ERROR = 0x02 ,
    NETIDE_MGMT = 0x03 ,
    MODULE_ANNOUNCEMENT = 0x04 ,
    MODULE_ACKNOWLEDGE = 0x05 ,
    NETIDE_HEARTBEAT = 0x06 ,
    NETIDE_OPENFLOW = 0x11 ,
    NETIDE_NETCONF = 0x12 ,
    NETIDE_OPFLEX = 0x13
};
```

- `datapath_id` is a 64-bit field that uniquely identifies the network elements.
- `module_id` is a 32-bits field that uniquely identifies Backends and application modules running on top of each client controller. The composition mechanism in the core layer leverages on this field to implement the correct execution flow of these modules.
- `xid` is the transaction identifier associated to the each message. Replies must use the same value to facilitate the pairing.

Module announcement

The first operation performed by a Backend is registering itself and the modules that it is running to the Core. This is done by using the `MODULE_ANNOUNCEMENT` and `MODULE_ACKNOWLEDGE` message types. As a result of this process, each Backend and application module can be recognized by the Core through an identifier (the `module_id`) placed in the NetIDE header. First, a Backend registers itself by using the following schema: backend-<platform name>-<pid>.

For example,odule a Ryu Backend will register by using the following name in the message backend-ryu-12345 where 12345 is the process ID of the registering instance of the Ryu platform. The format of the message is the following:

```
struct NetIDE_message {
    netide_ver = 0x03
    type = MODULE_ANNOUNCEMENT
    length = len(" backend -< platform_name >-<pid >")
    xid = 0
    module_id = 0
    datapath_id = 0
    data = " backend -< platform_name >-<pid >"
}
```

The answer generated by the Core will include a `module_id` number and the Backend name in the payload (the same indicated in the `MODULE_ANNOUNCEMENT` message):

```
struct NetIDE_message {
    netide_ver = 0x03
    type = MODULE_ACKNOWLEDGE
    length = len(" backend -< platform_name >-<pid >")
    xid = 0
    module_id = MODULE_ID
    datapath_id = 0
    data = " backend -< platform_name >-<pid >"
}
```

Once a Backend has successfully registered itself, it can start registering its modules with the same procedure described above by indicating the name of the module in the data (e.g. data="Firewall"). From this point on, the Backend will insert its own `module_id` in the header of the messages it generates (e.g. heartbeat, hello messages, OpenFlow echo messages from the client controllers, etc.). Otherwise, it will encapsulate the control/configuration messages (e.g. FlowMod, PacketOut, FeatureRequest, NETCONF request, etc.) generated by network application modules with the specific +module_id+s.

Heartbeat

The heartbeat mechanism has been introduced after the adoption of the ZeroMQ messaging queuing library to transmit the NetIDE messages. Unfortunately, the ZeroMQ library does not offer any mechanism to find out about disrupted connections (and also completely unresponsive peers). This limitation of the ZeroMQ library can be an issue for the Core's composition mechanism and for the tools connected to the Network Engine, as they cannot understand when an client controller disconnects or crashes. As a consequence, Backends must periodically send (let's say every 5 seconds) a "heartbeat" message to the Core. If the Core does not receive at least one "heartbeat" message from the Backend within a certain timeframe, the Core considers it disconnected, removes all the related data from its memory structures and informs the relevant tools. The format of the message is the following:

```
struct NetIDE_message {
    netide_ver = 0x03
    type = NETIDE_HEARTBEAT
    length = 0
    xid = 0
    module_id = backend -id
    datapath_id = 0
    data = 0
}
```

Handshake

Upon a successful connection with the Core, the client controller must immediately send a hello message with the list of the control and/or management protocols needed by the applications deployed on top of it.

```
struct NetIDE_message {
    struct netide_header header ;
    uint8 data [0]
};
```

The header contains the following values:

- `netide ver=0x03`
- `type=NETIDE_HELLO`
- `length=2*NR_PROTOCOLS`

- `data` contains one 2-byte word (in big endian order) for each protocol, with the first byte containing the code of the protocol according to the above enum, while the second byte in- dictates the version of the protocol (e.g. according to the ONF specification, 0x01 for OpenFlow v1.0, 0x02 for OpenFlow v1.1, etc.). NETCONF version is marked with 0x01 that refers to the specification in the RFC6241, while OpFlex version is marked with 0x00 since this protocol is still in work-in-progress stage.

The Core relays hello messages to the server controller which responds with another hello message containing the following:

- `netide ver=0x03`
- `type=NETIDE_HELLO`
- `length=2*NR_PROTOCOLS`

If at least one of the protocols requested by the client is supported. In particular, `data` contains the codes of the protocols that match the client's request (2-bytes words, big endian order). If the hand- shake fails because none of the requested protocols is supported by the server controller, the header of the answer is as follows:

- `netide ver=0x03`
- `type=NETIDE_ERROR`
- `length=2*NR_PROTOCOLS`
- `data` contains the codes of all the protocols supported by the server controller (2-bytes words, big endian order). In this case, the TCP session is terminated by the server controller just after the answer is received by the client. ‘

NetVirt Developer Guide

Neutron Service Developer Guide

Overview

This Karaf feature (`odl-neutron-service`) provides integration support for OpenStack Neutron via the OpenDaylight ML2 mechanism driver. The Neutron Service is only one of the components necessary for OpenStack integration. It defines YANG models for OpenStack Neutron data models and northbound API via REST API and YANG model RESTCONF.

Those developers who want to add new provider for new OpenStack Neutron extensions/services (Neutron constantly adds new extensions/services and OpenDaylight will keep up with those new things) need to communicate with this Neutron Service or add models to Neutron Service. If you want to add new extensions/services themselves to the Neutron Service, new YANG data models need to be added, but that is out of scope of this document because this guide is for a developer who will be *using* the feature to build something separate, but *not* somebody who will be developing code for this feature itself.

Neutron Service Architecture

The Neutron Service defines YANG models for OpenStack Neutron integration. When OpenStack admins/users request changes (creation/update/deletion) of Neutron resources, e.g., Neutron network, Neutron subnet, Neutron port, the corresponding YANG model within OpenDaylight will be modified. The OpenDaylight OpenStack will subscribe the changes on those models and will be notified those modification through MD-SAL when changes are made. Then the provider will do the necessary tasks to realize OpenStack integration. How to realize it (or even not realize it) is up to each provider. The Neutron Service itself does not take care of it.

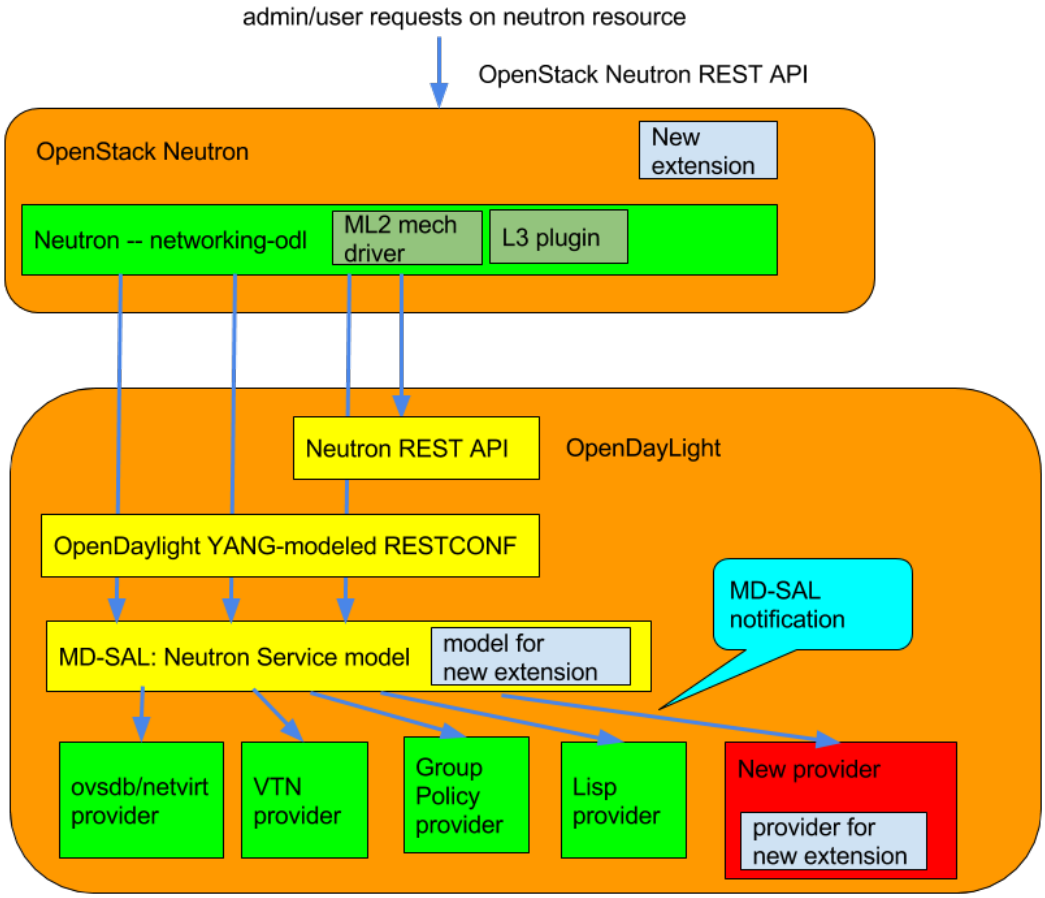

Fig. 21: Neutron Service Architecture

How to Write a SB Neutron Consumer

In Boron, there is only one options for SB Neutron Consumers:

- Listening for changes via the Neutron YANG model

Until Beryllium there was another way with the legacy I*Aware interface. From Boron, the interface was eliminated. So all the SB Neutron Consumers have to use Neutron YANG model.

Neutron YANG models

Neutron service defines YANG models for Neutron. The details can be found at

- https://git.opendaylight.org/gerrit/gitweb?p=neutron.git;a=tree;f=model/src/main/yang;hb=refs/heads/stable/boron

Basically those models are based on OpenStack Neutron API definitions. For exact definitions, OpenStack Neutron source code needs to be referred as the above documentation doesn't always cover the necessary details. There is nothing special to utilize those Neutron YANG models. The basic procedure will be:

1. subscribe for changes made to the the model
2. respond on the data change notification for each models

Note: Currently there is no way to refuse the request configuration at this point. That is left to future work.

```
public class NeutronNetworkChangeListener implements DataChangeListener,␣
↪AutoCloseable {
    private ListenerRegistration<DataChangeListener> registration;
    private DataBroker db;

    public NeutronNetworkChangeListener(DataBroker db){
        this.db = db;
        // create identity path to register on service startup
        InstanceIdentifier<Network> path = InstanceIdentifier
                .create(Neutron.class)
                .child(Networks.class)
                .child(Network.class);
        LOG.debug("Register listener for Neutron Network model data changes");
        // register for Data Change Notification
        registration =
                this.db.registerDataChangeListener(LogicalDatastoreType.CONFIGURATION,
↪ path, this, DataChangeScope.ONE);

    }

    @Override
    public void onDataChanged(
            AsyncDataChangeEvent<InstanceIdentifier<?>, DataObject> changes) {
        LOG.trace("Data changes : {}",changes);

        // handle data change notification
        Object[] subscribers = NeutronIAwareUtil.getInstances(INeutronNetworkAware.
↪class, this);
        createNetwork(changes, subscribers);
        updateNetwork(changes, subscribers);
```

(continues on next page)

(continued from previous page)

```
            deleteNetwork(changes, subscribers);
        }
    }
```

Neutron configuration

From Boron, new models of configuration for OpenDaylight to tell OpenStack neutron/networking-odl its configuration/capability.

hostconfig

This is for OpenDaylight to tell per-node configuration to Neutron. Especially this is used by pseudo agent port binding heavily.

The model definition can be found at

- https://git.opendaylight.org/gerrit/gitweb?p=neutron.git;a=blob;f=model/src/main/yang/neutron-hostconfig.yang;hb=refs/heads/stable/boron

How to populate this for pseudo agent port binding is documented at

- http://git.openstack.org/cgit/openstack/networking-odl/tree/doc/source/devref/hostconfig.rst

Neutron extension config

In Boron this is experimental. The model definition can be found at

- https://git.opendaylight.org/gerrit/gitweb?p=neutron.git;a=blob;f=model/src/main/yang/neutron-extensions.yang;hb=refs/heads/stable/boron

Each Neutron Service provider has its own feature set. Some support the full features of OpenStack, but others support only a subset. With same supported Neutron API, some functionality may or may not be supported. So there is a need for a way that OpenDaylight can tell networking-odl its capability. Thus networking-odl can initialize Neutron properly based on reported capability.

Neutorn Logger

There is another small Karaf feature, `odl-neutron-logger`, which logs changes of Neutron YANG models. which can be used for debug/audit.

It would also help to understand how to listen the change.

- https://git.opendaylight.org/gerrit/gitweb?p=neutron.git;a=blob;f=neutron-logger/src/main/java/org/opendaylight/neutron/logger/NeutronLogger.java;hb=refs/heads/stable/boron

API Reference Documentation

The OpenStack Neutron API references

- http://developer.openstack.org/api-ref-networking-v2.html
- http://developer.openstack.org/api-ref-networking-v2-ext.html

Neutron Northbound

How to add new API support

OpenStack Neutron is a moving target. It is continuously adding new features as new rest APIs. Here is a basic step to add new API support:

In the Neutron Northbound project:

- Add new YANG model for it under `neutron/model/src/main/yang` and `update neutron.yang`
- Add northbound API for it, and neutron-spi
  - Implement `Neutron<New API>Request.java` and `Neutron<New API>Norhtbound.java` under `neutron/northbound-api/src/main/java/org/opendaylight/neutron/northbound/api/`
  - Implement `INeutron<New API>CRUD.java` and new data structure if any under `neutron/neutron-spi/src/main/java/org/opendaylight/neutron/spi/`
  - update `neutron/neutron-spi/src/main/java/org/opendaylight/neutron/spi/NeutronCRUDInterfaces.java` to wire new CRUD interface
  - Add unit tests, `Neutron<New structure>JAXBTest.java` under `neutron/neutron-spi/src/test/java/org/opendaylight/neutron/spi/`
- update `neutron/northbound-api/src/main/java/org/opendaylight/neutron/northbound/api/NeutronNorthboundRSApplication.java` to wire new northbound api to `RSApplication`
- Add transcriber, `Neutron<New API>Interface.java` under `transcriber/src/main/java/org/opendaylight/neutron/transcriber/`
- update `transcriber/src/main/java/org/opendaylight/neutron/transcriber/NeutronTranscriberProvider.java` to wire a new transcriber
  - Add integration tests `Neutron<New API>Tests.java` under `integration/test/src/test/java/org/opendaylight/neutron/e2etest/`
  - update `integration/test/src/test/java/org/opendaylight/neutron/e2etest/ITNeutronE2E.java` to run a newly added tests.

In OpenStack networking-odl

- Add new driver (or plugin) for new API with tests.

In a southbound Neutron Provider

- implement actual backend to realize those new API by listening related YANG models.

How to write transcriber

For each Neutron data object, there is an `Neutron*Interface` defined within the transcriber artifact that will write that object to the MD-SAL configuration datastore.

All `Neutron*Interface` extend `AbstractNeutronInterface`, in which two methods are defined:

- one takes the Neutron object as input, and will create a data object from it.
- one takes an uuid as input, and will create a data object containing the uuid.

```
protected abstract T toMd(S neutronObject);
protected abstract T toMd(String uuid);
```

In addition the `AbstractNeutronInterface` class provides several other helper methods (`addMd`, `updateMd`, `removeMd`), which handle the actual writing to the configuration datastore.

The semantics of the `toMD()` methods

Each of the Neutron YANG models defines structures containing data. Further each YANG-modeled structure has it own builder. A particular `toMD()` method instantiates an instance of the correct builder, fills in the properties of the builder from the corresponding values of the Neutron object and then creates the YANG-modeled structures via the `build()` method.

As an example, one of the `toMD` code for Neutron Networks is presented below:

```
protected Network toMd(NeutronNetwork network) {
    NetworkBuilder networkBuilder = new NetworkBuilder();
    networkBuilder.setAdminStateUp(network.getAdminStateUp());
    if (network.getNetworkName() != null) {
        networkBuilder.setName(network.getNetworkName());
    }
    if (network.getShared() != null) {
        networkBuilder.setShared(network.getShared());
    }
    if (network.getStatus() != null) {
        networkBuilder.setStatus(network.getStatus());
    }
    if (network.getSubnets() != null) {
        List<Uuid> subnets = new ArrayList<Uuid>();
        for( String subnet : network.getSubnets()) {
            subnets.add(toUuid(subnet));
        }
        networkBuilder.setSubnets(subnets);
    }
    if (network.getTenantID() != null) {
        networkBuilder.setTenantId(toUuid(network.getTenantID()));
    }
    if (network.getNetworkUUID() != null) {
        networkBuilder.setUuid(toUuid(network.getNetworkUUID()));
    } else {
        logger.warn("Attempting to write neutron network without UUID");
    }
    return networkBuilder.build();
}
```

NeXt Developer Guide

Please see the NeXt documentation and tutorials here: https://github.com/NeXt-UI/next-tutorials

ODL Parent Developer Guide

Parent POMs

Overview

The ODL Parent component for OpenDaylight provides a number of Maven parent POMs which allow Maven projects to be easily integrated in the OpenDaylight ecosystem. Technically, the aim of projects in OpenDaylight is to produce Karaf features, and these parent projects provide common support for the different types of projects involved.

These parent projects are:

- `odlparent-lite` — the basic parent POM for Maven modules which don't produce artifacts (*e.g.* aggregator POMs)
- `odlparent` — the common parent POM for Maven modules containing Java code
- `bundle-parent` — the parent POM for Maven modules producing OSGi bundles
- `features-parent` — the parent POM for Maven modules producing Karaf features

odlparent-lite

This is the base parent for all OpenDaylight Maven projects and modules. It provides the following, notably to allow publishing artifacts to Maven Central:

- license information;
- organization information;
- issue management information (a link to our Bugzilla);
- continuous integration information (a link to our Jenkins setup);
- default Maven plugins (`maven-clean-plugin`, `maven-deploy-plugin`, `maven-install-plugin`, `maven-javadoc-plugin` with HelpMojo support, `maven-project-info-reports-plugin`, `maven-site-plugin` with Asciidoc support, `jdepend-maven-plugin`);
- distribution management information.

It also defines two profiles which help during development:

- `q` (`-Pq`), the quick profile, which disables tests, code coverage, Javadoc generation, code analysis, etc. — anything which isn't necessary to build the bundles and features (see this blog post for details);
- `addInstallRepositoryPath` (`-DaddInstallRepositoryPath=.../karaf/system`) which can be used to drop a bundle in the appropriate Karaf location, to enable hot-reloading of bundles during development (see this blog post for details).

For modules which don't produce any useful artifacts (*e.g.* aggregator POMs), you should add the following to avoid processing artifacts:

```
<build>
    <plugins>
        <plugin>
```

(continues on next page)

(continued from previous page)

```
            <groupId>org.apache.maven.plugins</groupId>
            <artifactId>maven-deploy-plugin</artifactId>
            <configuration>
                <skip>true</skip>
            </configuration>
        </plugin>
        <plugin>
            <groupId>org.apache.maven.plugins</groupId>
            <artifactId>maven-install-plugin</artifactId>
            <configuration>
                <skip>true</skip>
            </configuration>
        </plugin>
    </plugins>
</build>
```

odlparent

This inherits from `odlparent-lite` and mainly provides dependency and plugin management for OpenDaylight projects.

If you use any of the following libraries, you should rely on `odlparent` to provide the appropriate versions:

- Akka (and Scala)
- Apache Commons:
  - `commons-codec`
  - `commons-fileupload`
  - `commons-io`
  - `commons-lang`
  - `commons-lang3`
  - `commons-net`
- Apache Shiro
- Guava
- JAX-RS with Jersey
- JSON processing:
  - GSON
  - Jackson
- Logging:
  - Logback
  - SLF4J
- Netty
- OSGi:

  - Apache Felix
  - core OSGi dependencies (`core`, `compendium`...)
- Testing:
  - Hamcrest
  - JSON assert
  - JUnit
  - Mockito
  - Pax Exam
  - PowerMock
- XML/XSL:
  - Xerces
  - XML APIs

Note: This list isn't exhaustive. It's also not cast in stone; if you'd like to add a new dependency (or migrate a dependency), please contact the mailing list.

`odlparent` also enforces some Checkstyle verification rules. In particular, it enforces the common license header used in all OpenDaylight code:

```
/*
 * Copyright © ${year} ${holder} and others.  All rights reserved.
 *
 * This program and the accompanying materials are made available under the
 * terms of the Eclipse Public License v1.0 which accompanies this distribution,
 * and is available at http://www.eclipse.org/legal/epl-v10.html
 */
```

where "`${year}`" is initially the first year of publication, then (after a year has passed) the first and latest years of publication, separated by commas (*e.g.* "2014, 2016"), and "`${holder}`" is the initial copyright holder (typically, the first author's employer). "All rights reserved" is optional.

If you need to disable this license check, *e.g.* for files imported under another license (EPL-compatible of course), you can override the `maven-checkstyle-plugin` configuration. `features-test` does this for its `CustomBundleUrlStreamHandlerFactory` class, which is ASL-licensed:

```
<plugin>
    <artifactId>maven-checkstyle-plugin</artifactId>
    <executions>
        <execution>
            <id>check-license</id>
            <goals>
                <goal>check</goal>
            </goals>
            <phase>process-sources</phase>
            <configuration>
                <configLocation>check-license.xml</configLocation>
                <headerLocation>EPL-LICENSE.regexp.txt</headerLocation>
                <includeResources>false</includeResources>
                <includeTestResources>false</includeTestResources>
```

(continues on next page)

(continued from previous page)

```
                <sourceDirectory>${project.build.sourceDirectory}</sourceDirectory>
                <excludes>
                    <!-- Skip Apache Licensed files -->
                    org/opendaylight/odlparent/featuretest/
↪CustomBundleUrlStreamHandlerFactory.java
                </excludes>
                <failsOnError>false</failsOnError>
                <consoleOutput>true</consoleOutput>
            </configuration>
        </execution>
    </executions>
</plugin>
```

bundle-parent

This inherits from `odlparent` and enables functionality useful for OSGi bundles:

- `maven-javadoc-plugin` is activated, to build the Javadoc JAR;
- `maven-source-plugin` is activated, to build the source JAR;
- `maven-bundle-plugin` is activated (including extensions), to build OSGi bundles (using the "bundle" packaging).

In addition to this, JUnit is included as a default dependency in "test" scope.

features-parent

This inherits from `odlparent` and enables functionality useful for Karaf features:

- `karaf-maven-plugin` is activated, to build Karaf features — but for OpenDaylight, projects need to use "jar" packaging (**not** "kar");
- `features.xml` files are processed from templates stored in `src/main/features/features.xml`;
- Karaf features are tested after build to ensure they can be activated in a Karaf container.

The `features.xml` processing allows versions to be ommitted from certain feature dependencies, and replaced with "`{{version}}`". For example:

```
<features name="odl-mdsal-${project.version}" xmlns="http://karaf.apache.org/xmlns/
↪features/v1.2.0"
   xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
   xsi:schemaLocation="http://karaf.apache.org/xmlns/features/v1.2.0 http://karaf.
↪apache.org/xmlns/features/v1.2.0">

    <repository>mvn:org.opendaylight.odlparent/features-odlparent/{{VERSION}}/xml/
↪features</repository>

    [...]
    <feature name='odl-mdsal-broker-local' version='${project.version}' description=
↪"OpenDaylight :: MDSAL :: Broker">
        <feature version='${yangtools.version}'>odl-yangtools-common</feature>
        <feature version='${mdsal.version}'>odl-mdsal-binding-dom-adapter</feature>
```

(continues on next page)

(continued from previous page)

```
        <feature version='${mdsal.model.version}'>odl-mdsal-models</feature>
        <feature version='${project.version}'>odl-mdsal-common</feature>
        <feature version='${config.version}'>odl-config-startup</feature>
        <feature version='${config.version}'>odl-config-netty</feature>
        <feature version='[3.3.0,4.0.0)'>odl-lmax</feature>
        [...]
        <bundle>mvn:org.opendaylight.controller/sal-dom-broker-config/{{VERSION}}</
↪bundle>
        <bundle start-level="40">mvn:org.opendaylight.controller/blueprint/{{VERSION}}
↪</bundle>
        <configfile finalname="${config.configfile.directory}/${config.mdsal.
↪configfile}">mvn:org.opendaylight.controller/md-sal-config/{{VERSION}}/xml/config</
↪configfile>
    </feature>
```

As illustrated, versions can be ommitted in this way for repository dependencies, bundle dependencies and configuration files. They must be specified traditionally (either hard-coded, or using Maven properties) for feature dependencies.

Features

The ODL Parent component for OpenDaylight provides a number of Karaf features which can be used by other Karaf features to use certain third-party upstream dependencies.

These features are:

- Akka features (in the `features-akka` repository):
  - `odl-akka-all` — all Akka bundles;
  - `odl-akka-scala` — Scala runtime for OpenDaylight;
  - `odl-akka-system` — Akka actor framework bundles;
  - `odl-akka-clustering` — Akka clustering bundles and dependencies;
  - `odl-akka-leveldb` — LevelDB;
  - `odl-akka-persistence` — Akka persistence;
- general third-party features (in the `features-odlparent` repository):
  - `odl-netty` — all Netty bundles;
  - `odl-guava` — Guava;
  - `odl-lmax` — LMAX Disruptor.

To use these, you need to declare a dependency on the appropriate repository in your `features.xml` file:

```
<repository>mvn:org.opendaylight.odlparent/features-odlparent/{{VERSION}}/xml/features
↪</repository>
```

and then include the feature, *e.g.*:

```
<feature name='odl-mdsal-broker-local' version='${project.version}' description=
↪"OpenDaylight :: MDSAL :: Broker">
    [...]
```

(continues on next page)

(continued from previous page)

```
    <feature version='[3.3.0,4.0.0)'>odl-lmax</feature>
    [...]
</feature>
```

You also need to depend on the features repository in your POM:

```
<dependency>
    <groupId>org.opendaylight.odlparent</groupId>
    <artifactId>features-odlparent</artifactId>
    <classifier>features</classifier>
    <type>xml</type>
</dependency>
```

assuming the appropriate dependency management:

```
<dependencyManagement>
    <dependencies>
        <dependency>
            <groupId>org.opendaylight.odlparent</groupId>
            <artifactId>odlparent-artifacts</artifactId>
            <version>1.7.0-SNAPSHOT</version>
            <scope>import</scope>
            <type>pom</type>
        </dependency>
    </dependencies>
</dependencyManagement>
```

(the version number there is appropriate for Boron). For the time being you also need to depend separately on the individual JARs as compile-time dependencies to build your dependent code; the relevant dependencies are managed in `odlparent`'s dependency management.

The suggested version ranges are as follows:

- `odl-netty`: `[4.0.37,4.1.0)` or `[4.0.37,5.0.0)`;
- `odl-guava`: `[18,19)` (if your code is ready for it, `[19,20)` is also available, but the current default version of Guava in OpenDaylight is 18);
- `odl-lmax`: `[3.3.4,4.0.0)`

OCP Plugin Developer Guide

This document is intended for both OCP (ORI [Open Radio Interface] C&M [Control and Management] Protocol) agent developers and OpenDaylight service/application developers. It describes essential information needed to implement an OCP agent that is capable of interoperating with the OCP plugin running in OpenDaylight, including the OCP connection establishment and state machines used on both ends of the connection. It also provides a detailed description of the northbound/southbound APIs that the OCP plugin exposes to allow automation and programmability.

Overview

OCP is an ETSI standard protocol for control and management of Remote Radio Head (RRH) equipment. The OCP Project addresses the need for a southbound plugin that allows applications and controller services to interact with

RRHs using OCP. The OCP southbound plugin will allow applications acting as a Radio Equipment Control (REC) to interact with RRHs that support an OCP agent.

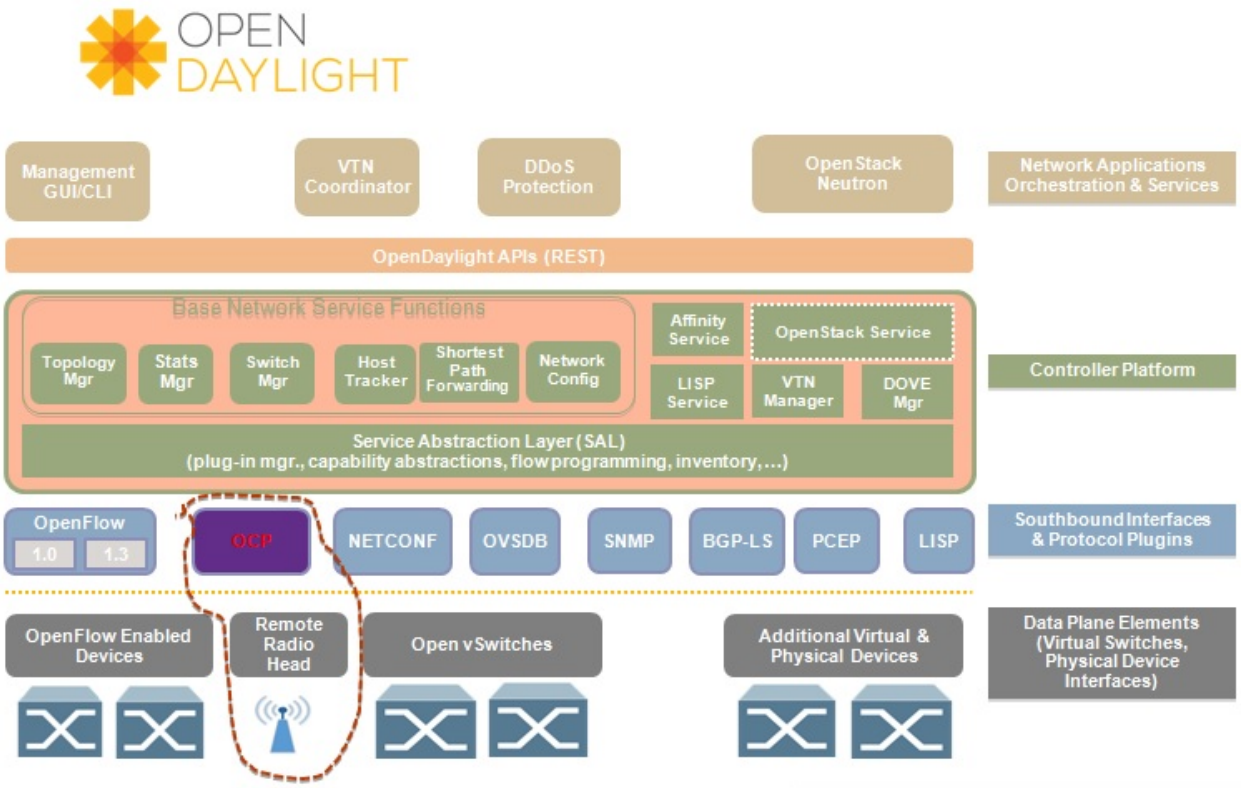

Fig. 22: OCP southbound plugin

Architecture

OCP is a vendor-neutral standard communications interface defined to enable control and management between RE and REC of an ORI architecture. The OCP Plugin supports the implementation of the OCP specification; it is based on the Model Driven Service Abstraction Layer (MD-SAL) architecture.

The OCP Plugin project consists of three main components: OCP southbound plugin, OCP protocol library and OCP service. For details on each of them, refer to the OCP Plugin User Guide.

Connection Establishment

The OCP layer is transported over a TCP/IP connection established between the RE and the REC. OCP provides the following functions:

- Control & Management of the RE by the REC
- Transport of AISG/3GPP Iuant Layer 7 messages and alarms between REC and RE

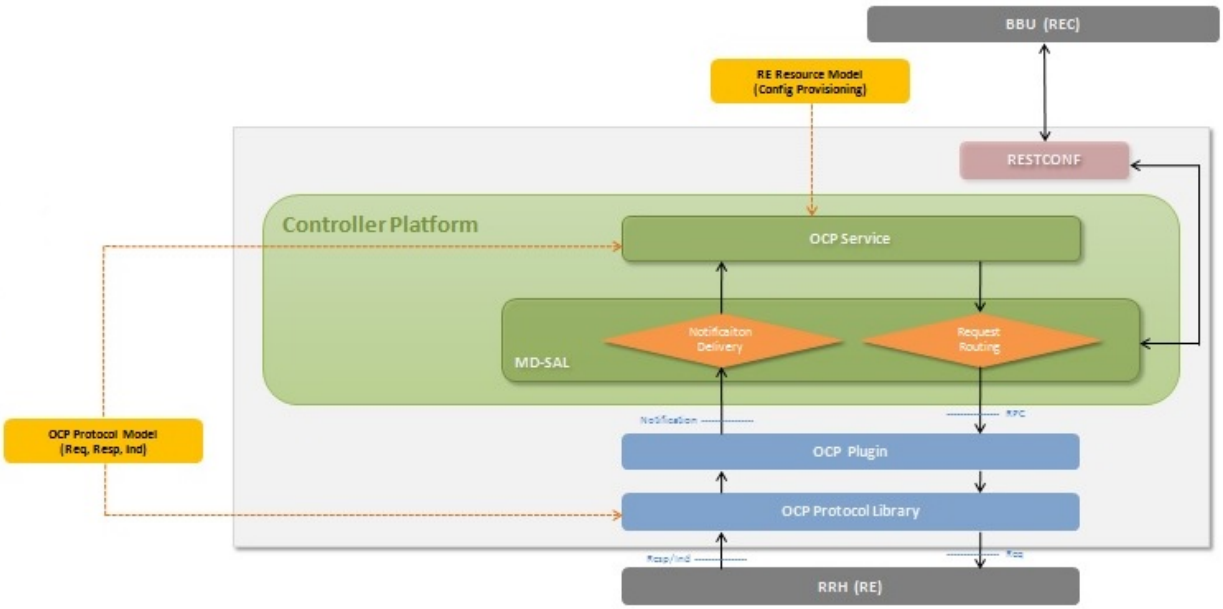

Fig. 23: Overall architecture

Hello Message

Hello message is used by the OCP agent during connection setup. It is used for version negotiation. When the connection is established, the OCP agent immediately sends a Hello message with the version field set to highest OCP version supported by itself, along with the verdor ID and serial number of the radio head it is running on.

The combinaiton of the verdor ID and serial number will be used by the OCP plugin to uniquely identify a managed radio head. When not receiving reply from the OCP plugin, the OCP agent can resend Hello message with pre-defined Hello timeout (THLO) and Hello resend times (NHLO).

According to ORI spec, the default value of TCP Link Monitoring Timer (TTLM) is 50 seconds. The RE shall trigger an OCP layer restart while TTLM expires in RE or the RE detects a TCP link failure. So we may define NHLO * THLO = 50 seconds (e.g. NHLO = 10, THLO = 5 seconds).

By nature the Hello message is a new type of indication, and it contains supported OCP version, vendor ID and serial number as shown below.

Hello message.

```
<?xml version="1.0" encoding="UTF-8"?>
<msg xmlns="http://uri.etsi.org/ori/002-2/v4.1.1">
  <header>
    <msgType>IND</msgType>
    <msgUID>0</msgUID>
  </header>
  <body>
    <helloInd>
      <version>4.1.1</version>
      <vendorId>XYZ</vendorId>
      <serialNumber>ABC123</serialNumber>
    </helloInd>
  </body>
</msg>
```

Ack Message

Hello from the OCP agent will always make the OCP plugin respond with ACK. In case everything is OK, it will be ACK(OK). In case something is wrong, it will be ACK(FAIL).

If the OCP agent receives ACK(OK), it goes to Established state. If the OCP agent receives ACK(FAIL), it goes to Maintenance state. The failure code and reason of ACK(FAIL) are defined as below:

- FAIL_OCP_VERSION (OCP version not supported)
- FAIL_NO_MORE_CAPACITY (OCP plugin cannot control any more radio heads)

The result inside Ack message indicates OK or FAIL with different reasons.

Ack message.

```
<?xml version="1.0" encoding="UTF-8"?>
<msg xmlns="http://uri.etsi.org/ori/002-2/v4.1.1">
  <header>
    <msgType>ACK</msgType>
    <msgUID>0</msgUID>
  </header>
  <body>
    <helloAck>
      <result>FAIL_OCP_VERSION</result>
    </helloAck>
  </body>
</msg>
```

State Machines

The following figures illustrate the Finite State Machine (FSM) of the OCP agent and OCP plugin for new connection procedure.

Northbound APIs

There are ten exposed northbound APIs: health-check, set-time, re-reset, get-param, modify-param, create-obj, delete-obj, get-state, modify-state and get-fault

health-check

The Health Check procedure allows the application to verify that the OCP layer is functioning correctly at the RE.

Default URL: http://localhost:8181/restconf/operations/ocp-service:health-check-nb

POST Input

Field Name	Type	Description	Example	Required ?
nodeId	String	Inventory node reference for OCP radio head	ocp:MTI-101-200	Yes
tcpLinkMonTimeout	unsigned Short	TCP Link Monitoring Timeout (unit: seconds)	50	Yes

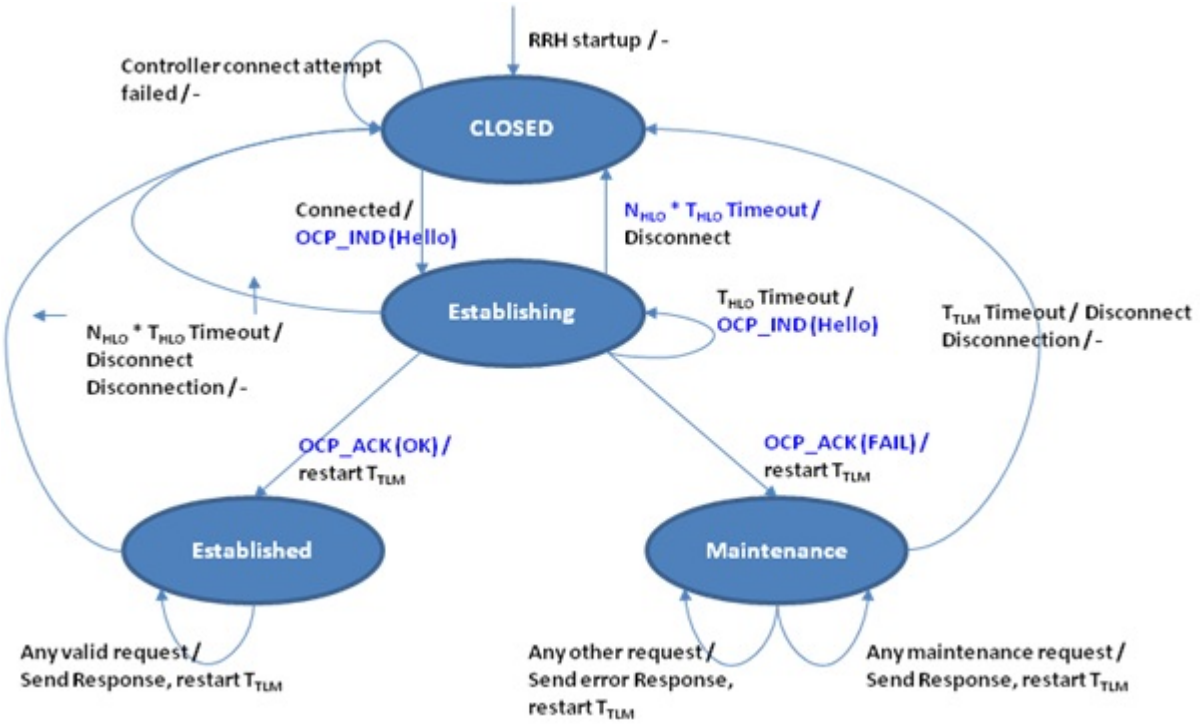

Fig. 24: OCP agent state machine

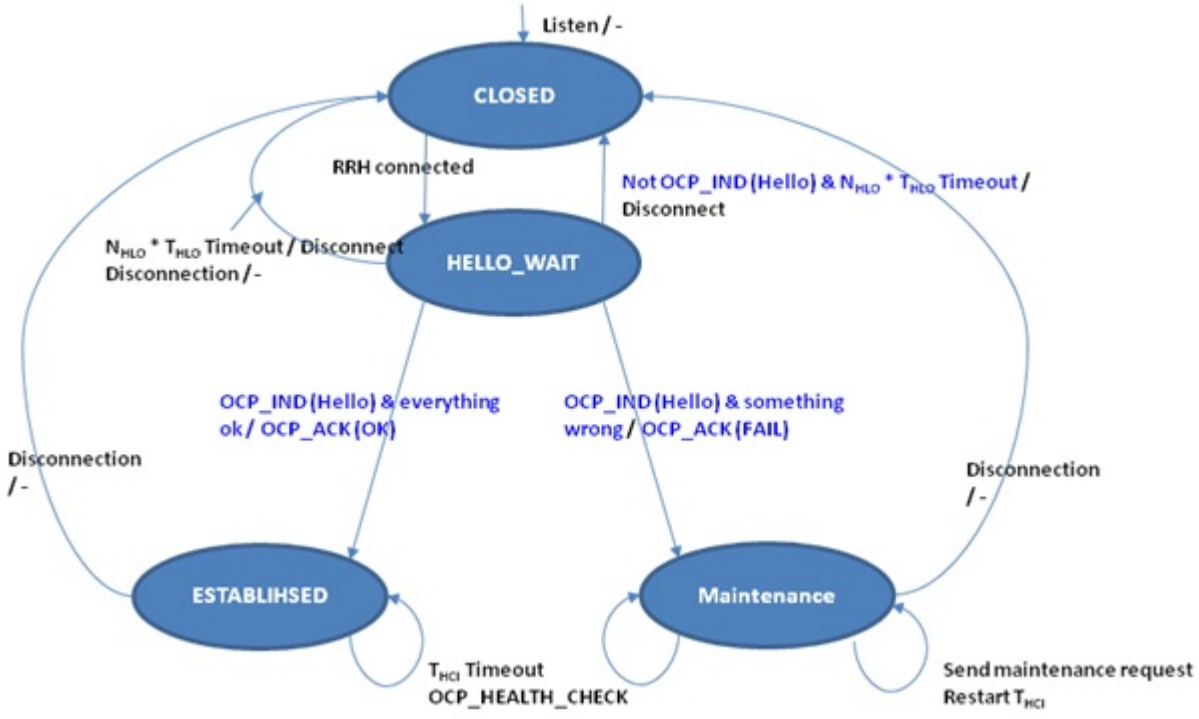

Fig. 25: OCP plugin state machine

Example.

```
{
    "health-check-nb": {
        "input": {
            "nodeId": "ocp:MTI-101-200",
            "tcpLinkMonTimeout": "50"
        }
    }
}
```

POST Output

Field Name	Type	Description
result	String, enumerated	Common default result codes

Example.

```
{
    "output": {
        "result": "SUCCESS"
    }
}
```

set-time

The Set Time procedure allows the application to set/update the absolute time reference that shall be used by the RE.

Default URL: http://localhost:8181/restconf/operations/ocp-service:set-time-nb

POST Input

Field Name	Type	Description	Example	Required?
nodeId	String	Inventory node reference for OCP radio head	ocp:MTI-101-200	Yes
newTime	dateTime	New datetime setting for radio head	2016-04-26T10:23:00-05:00	Yes

Example.

```
{
    "set-time-nb": {
        "input": {
            "nodeId": "ocp:MTI-101-200",
            "newTime": "2016-04-26T10:23:00-05:00"
        }
    }
}
```

POST Output

Field Name	Type	Description
result	String, enumerated	Common default result codes + FAIL_INVALID_TIMEDATA

Example.

```
{
    "output": {
        "result": "SUCCESS"
    }
}
```

re-reset

The RE Reset procedure allows the application to reset a specific RE.

Default URL: http://localhost:8181/restconf/operations/ocp-service:re-reset-nb

POST Input

Field Name	Type	Description	Example	Required?
nodeId	String	Inventory node reference for OCP radio head	ocp:MTI-101-200	Yes

Example.

```
{
    "re-reset-nb": {
        "input": {
            "nodeId": "ocp:MTI-101-200"
        }
    }
}
```

POST Output

Field Name	Type	Description
result	String, enumerated	Common default result codes

Example.

```
{
    "output": {
        "result": "SUCCESS"
    }
}
```

get-param

The Object Parameter Reporting procedure allows the application to retrieve the following information:

1. the defined object types and instances within the Resource Model of the RE
2. the values of the parameters of the objects

Default URL: http://localhost:8181/restconf/operations/ocp-service:get-param-nb

POST Input

Field Name	Type	Description	Example	Required?
nodeId	String	Inventory node reference for OCP radio head	ocp:MTI-101-200	Yes
objId	String	Object ID	RxSigPath_5G:1	Yes
paramName	String	Parameter name	dataLink	Yes

Example.

```
{
    "get-param-nb": {
        "input": {
            "nodeId": "ocp:MTI-101-200",
            "objId": "RxSigPath_5G:1",
            "paramName": "dataLink"
        }
    }
}
```

POST Output

Field Name	Type	Description
id	String	Object ID
name	String	Object parameter name
value	String	Object parameter value
result	String, enumerated	Common default result codes + “FAIL_UNKNOWN_OBJECT”, “FAIL_UNKNOWN_PARAM”

Example.

```
{
    "output": {
        "obj": [
            {
                "id": "RxSigPath_5G:1",
                "param": [
                    {
                        "name": "dataLink",
                        "value": "dataLink:1"
                    }
                ]
```

(continues on next page)

(continued from previous page)

```
            }
        ],
        "result": "SUCCESS"
    }
}
```

modify-param

The Object Parameter Modification procedure allows the application to configure the values of the parameters of the objects identified by the Resource Model.

Default URL: http://localhost:8181/restconf/operations/ocp-service:modify-param-nb

POST Input

Field Name	Type	Description	Example	Required?
nodeId	String	Inventory node reference for OCP radio head	ocp:MTI-101-200	Yes
objId	String	Object ID	RxSigPath_5G:1	Yes
name	String	Object parameter name	dataLink	Yes
value	String	Object parameter value	dataLink:1	Yes

Example.

```
{
    "modify-param-nb": {
        "input": {
            "nodeId": "ocp:MTI-101-200",
            "objId": "RxSigPath_5G:1",
            "param": [
                {
                    "name": "dataLink",
                    "value": "dataLink:1"
                }
            ]
        }
    }
}
```

POST Output

Field Name	Type	Description
objId	String	Object ID
globResult	String, enumerated	Common default result codes + "FAIL_UNKNOWN_OBJECT", "FAIL_PARAMETER_FAIL", "FAIL_NOSUCH_RESOURCE"
name	String	Object parameter name
result	String, enumerated	"SUCCESS", "FAIL_UNKNOWN_PARAM", "FAIL_PARAM_READONLY", "FAIL_PARAM_LOCKREQUIRED", "FAIL_VALUE_OUTOF_RANGE", "FAIL_VALUE_TYPE_ERROR"

Example.

```
{
    "output": {
        "objId": "RxSigPath_5G:1",
        "globResult": "SUCCESS",
        "param": [
            {
                "name": "dataLink",
                "result": "SUCCESS"
            }
        ]
    }
}
```

create-obj

The Object Creation procedure allows the application to create and initialize a new instance of the given object type on the RE.

Default URL: http://localhost:8181/restconf/operations/ocp-service:create-obj-nb

POST Input

Field Name	Type	Description	Example	Required?
nodeId	String	Inventory node reference for OCP radio head	ocp:MTI-101-200	Yes
objType	String	Object type	RxSigPath_5G	Yes
name	String	Object parameter name	dataLink	No
value	String	Object parameter value	dataLink:1	No

Example.

```
{
    "create-obj-nb": {
        "input": {
            "nodeId": "ocp:MTI-101-200",
            "objType": "RxSigPath_5G",
```

(continues on next page)

(continued from previous page)

```
            "param": [
                {
                    "name": "dataLink",
                    "value": "dataLink:1"
                }
            ]
        }
    }
}
```

POST Output

Field Name	Type	Description
objId	String	Object ID
globResult	String, enumerated	Common default result codes + “FAIL_UNKNOWN_OBJTYPE”, “FAIL_STATIC_OBJTYPE”, “FAIL_UNKNOWN_OBJECT”, “FAIL_CHILD_NOTALLOWED”, “FAIL_OUTOF_RESOURCES” “FAIL_PARAMETER_FAIL”, “FAIL_NOSUCH_RESOURCE”
name	String	Object parameter name
result	String, enumerated	“SUCCESS”, “FAIL_UNKNOWN_PARAM”, “FAIL_PARAM_READONLY”, “FAIL_PARAM_LOCKREQUIRED”, “FAIL_VALUE_OUTOF_RANGE”, “FAIL_VALUE_TYPE_ERROR”

Example.

```
{
    "output": {
        "objId": "RxSigPath_5G:0",
        "globResult": "SUCCESS",
        "param": [
            {
                "name": "dataLink",
                "result": "SUCCESS"
            }
        ]
    }
}
```

delete-obj

The Object Deletion procedure allows the application to delete a given object instance and recursively its entire child objects on the RE.

Default URL: http://localhost:8181/restconf/operations/ocp-service:delete-obj-nb

POST Input

Field Name	Type	Description	Example	Required?
nodeId	String	Inventory node reference for OCP radio head	ocp:MTI-101-200	Yes
objId	String	Object ID	RxSigPath_5G:1	Yes

Example.

```
{
    "delete-obj-nb": {
        "input": {
            "nodeId": "ocp:MTI-101-200",
            "obj-id": "RxSigPath_5G:0"
        }
    }
}
```

POST Output

Field Name	Type	Description
result	String, enumerated	Common default result codes + "FAIL_UNKNOWN_OBJECT", "FAIL_STATIC_OBJTYPE", "FAIL_LOCKREQUIRED"

Example.

```
{
    "output": {
        "result": "SUCCESS"
    }
}
```

get-state

The Object State Reporting procedure allows the application to acquire the current state (for the requested state type) of one or more objects of the RE resource model, and additionally configure event-triggered reporting of the detected state changes for all state types of the indicated objects.

Default URL: http://localhost:8181/restconf/operations/ocp-service:get-state-nb

POST Input

Field Name	Type	Description	Example	Required ?
nodeId	String	Inventory node reference for OCP radio head	ocp:MTI-101-200	Yes
objId	String	Object ID	RxSig-Path_5G:1	Yes
stateType	String, enumerated	Valid values: “AST”, “FST”, “ALL”	ALL	Yes
eventDrivenReporting	Boolean	Event-triggered reporting of state change	true	Yes

Example.

```
{
    "get-state-nb": {
        "input": {
            "nodeId": "ocp:MTI-101-200",
            "objId": "antPort:0",
            "stateType": "ALL",
            "eventDrivenReporting": "true"
        }
    }
}
```

POST Output

Field Name	Type	Description
id	String	Object ID
type	String, enumerated	State type. Valid values: “AST”, “FST
value	String, enumerated	State value. Valid values: For state type = “AST”: “LOCKED”, “UNLOCKED”. For state type = “FST”: “PRE_OPERATIONAL”, “OPERATIONAL”, “DEGRADED”, “FAILED”, “NOT_OPERATIONAL”, “DISABLED”
result	String, enumerated	Common default result codes + “FAIL_UNKNOWN_OBJECT”, “FAIL_UNKNOWN_STATETYPE”, “FAIL_VALUE_OUTOF_RANGE”

Example.

```
{
    "output": {
        "obj": [
            {
                "id": "antPort:0",
```

(continues on next page)

(continued from previous page)

```
                "state": [
                    {
                        "type": "FST",
                        "value": "DISABLED"
                    },
                    {
                        "type": "AST",
                        "value": "LOCKED"
                    }
                ]
            }
        ],
        "result": "SUCCESS"
    }
}
```

modify-state

The Object State Modification procedure allows the application to trigger a change in the state of an object of the RE Resource Model.

Default URL: http://localhost:8181/restconf/operations/ocp-service:modify-state-nb

POST Input

Field Name	Type	Description	Example	Required?
nodeId	String	Inventory node reference for OCP radio head	ocp:MTI-101-200	Yes
objId	String	Object ID	RxSigPath_5G:1	Yes
stateType	String, enumerated	Valid values: “AST”, “FST”, “ALL”	AST	Yes
stateValue	String, enumerated	Valid values: For state type = “AST”: “LOCKED”, “UNLOCKED”. For state type = “FST”: “PRE_OPERATIONAL”, “OPERATIONAL”, “DEGRADED”, “FAILED”, “NOT_OPERATIONAL”, “DISABLED”	LOCKED	Yes

Example.

```
{
    "modify-state-nb": {
        "input": {
            "nodeId": "ocp:MTI-101-200",
            "objId": "RxSigPath_5G:1",
            "stateType": "AST",
            "stateValue": "LOCKED"
```

(continues on next page)

(continued from previous page)

```
        }
    }
}
```

POST Output

Field Name	Type	Description
objId	String	Object ID
stateType	String, enumerated	State type. Valid values: “AST”, “FST
stateValue	String, enumerated	State value. Valid values: For state type = “AST”: “LOCKED”, “UNLOCKED”. For state type = “FST”: “PRE_OPERATIONAL”, “OPERATIONAL”, “DEGRADED”, “FAILED”, “NOT_OPERATIONAL”, “DISABLED”
result	String, enumerated	Common default result codes + “FAIL_UNKNOWN_OBJECT”, “FAIL_UNKNOWN_STATETYPE”, “FAIL_UNKNOWN_STATEVALUE”, “FAIL_STATE_READONLY”, “FAIL_RESOURCE_UNAVAILABLE”, “FAIL_RESOURCE_INUSE”, “FAIL_PARENT_CHILD_CONFLICT”, “FAIL_PRECONDITION_NOTMET

Example.

```
{
    "output": {
        "objId": "RxSigPath_5G:1",
        "stateType": "AST",
        "stateValue": "LOCKED",
        "result": "SUCCESS",
    }
}
```

get-fault

The Fault Reporting procedure allows the application to acquire information about all current active faults associated with a primary object, as well as configure the RE to report when the fault status changes for any of faults associated with the indicated primary object.

Default URL: http://localhost:8181/restconf/operations/ocp-service:get-fault-nb

POST Input

Field Name	Type	Description	Example	Required?
nodeId	String	Inventory node reference for OCP radio head	ocp:MTI-101-200	Yes
objId	String	Object ID	RE:0	Yes
eventDrive nReporting	Boolean	Event-triggered reporting of fault	true	Yes

Example.

```
{
    "get-fault-nb": {
        "input": {
            "nodeId": "ocp:MTI-101-200",
            "objId": "RE:0",
            "eventDrivenReporting": "true"
        }
    }
}
```

POST Output

Field Name	Type	Description
result	String, enumerated	Common default result codes + "FAIL_UNKNOWN_OBJECT", "FAIL_VALUE_OUTOF_RANGE"
id (obj)	String	Object ID
id (fault)	String	Fault ID
severity	String	Fault severity
timestamp	dateTime	Time stamp
descr	String	Text description
affectedObj	String	Affected object

Example.

```
{
    "output": {
        "result": "SUCCESS",
        "obj": [
            {
                "id": "RE:0",
                "fault": [
                    {
                        "id": "FAULT_OVERTEMP",
                        "severity": "DEGRADED",
                        "timestamp": "2012-02-12T16:35:00",
                        "descr": "PA temp too high; Pout reduced",
```

(continues on next page)

(continued from previous page)

```
                        "affectedObj": [
                            "TxSigPath_EUTRA:0",
                            "TxSigPath_EUTRA:1"
                        ]
                    },
                    {
                        "id": "FAULT_VSWR_OUTOF_RANGE",
                        "severity": "WARNING",
                        "timestamp": "2012-02-12T16:01:05",
                    }
                ]
            }
        ],
    }
}
```

Note: The northbound APIs described above wrap the southbound APIs to make them accessible to external applications via RESTCONF, as well as take care of synchronizing the RE resource model between radio heads and the controller's datastore. See applications/ocp-service/src/main/yang/ocp-resourcemodel.yang for the yang representation of the RE resource model.

Java Interfaces (Southbound APIs)

The southbound APIs provide concrete implementation of the following OCP elementary functions: health-check, set-time, re-reset, get-param, modify-param, create-obj, delete-obj, get-state, modify-state and get-fault. Any OpenDaylight services/applications (of course, including OCP service) wanting to speak OCP to radio heads will need to use them.

SalDeviceMgmtService

Interface SalDeviceMgmtService defines three methods corresponding to health-check, set-time and re-reset.

SalDeviceMgmtService.java.

```
package org.opendaylight.yang.gen.v1.urn.opendaylight.ocp.device.mgmt.rev150811;

public interface SalDeviceMgmtService
    extends
    RpcService
{

    Future<RpcResult<HealthCheckOutput>> healthCheck(HealthCheckInput input);

    Future<RpcResult<SetTimeOutput>> setTime(SetTimeInput input);

    Future<RpcResult<ReResetOutput>> reReset(ReResetInput input);

}
```

SalConfigMgmtService

Interface SalConfigMgmtService defines two methods corresponding to get-param and modify-param.

SalConfigMgmtService.java.

```
package org.opendaylight.yang.gen.v1.urn.opendaylight.ocp.config.mgmt.rev150811;

public interface SalConfigMgmtService
    extends
    RpcService
{

    Future<RpcResult<GetParamOutput>> getParam(GetParamInput input);

    Future<RpcResult<ModifyParamOutput>> modifyParam(ModifyParamInput input);

}
```

SalObjectLifecycleService

Interface SalObjectLifecycleService defines two methods corresponding to create-obj and delete-obj.

SalObjectLifecycleService.java.

```
package org.opendaylight.yang.gen.v1.urn.opendaylight.ocp.object.lifecycle.rev150811;

public interface SalObjectLifecycleService
    extends
    RpcService
{

    Future<RpcResult<CreateObjOutput>> createObj(CreateObjInput input);

    Future<RpcResult<DeleteObjOutput>> deleteObj(DeleteObjInput input);

}
```

SalObjectStateMgmtService

Interface SalObjectStateMgmtService defines two methods corresponding to get-state and modify-state.

SalObjectStateMgmtService.java.

```
package org.opendaylight.yang.gen.v1.urn.opendaylight.ocp.object.state.mgmt.rev150811;

public interface SalObjectStateMgmtService
    extends
    RpcService
{

    Future<RpcResult<GetStateOutput>> getState(GetStateInput input);

    Future<RpcResult<ModifyStateOutput>> modifyState(ModifyStateInput input);
```

(continues on next page)

(continued from previous page)

```
}
```

SalFaultMgmtService

Interface SalFaultMgmtService defines only one method corresponding to get-fault.

SalFaultMgmtService.java.

```
package org.opendaylight.yang.gen.v1.urn.opendaylight.ocp.fault.mgmt.rev150811;

public interface SalFaultMgmtService
    extends
    RpcService
{

    Future<RpcResult<GetFaultOutput>> getFault(GetFaultInput input);

}
```

Notifications

In addition to indication messages, the OCP southbound plugin will translate specific events (e.g., connect, disconnect) coming up from the OCP protocol library into MD-SAL Notification objects and then publish them to the MD-SAL. Also, the OCP service will notify the completion of certain operation via Notification as well.

SalDeviceMgmtListener

An onDeviceConnected Notification will be published to the MD-SAL as soon as a radio head is connected to the controller, and when that radio head is disconnected the OCP southbound plugin will publish an onDeviceDisconnected Notification in response to the disconnect event propagated from the OCP protocol library.

SalDeviceMgmtListener.java.

```
package org.opendaylight.yang.gen.v1.urn.opendaylight.ocp.device.mgmt.rev150811;

public interface SalDeviceMgmtListener
    extends
    NotificationListener
{

    void onDeviceConnected(DeviceConnected notification);

    void onDeviceDisconnected(DeviceDisconnected notification);

}
```

OcpServiceListener

The OCP service will publish an onAlignmentCompleted Notification to the MD-SAL once it has completed the OCP alignment procedure with the radio head.

OcpServiceListener.java.

```
package org.opendaylight.yang.gen.v1.urn.opendaylight.params.xml.ns.yang.ocp.
↪applications.ocp.service.rev150811;

public interface OcpServiceListener
    extends
    NotificationListener
{

    void onAlignmentCompleted(AlignmentCompleted notification);

}
```

SalObjectStateMgmtListener

When receiving a state change indication message, the OCP southbound plugin will propagate the indication message to upper layer services/applications by publishing a corresponding onStateChangeInd Notification to the MD-SAL.

SalObjectStateMgmtListener.java.

```
package org.opendaylight.yang.gen.v1.urn.opendaylight.ocp.object.state.mgmt.rev150811;

public interface SalObjectStateMgmtListener
    extends
    NotificationListener
{

    void onStateChangeInd(StateChangeInd notification);

}
```

SalFaultMgmtListener

When receiving a fault indication message, the OCP southbound plugin will propagate the indication message to upper layer services/applications by publishing a corresponding onFaultInd Notification to the MD-SAL.

SalFaultMgmtListener.java.

```
package org.opendaylight.yang.gen.v1.urn.opendaylight.ocp.fault.mgmt.rev150811;

public interface SalFaultMgmtListener
    extends
    NotificationListener
{

    void onFaultInd(FaultInd notification);

}
```

ODL-SDNi Developer Guide

Overview

This project aims at enabling inter-SDN controller communication by developing SDNi (Software Defined Networking interface) as an application (ODL-SDNi App).

ODL-SDNi Architecture

- SDNi Aggregator: Northbound SDNi plugin acts as an aggregator for collecting network information such as topology, stats, host etc. This plugin can be evolving as per needs of network data requested to be shared across federated SDN controllers.
- SDNi API: API view autogenerated and accessible through RESTCONF to fetch the aggregated information from the northbound plugin – SDNi aggregator.The RESTCONF protocol operates on a conceptual datastore defined with the YANG data modeling language.
- SDNi Wrapper: SDNi BGP Wrapper will be responsible for the sharing and collecting information to/from federated controllers.
- SDNi UI:This component displays the SDN controllers connected to each other.

SDNi Aggregator

- SDNiAggregator connects with the Base Network Service Functions of the controller. Currently it is querying network topology through MD-SAL for creating SDNi network capability.
- SDNiAggregator is customized to retrieve the host controller's details, while running the controller in cluster mode. Rest of the northbound APIs of controller will retrieve the entire topology information of all the connected controllers.
- The SDNiAggregator creates a topology structure.This structure is populated by the various network funtions.

SDNi API

Topology and QoS data is fetched from SDNiAggregator through RESTCONF.

http://\protect\T1\textdollar\protect\T1\textbraceleftcontrolleripaddress\protect\T1\textbraceright:8181/apidoc/explorer/index.html

http://\protect\T1\textdollar\protect\T1\textbraceleftipaddress\protect\T1\textbraceright:8181/restconf/operations/opendaylight-sdni-topology-msg:getAllPeerTopology

Peer Topology Data: Controller IP Address, Links, Nodes, Link Bandwidths, MAC Address of switches, Latency, Host IP address.

http://\protect\T1\textdollar\protect\T1\textbraceleftipaddress\protect\T1\textbraceright:8181/restconf/operations/opendaylight-sdni-qos-msg:get-all-node-connectors-statistics

QOS Data: Node, Port, Transmit Packets, Receive Packets, Collision Count, Receive Frame Error, Receive Over Run Error, Receive Crc Error

http://\protect\T1\textdollar\protect\T1\textbraceleftipaddress\protect\T1\textbraceright:8181/restconf/operations/opendaylight-sdni-qos-msg:get-all-peer-node-connectors-statistics

Peer QOS Data: Node, Port, Transmit Packets, Receive Packets, Collision Count, Receive Frame Error, Receive Over Run Error, Receive Crc Error

SDNi Wrapper

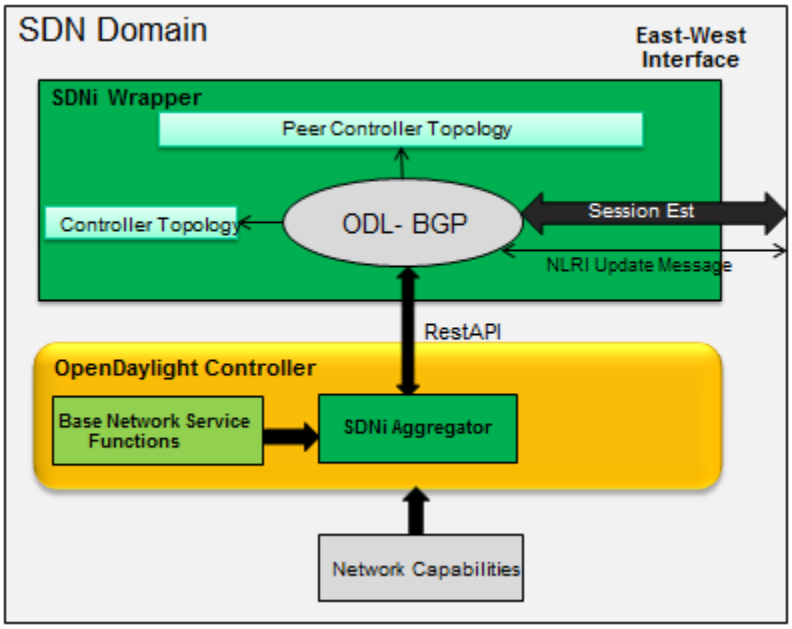

Fig. 26: SDNiWrapper

- SDNiWrapper is an extension of ODL-BGPCEP where SDNi topology data is exchange along with the Update NLRI message. Refer http://tools.ietf.org/html/draft-ietf-idr-ls-distribution-04 for more information on NLRI.
- SDNiWrapper gets the controller's network capabilities through SDNi Aggregator and serialize it in Update NLRI message. This NLRI message will get exchange between the clustered controllers through BGP-UPDATE message. Similarly peer controller's UPDATE message is received and unpacked then format to SDNi Network capability data, which will be stored for further purpose.

SDNi UI

This component displays the SDN controllers connected to each other.

http://localhost:8181/index.html#/sdniUI/sdnController

API Reference Documentation

Go to http://\protect\T1\textdollar\protect\T1\textbraceleftcontrolleripaddress\protect\T1\textbraceright:8181/apidoc/explorer/index.html, sign in, and expand the opendaylight-sdni panel. From there, users can execute various API calls to test their SDNi deployment.

OF-CONFIG Developer Guide

Overview

OF-CONFIG defines an OpenFlow switch as an abstraction called an OpenFlow Logical Switch. The OF-CONFIG protocol enables configuration of essential artifacts of an OpenFlow Logical Switch so that an OpenFlow controller

can communicate and control the OpenFlow Logical switch via the OpenFlow protocol. OF-CONFIG introduces an operating context for one or more OpenFlow data paths called an OpenFlow Capable Switch for one or more switches. An OpenFlow Capable Switch is intended to be equivalent to an actual physical or virtual network element (e.g. an Ethernet switch) which is hosting one or more OpenFlow data paths by partitioning a set of OpenFlow related resources such as ports and queues among the hosted OpenFlow data paths. The OF-CONFIG protocol enables dynamic association of the OpenFlow related resources of an OpenFlow Capable Switch with specific OpenFlow Logical Switches which are being hosted on the OpenFlow Capable Switch. OF-CONFIG does not specify or report how the partitioning of resources on an OpenFlow Capable Switch is achieved. OF-CONFIG assumes that resources such as ports and queues are partitioned amongst multiple OpenFlow Logical Switches such that each OpenFlow Logical Switch can assume full control over the resources that is assigned to it.

How to start

- start OF-CONFIG feature as below:

```
feature:install odl-of-config-all
```

Compatible with NETCONF

- Config OpenFlow Capable Switch via OpenFlow Configuration Points

  Method: POST

  URI: http://localhost:8181/restconf/config/network-topology:network-topology/topology/topology-netconf/node/controller-config/yang-ext:mount/config:modules

  Headers: Content-Type" and "Accept" header attributes set to application/xml

  Payload:

```
<module xmlns="urn:opendaylight:params:xml:ns:yang:controller:config">
  <type xmlns:prefix=
↪"urn:opendaylight:params:xml:ns:yang:controller:md:sal:connector:netconf">
↪prefix:sal-netconf-connector</type>
  <name>testtool</name>
  <address xmlns=
↪"urn:opendaylight:params:xml:ns:yang:controller:md:sal:connector:netconf">10.74.
↪151.67</address>
  <port xmlns=
↪"urn:opendaylight:params:xml:ns:yang:controller:md:sal:connector:netconf">830</
↪port>
  <username xmlns=
↪"urn:opendaylight:params:xml:ns:yang:controller:md:sal:connector:netconf">
↪mininet</username>
  <password xmlns=
↪"urn:opendaylight:params:xml:ns:yang:controller:md:sal:connector:netconf">
↪mininet</password>
  <tcp-only xmlns=
↪"urn:opendaylight:params:xml:ns:yang:controller:md:sal:connector:netconf">false
↪</tcp-only>
  <event-executor xmlns=
↪"urn:opendaylight:params:xml:ns:yang:controller:md:sal:connector:netconf">
    <type xmlns:prefix="urn:opendaylight:params:xml:ns:yang:controller:netty">
↪prefix:netty-event-executor</type>
    <name>global-event-executor</name>
```

(continues on next page)

(continued from previous page)

```
  </event-executor>
  <binding-registry xmlns=
↪"urn:opendaylight:params:xml:ns:yang:controller:md:sal:connector:netconf">
    <type xmlns:prefix=
↪"urn:opendaylight:params:xml:ns:yang:controller:md:sal:binding">prefix:binding-
↪broker-osgi-registry</type>
    <name>binding-osgi-broker</name>
  </binding-registry>
  <dom-registry xmlns=
↪"urn:opendaylight:params:xml:ns:yang:controller:md:sal:connector:netconf">
    <type xmlns:prefix="urn:opendaylight:params:xml:ns:yang:controller:md:sal:dom
↪">prefix:dom-broker-osgi-registry</type>
    <name>dom-broker</name>
  </dom-registry>
  <client-dispatcher xmlns=
↪"urn:opendaylight:params:xml:ns:yang:controller:md:sal:connector:netconf">
    <type xmlns:prefix=
↪"urn:opendaylight:params:xml:ns:yang:controller:config:netconf">prefix:netconf-
↪client-dispatcher</type>
    <name>global-netconf-dispatcher</name>
  </client-dispatcher>
  <processing-executor xmlns=
↪"urn:opendaylight:params:xml:ns:yang:controller:md:sal:connector:netconf">
    <type xmlns:prefix="urn:opendaylight:params:xml:ns:yang:controller:threadpool
↪">prefix:threadpool</type>
    <name>global-netconf-processing-executor</name>
  </processing-executor>
</module>
```

- NETCONF establishes the connections with OpenFlow Capable Switches using the parameters in the previous step. NETCONF also gets the information of whether the OpenFlow Switch supports NETCONF during the signal handshaking. The information will be stored in the NETCONF topology as prosperity of a node.
- OF-CONFIG can be aware of the switches accessing and leaving by monitoring the data changes in the NETCONF topology. For the detailed information it can be refered to the implementation.

The establishment of OF-CONFIG topology

Firstly, OF-CONFIG will check whether the newly accessed switch supports OF-CONFIG by querying the NETCONF interface.

1. During the NETCONF connection's establishment, the NETCONF and the switches will exchange the their capabilities via the "hello" message.
2. OF-CONFIG gets the connection information between the NETCONF and switches by monitoring the data changes via the interface of DataChangeListener.
3. After the NETCONF connection established, the OF-CONFIG module will check whether OF-CONFIG capability is in the switch's capabilities list which is got in step 1.
4. If the result of step 3 is yes, the OF-CONFIG will call the following processing steps to create the topology database.

For the detailed information it can be referred to the implementation.

Secondly, the capable switch node and logical switch node are added in the OF-CONFIG topology if the switch supports OF-CONFIG.

OF-CONFIG's topology compromise: Capable Switch topology (underlay) and logical Switch topology (overlay). Both of them are enhanced (augment) on

/topo:network-topology/topo:topology/topo:node

The NETCONF will add the nodes in the Topology via the path of "/topo:network-topology/topo:topology/topo:node" if it gets the configuration information of the switches.

For the detailed information it can be referred to the implementation.

OpenFlow Protocol Library Developer Guide

Introduction

OpenFlow Protocol Library is component in OpenDaylight, that mediates communication between OpenDaylight controller and hardware devices supporting OpenFlow protocol. Primary goal is to provide user (or upper layers of OpenDaylight) communication channel, that can be used for managing network hardware devices.

Features Overview

There are three features inside openflowjava:

- **odl-openflowjava-protocol** provides all openflowjava bundles, that are needed for communication with openflow devices. It ensures message translation and handles network connections. It also provides openflow protocol specific model.
- **odl-openflowjava-all** currently contains only odl-openflowjava-protocol feature.
- **odl-openflowjava-stats** provides mechanism for message counting and reporting. Can be used for performance analysis.

odl-openflowjava-protocol Architecture

Basic bundles contained in this feature are openflow-protocol-api, openflow-protocol-impl, openflow-protocol-spi and util.

- **openflow-protocol-api** - contains openflow model, constants and keys used for (de)serializer registration.
- **openflow-protocol-impl** - contains message factories, that translate binary messages into DataObjects and vice versa. Bundle also contains network connection handlers - servers, netty pipeline handlers, ...
- **openflow-protocol-spi** - entry point for openflowjava configuration, startup and close. Basically starts implementation.
- **util** - utility classes for binary-Java conversions and to ease experimenter key creation

odl-openflowjava-stats Feature

Runs over odl-openflowjava-protocol. It counts various message types / events and reports counts in specified time periods. Statistics collection can be configured in openflowjava-config/src/main/resources/45-openflowjava-stats.xml

Key APIs and Interfaces

Basic API / SPI classes are ConnectionAdapter (Rpc/notifications) and SwitchConnectionProcider (configure, start, shutdown)

Installation

Pull the code and import project into your IDE.

```
git clone ssh://<username>@git.opendaylight.org:29418/openflowjava.git
```

Configuration

Current implementation allows to configure:

- listening port (mandatory)
- transfer protocol (mandatory)
- switch idle timeout (mandatory)
- TLS configuration (optional)
- thread count (optional)

You can find exemplary Openflow Protocol Library instance configuration below:

```
<data xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <modules xmlns="urn:opendaylight:params:xml:ns:yang:controller:config">
    <!-- default OF-switch-connection-provider (port 6633) -->
    <module>
      <type xmlns:prefix=
↪"urn:opendaylight:params:xml:ns:yang:openflow:switch:connection:provider:impl">
↪prefix:openflow-switch-connection-provider-impl</type>
      <name>openflow-switch-connection-provider-default-impl</name>
      <port>6633</port>
<!--  Possible transport-protocol options: TCP, TLS, UDP -->
      <transport-protocol>TCP</transport-protocol>
      <switch-idle-timeout>15000</switch-idle-timeout>
<!--       Exemplary TLS configuration:
            - uncomment the <tls> tag
            - copy exemplary-switch-privkey.pem, exemplary-switch-cert.pem and␣
↪exemplary-cacert.pem
              files into your virtual machine
            - set VM encryption options to use copied keys
            - start communication
           Please visit OpenflowPlugin or Openflow Protocol Library#Documentation␣
↪wiki pages
           for detailed information regarding TLS -->
<!--       <tls>
             <keystore>/exemplary-ctlKeystore</keystore>
             <keystore-type>JKS</keystore-type>
             <keystore-path-type>CLASSPATH</keystore-path-type>
             <keystore-password>opendaylight</keystore-password>
             <truststore>/exemplary-ctlTrustStore</truststore>
             <truststore-type>JKS</truststore-type>
```

(continues on next page)

(continued from previous page)

```
            <truststore-path-type>CLASSPATH</truststore-path-type>
            <truststore-password>opendaylight</truststore-password>
            <certificate-password>opendaylight</certificate-password>
          </tls> -->
<!--      Exemplary thread model configuration. Uncomment <threads> tag below to␣
↪adjust default thread model -->
<!--      <threads>
            <boss-threads>2</boss-threads>
            <worker-threads>8</worker-threads>
          </threads> -->
    </module>
```

```
    <!-- default OF-switch-connection-provider (port 6653) -->
    <module>
      <type xmlns:prefix=
↪"urn:opendaylight:params:xml:ns:yang:openflow:switch:connection:provider:impl">
↪prefix:openflow-switch-connection-provider-impl</type>
      <name>openflow-switch-connection-provider-legacy-impl</name>
      <port>6653</port>
<!--  Possible transport-protocol options: TCP, TLS, UDP -->
      <transport-protocol>TCP</transport-protocol>
      <switch-idle-timeout>15000</switch-idle-timeout>
<!--      Exemplary TLS configuration:
           - uncomment the <tls> tag
           - copy exemplary-switch-privkey.pem, exemplary-switch-cert.pem and␣
↪exemplary-cacert.pem
             files into your virtual machine
           - set VM encryption options to use copied keys
           - start communication
          Please visit OpenflowPlugin or Openflow Protocol Library#Documentation␣
↪wiki pages
          for detailed information regarding TLS -->
<!--      <tls>
            <keystore>/exemplary-ctlKeystore</keystore>
            <keystore-type>JKS</keystore-type>
            <keystore-path-type>CLASSPATH</keystore-path-type>
            <keystore-password>opendaylight</keystore-password>
            <truststore>/exemplary-ctlTrustStore</truststore>
            <truststore-type>JKS</truststore-type>
            <truststore-path-type>CLASSPATH</truststore-path-type>
            <truststore-password>opendaylight</truststore-password>
            <certificate-password>opendaylight</certificate-password>
          </tls> -->
<!--      Exemplary thread model configuration. Uncomment <threads> tag below to␣
↪adjust default thread model -->
<!--      <threads>
            <boss-threads>2</boss-threads>
            <worker-threads>8</worker-threads>
          </threads> -->
    </module>
```

```
  <module>
    <type xmlns:prefix=
↪"urn:opendaylight:params:xml:ns:yang:openflow:common:config:impl">prefix:openflow-
↪provider-impl</type>
    <name>openflow-provider-impl</name>
```

(continues on next page)

(continued from previous page)

```
    <openflow-switch-connection-provider>
      <type xmlns:ofSwitch=
↪"urn:opendaylight:params:xml:ns:yang:openflow:switch:connection:provider">
↪ofSwitch:openflow-switch-connection-provider</type>
      <name>openflow-switch-connection-provider-default</name>
    </openflow-switch-connection-provider>
    <openflow-switch-connection-provider>
      <type xmlns:ofSwitch=
↪"urn:opendaylight:params:xml:ns:yang:openflow:switch:connection:provider">
↪ofSwitch:openflow-switch-connection-provider</type>
      <name>openflow-switch-connection-provider-legacy</name>
    </openflow-switch-connection-provider>
    <binding-aware-broker>
      <type xmlns:binding=
↪"urn:opendaylight:params:xml:ns:yang:controller:md:sal:binding">binding:binding-
↪broker-osgi-registry</type>
      <name>binding-osgi-broker</name>
    </binding-aware-broker>
  </module>
</modules>
```

Possible transport-protocol options:

- TCP
- TLS
- UDP

Switch-idle timeout specifies time needed to detect idle state of switch. When no message is received from switch within this time, upper layers are notified on switch idleness. To be able to use this exemplary TLS configuration:

- uncomment the `<tls>` tag
- copy *exemplary-switch-privkey.pem*, *exemplary-switch-cert.pem* and *exemplary-cacert.pem* files into your virtual machine
- set VM encryption options to use copied keys (please visit TLS support wiki page for detailed information regarding TLS)
- start communication

Thread model configuration specifies how many threads are desired to perform Netty's I/O operations.

- boss-threads specifies the number of threads that register incoming connections
- worker-threads specifies the number of threads performing read / write (+ serialization / deserialization) operations.

Architecture

Public API `(openflow-protocol-api)`

Set of interfaces and builders for immutable data transfer objects representing Openflow Protocol structures.

Transfer objects and service APIs are infered from several YANG models using code generator to reduce verbosity of definition and repeatibility of code.

The following YANG modules are defined:

- openflow-types - defines common Openflow specific types
- openflow-instruction - defines base Openflow instructions
- openflow-action - defines base Openflow actions
- openflow-augments - defines object augmentations
- openflow-extensible-match - defines Openflow OXM match
- openflow-protocol - defines Openflow Protocol messages
- system-notifications - defines system notification objects
- openflow-configuration - defines structures used in ConfigSubsystem

This modules also reuse types from following YANG modules:

- ietf-inet-types - IP adresses, IP prefixes, IP-protocol related types
- ietf-yang-types - Mac Address, etc.

The use of predefined types is to make APIs contracts more safe, better readable and documented (e.g using MacAddress instead of byte array...)

TCP Channel pipeline `(openflow-protocol-impl)`

Creates channel processing pipeline based on configuration and support.

TCP Channel pipeline.

imageopenflowjava/500px-TCPChannelPipeline.png[width=500]

Switch Connection Provider.

Implementation of connection point for other projects. Library exposes its functionality through this class. Library can be configured, started and shutdowned here. There are also methods for custom (de)serializer registration.

Tcp Connection Initializer.

In order to initialize TCP connection to a device (switch), OF Plugin calls method `initiateConnection()` in `SwitchConnectionProvider`. This method in turn initializes (Bootstrap) server side channel towards the device.

TCP Handler.

Represents single server that is handling incoming connections over TCP / TLS protocol. TCP Handler creates a single instance of TCP Channel Initializer that will initialize channels. After that it binds to configured InetAddress and port. When a new device connects, TCP Handler registers its channel and passes control to TCP Channel Initializer.

TCP Channel Initializer.

This class is used for channel initialization / rejection and passing arguments. After a new channel has been registered it calls Switch Connection Handler's (OF Plugin) accept method to decide if the library should keep the newly registered channel or if the channel should be closed. If the channel has been accepted, TCP Channel Initializer creates the whole pipeline with needed handlers and also with ConnectionAdapter instance. After the channel pipeline is ready, Switch Connection Handler is notified with `onConnectionReady` notification. OpenFlow Plugin can now start sending messages downstream.

Idle Handler.

If there are no messages received for more than time specified, this handler triggers idle state notification. The switch idle timeout is received as a parameter from ConnectionConfiguration settings. Idle State Handler is inactive while

there are messages received within the switch idle timeout. If there are no messages received for more than timeout specified, handler creates SwitchIdleEvent message and sends it upstream.

TLS Handler.

It encrypts and decrypts messages over TLS protocol. Engaging TLS Handler into pipeline is matter of configuration (`<tls>` tag). TLS communication is either unsupported or required. TLS Handler is represented as a Netty's SslHandler.

OF Frame Decoder.

Parses input stream into correct length message frames for further processing. Framing is based on Openflow header length. If received message is shorter than minimal length of OpenFlow message (8 bytes), OF Frame Decoder waits for more data. After receiving at least 8 bytes the decoder checks length in OpenFlow header. If there are still some bytes missing, the decoder waits for them. Else the OF Frame Decoder sends correct length message to next handler in the channel pipeline.

OF Version Detector.

Detects version of used OpenFlow Protocol and discards unsupported version messages. If the detected version is supported, OF Version Detector creates `VersionMessageWrapper` object containing the detected version and byte message and sends this object upstream.

OF Decoder.

Chooses correct deserilization factory (based on message type) and deserializes messages into generated DTOs (Data Transfer Object). OF Decoder receives `VersionMessageWrapper` object and passes it to `DeserializationFactory` which will return translated DTO. `DeserializationFactory` creates `MessageCodeKey` object with version and type of received message and Class of object that will be the received message deserialized into. This object is used as key when searching for appropriate decoder in `DecoderTable`. `DecoderTable` is basically a map storing decoders. Found decoder translates received message into DTO. If there was no decoder found, null is returned. After returning translated DTO back to OF Decoder, the decoder checks if it is null or not. When the DTO is null, the decoder logs this state and throws an Exception. Else it passes the DTO further upstream. Finally, the OF Decoder releases ByteBuf containing received and decoded byte message.

OF Encoder.

Chooses correct serialization factory (based on type of DTO) and serializes DTOs into byte messages. OF Encoder does the opposite than the OF Decoder using the same principle. OF Encoder receives DTO, passes it for translation and if the result is not null, it sends translated DTO downstream as a ByteBuf. Searching for appropriate encoder is done via MessageTypeKey, based on version and class of received DTO.

Delegating Inbound Handler.

Delegates received DTOs to Connection Adapter. It also reacts on channelInactive and channelUnregistered events. Upon one of these events is triggered, DelegatingInboundHandler creates DisconnectEvent message and sends it upstream, notifying upper layers about switch disconnection.

Channel Outbound Queue.

Message flushing handler. Stores outgoing messages (DTOs) and flushes them. Flush is performed based on time expired and on the number of messages enqueued.

Connection Adapter.

Provides a facade on top of pipeline, which hides netty.io specifics. Provides a set of methods to register for incoming messages and to send messages to particular channel / session. ConnectionAdapterImpl basically implements three interfaces (unified in one superinterface ConnectionFacade):

- ConnectionAdapter
- MessageConsumer

- OpenflowProtocolService

ConnectionAdapter interface has methods for setting up listeners (message, system and connection ready listener), method to check if all listeners are set, checking if the channel is alive and disconnect method. Disconnect method clears responseCache and disables consuming of new messages.

MessageConsumer interface holds only one method: `consume()`. `Consume()` method is called from DelegatingInboundHandler. This method processes received DTO's based on their type. There are three types of received objects:

- System notifications - invoke system notifications in OpenFlow Plugin (systemListener set). In case of `DisconnectEvent` message, the Connection Adapter clears response cache and disables consume() method processing,
- OpenFlow asynchronous messages (from switch) - invoke corresponding notifications in OpenFlow Plugin,
- OpenFlow symmetric messages (replies to requests) - create `RpcResponseKey` with XID and DTO's class set. This `RpcResponseKey` is then used to find corresponding future object in responseCache. Future object is set with success flag, received message and errors (if any occurred). In case no corresponding future was found in responseCache, Connection Adapter logs warning and discards the message. Connection Adapter also logs warning when an unknown DTO is received.

OpenflowProtocolService interface contains all rpc-methods for sending messages from upper layers (OpenFlow Plugin) downstream and responding. Request messages return Future filled with expected reply message, otherwise the expected Future is of type Void.

NOTE: MultipartRequest message is the only exception. Basically it is request - reply Message type, but it wouldn't be able to process more following MultipartReply messages if this was implemented as rpc (only one Future). This is why MultipartReply is implemented as notification. OpenFlow Plugin takes care of correct message processing.

UDP Channel pipeline (openflow-protocol-impl)

Creates UDP channel processing pipeline based on configuration and support. **Switch Connection Provider**, **Channel Outbound Queue** and **Connection Adapter** fulfill the same role as in case of TCP connection / channel pipeline (please see above).

UDP Handler.

Represents single server that is handling incoming connections over UDP (DTLS) protocol. UDP Handler creates a single instance of UDP Channel Initializer that will initialize channels. After that it binds to configured InetAddress and port. When a new device connects, UDP Handler registers its channel and passes control to UDP Channel Initializer.

UDP Channel Initializer.

This class is used for channel initialization and passing arguments. After a new channel has been registered (for UDP there is always only one channel) UDP Channel Initializer creates whole pipeline with needed handlers.

DTLS Handler.

Haven't been implemented yet. Will take care of secure DTLS connections.

OF Datagram Packet Handler.

Combines functionality of OF Frame Decoder and OF Version Detector. Extracts messages from received datagram packets and checks if message version is supported. If there is a message received from yet unknown sender, OF Datagram Packet Handler creates Connection Adapter for this sender and stores it under sender's address in `UdpConnectionMap`. This map is also used for sending the messages and for correct Connection Adapter lookup - to delegate messages from one channel to multiple sessions.

OF Datagram Packet Decoder.

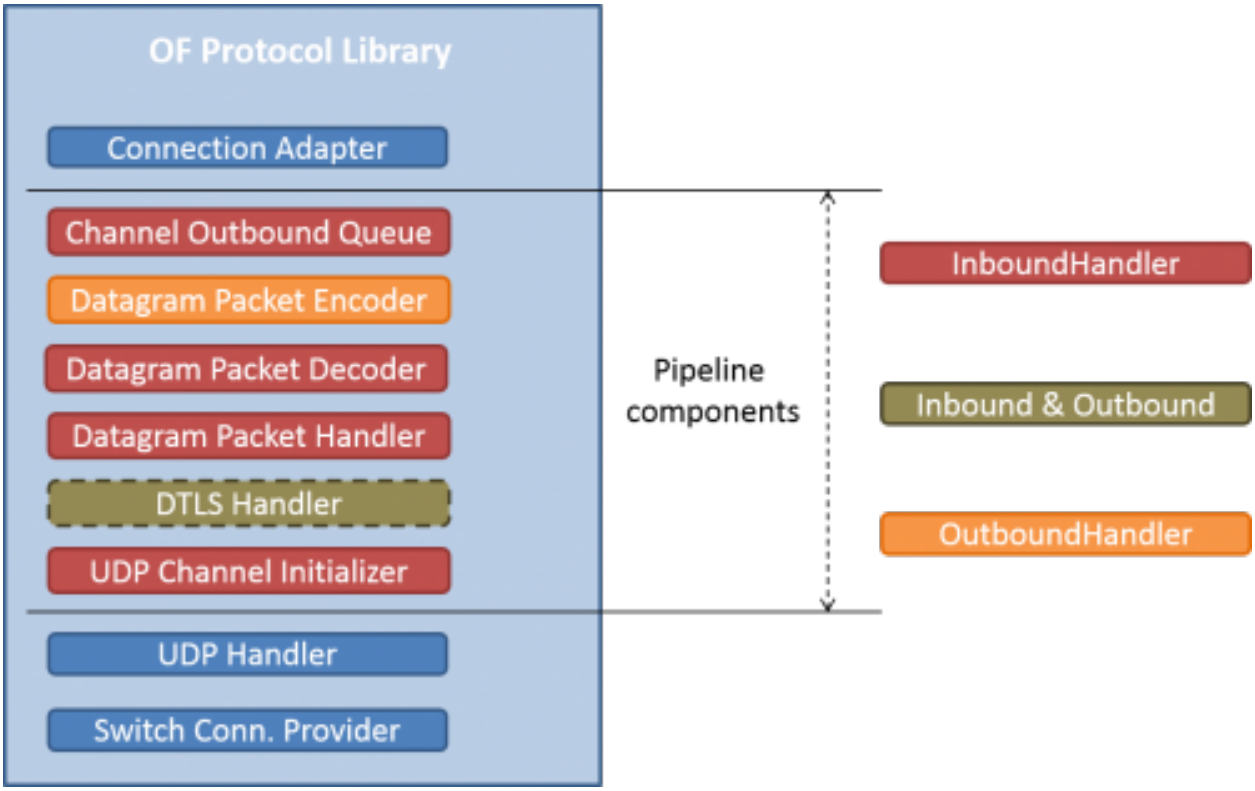

Fig. 27: UDP Channel pipeline

Chooses correct deserilization factory (based on message type) and deserializes messages into generated DTOs. OF Decoder receives `VersionMessageUdpWrapper` object and passes it to `DeserializationFactory` which will return translated DTO. `DeserializationFactory` creates `MessageCodeKey` object with version and type of received message and Class of object that will be the received message deserialized into. This object is used as key when searching for appropriate decoder in `DecoderTable`. `DecoderTable` is basically a map storing decoders. Found decoder translates received message into DTO (DataTransferObject). If there was no decoder found, null is returned. After returning translated DTO back to OF Datagram Packet Decoder, the decoder checks if it is null or not. When the DTO is null, the decoder logs this state. Else it looks up appropriate Connection Adapter in `UdpConnectionMap` and passes the DTO to found Connection Adapter. Finally, the OF Decoder releases `ByteBuf` containing received and decoded byte message.

OF Datagram Packet Encoder.

Chooses correct serialization factory (based on type of DTO) and serializes DTOs into byte messages. OF Datagram Packet Encoder does the opposite than the OF Datagram Packet Decoder using the same principle. OF Encoder receives DTO, passes it for translation and if the result is not null, it sends translated DTO downstream as a datagram packet. Searching for appropriate encoder is done via MessageTypeKey, based on version and class of received DTO.

SPI (openflow-protocol-spi)

Defines interface for library's connection point for other projects. Library exposes its functionality through this interface.

Integration test (openflow-protocol-it)

Testing communication with simple client.

Simple client(simple-client)

Lightweight switch simulator - programmable with desired scenarios.

Utility (util)

Contains utility classes, mainly for work with ByteBuf.

Library's lifecycle

Steps (after the library's bundle is started):

- [1] Library is configured by ConfigSubsystem (adress, ports, encryption, . . .)
- [2] Plugin injects its SwitchConnectionHandler into the Library
- [3] Plugin starts the Library
- [4] Library creates configured protocol handler (e.g. TCP Handler)
- [5] Protocol Handler creates Channel Initializer
- [6] Channel Initializer asks plugin whether to accept incoming connection on each new switch connection
- [7] Plugin responds:
  - true - continue building pipeline
  - false - reject connection / disconnect channel
- [8] Library notifies Plugin with onSwitchConnected(ConnectionAdapter) notification, passing reference to ConnectionAdapter, that will handle the connection
- [9] Plugin registers its system and message listeners
- [10] FireConnectionReadyNotification() is triggered, announcing that pipeline handlers needed for communication have been created and Plugin can start communication
- [11] Plugin shutdowns the Library when desired

Statistics collection

Introduction

Statistics collection collects message statistics. Current collected statistics (`DS` - downstream, `US` - upstream):

- `DS_ENTERED_OFJAVA` - all messages that entered openflowjava (picked up from openflowplugin)
- `DS_ENCODE_SUCCESS` - successfully encoded messages
- `DS_ENCODE_FAIL` - messages that failed during encoding (serialization) process
- `DS_FLOW_MODS_ENTERED` - all flow-mod messages that entered openflowjava

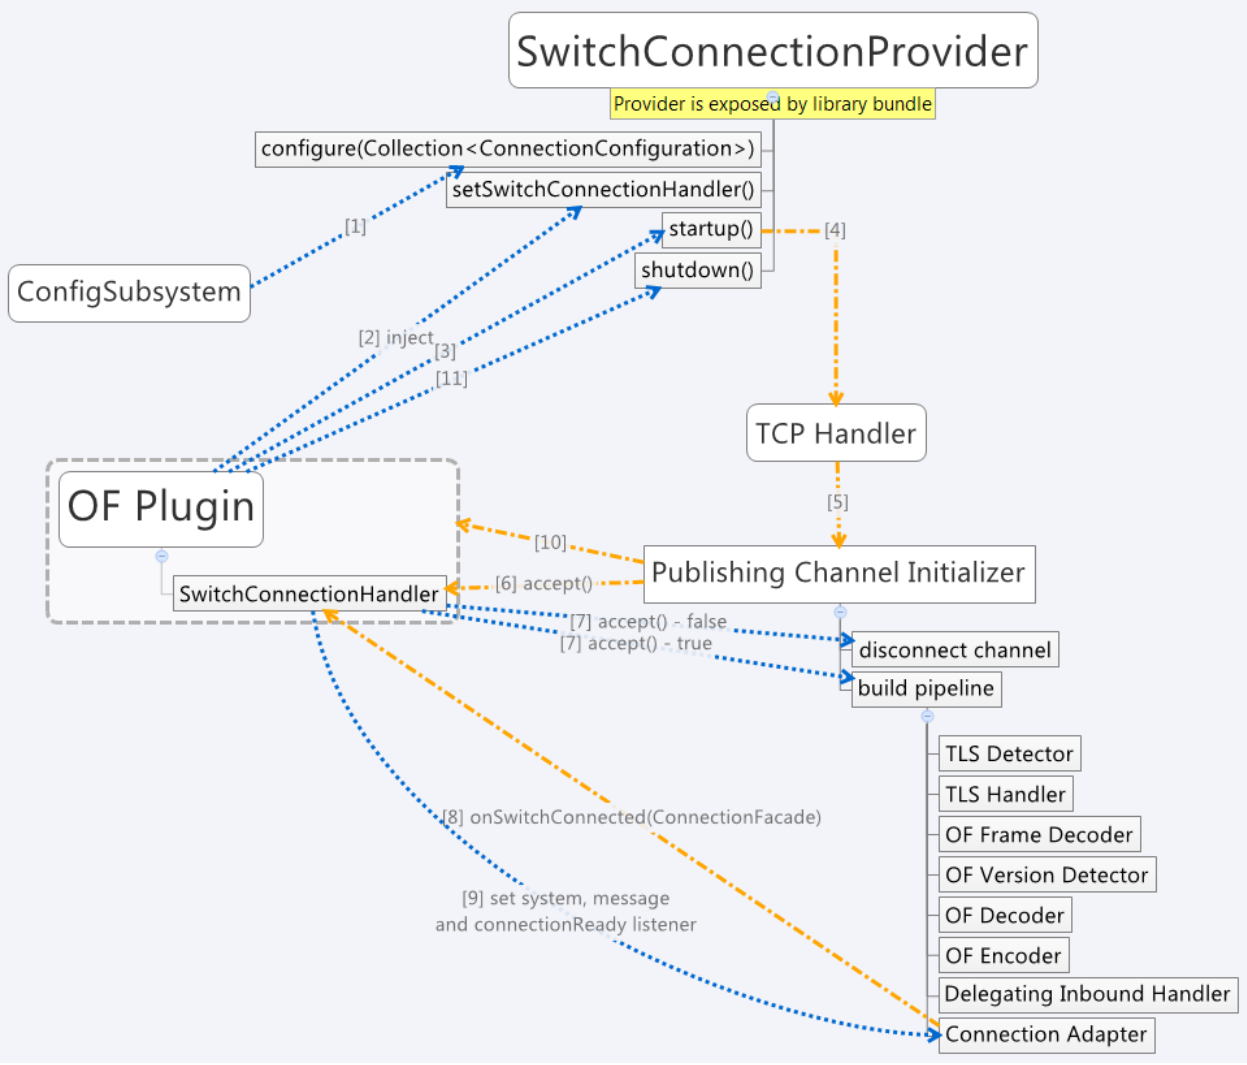

Fig. 28: Library lifecycle

- `DS_FLOW_MODS_SENT` - all flow-mod messages that were successfully sent
- `US_RECEIVED_IN_OFJAVA` - messages received from switch
- `US_DECODE_SUCCESS` - successfully decoded messages
- `US_DECODE_FAIL` - messages that failed during decoding (deserialization) process
- `US_MESSAGE_PASS` - messages handed over to openflowplugin

Karaf

In orded to start statistics, it is needed to feature:install odl-openflowjava-stats. To see the logs one should use log:set DEBUG org.opendaylight.openflowjava.statistics and than probably log:display (you can log:list to see if the logging has been set). To adjust collection settings it is enough to modify 45-openflowjava-stats.xml.

JConsole

JConsole provides two commands for the statistics collection:

- printing current statistics
- resetting statistic counters

After attaching JConsole to correct process, one only needs to go into MBeans `tab` → `org.opendaylight.controller` → `RuntimeBean` → `statistics-collection-service-impl` → `statistics-collection-service-impl` → `Operations` to be able to use this commands.

TLS Support

Note: see OpenFlow Plugin Developper Guide

Extensibility

Introduction

Entry point for the extensibility is `SwitchConnectionProvider`. `SwitchConnectionProvider` contains methods for (de)serializer registration. To register deserializer it is needed to use .register*Deserializer(key, impl). To register serializer one must use .register*Serializer(key, impl). Registration can occur either during configuration or at runtime.

NOTE: In case when experimenter message is received and no (de)serializer was registered, the library will throw `IllegalArgumentException`.

Basic Principle

In order to use extensions it is needed to augment existing model and register new (de)serializers.

Augmenting the model: 1. Create new augmentation

Register (de)serializers: 1. Create your (de)serializer 2. Let it implement `OFDeserializer<>` / `OFSerializer<>` - in case the structure you are (de)serializing needs to be used in Multipart TableFeatures messages, let it implement `HeaderDeserializer<>` / `HeaderSerializer` 3. Implement prescribed methods 4. Register your deserializer under appropriate key (in our case `ExperimenterActionDeserializerKey`) 5. Register your serializer under appropriate key (in our case `ExperimenterActionSerializerKey`) 6. Done, test your implementation

NOTE: If you don't know what key should be used with your (de)serializer implementation, please visit *Registration keys* page.

Example

Let's say we have vendor / experimenter action represented by this structure:

```
struct foo_action {
    uint16_t type;
    uint16_t length;
    uint32_t experimenter;
    uint16_t first;
    uint16_t second;
    uint8_t  pad[4];
}
```

First, we have to augment existing model. We create new module, which imports "`openflow-types.yang`" (don't forget to update your `pom.xml` with api dependency). Now we create foo action identity:

```
import openflow-types {prefix oft;}
identity foo {
    description "Foo action description";
    base oft:action-base;
}
```

This will be used as type in our structure. Now we must augment existing action structure, so that we will have the desired fields first and second. In order to create new augmentation, our module has to import "`openflow-action.yang`". The augment should look like this:

```
import openflow-action {prefix ofaction;}
augment "/ofaction:actions-container/ofaction:action" {
    ext:augment-identifier "foo-action";
        leaf first {
            type uint16;
        }
        leaf second {
            type uint16;
        }
    }
```

We are finished with model changes. Run mvn clean compile to generate sources. After generation is done, we need to implement our (de)serializer.

Deserializer:

```
public class FooActionDeserializer extends OFDeserializer<Action> {
   @Override
   public Action deserialize(ByteBuf input) {
       ActionBuilder builder = new ActionBuilder();
```

(continues on next page)

(continued from previous page)

```
        input.skipBytes(SIZE_OF_SHORT_IN_BYTES); *// we know the type of action*
        builder.setType(Foo.class);
        input.skipBytes(SIZE_OF_SHORT_IN_BYTES); *// we don't need length*
        *// now create experimenterIdAugmentation - so that openflowplugin can
        differentiate correct vendor codec*
        ExperimenterIdActionBuilder expIdBuilder = new ExperimenterIdActionBuilder();
        expIdBuilder.setExperimenter(new ExperimenterId(input.readUnsignedInt()));
        builder.addAugmentation(ExperimenterIdAction.class, expIdBuilder.build());
        FooActionBuilder fooBuilder = new FooActionBuilder();
        fooBuilder.setFirst(input.readUnsignedShort());
        fooBuilder.setSecond(input.readUnsignedShort());
        builder.addAugmentation(FooAction.class, fooBuilder.build());
        input.skipBytes(4); *// padding*
        return builder.build();
    }
}
```

Serializer:

```
public class FooActionSerializer extends OFSerializer<Action> {
    @Override
    public void serialize(Action action, ByteBuf outBuffer) {
        outBuffer.writeShort(FOO_CODE);
        outBuffer.writeShort(16);
        *// we don't have to check for ExperimenterIdAction augmentation - our
        serializer*
        *// was called based on the vendor / experimenter ID, so we simply write
        it to buffer*
        outBuffer.writeInt(VENDOR / EXPERIMENTER ID);
        FooAction foo = action.getAugmentation(FooAction.class);
        outBuffer.writeShort(foo.getFirst());
        outBuffer.writeShort(foo.getSecond());
        outBuffer.writeZero(4); //write padding
    }
}
```

Register both deserializer and serializer: `SwitchConnectionProvider.registerDeserializer(new ExperimenterActionDeserializerKey(0x04, VENDOR / EXPERIMENTER ID), new FooActionDeserializer());` `SwitchConnectionProvider.registerSerializer(new ExperimenterActionSerializerKey(0x04, VENDOR / EXPERIMENTER ID), new FooActionSerializer());`

We are ready to test our implementation.

NOTE: Vendor / Experimenter structures define only vendor / experimenter ID as common distinguisher (besides action type). Vendor / Experimenter ID is unique for all vendor messages - that's why vendor is able to register only one class under ExperimenterAction(De)SerializerKey. And that's why vendor has to switch / choose between his subclasses / subtypes on his own.

Detailed walkthrough: Deserialization extensibility

External interface & class description.

OFGeneralDeserializer:

- `OFDeserializer<E extends DataObject>`

  - *deserialize(ByteBuf)* - deserializes given ByteBuf

- `HeaderDeserializer<E extends DataObject>`
  - *deserializeHeaders(ByteBuf)* - deserializes only E headers (used in Multipart TableFeatures messages)

DeserializerRegistryInjector

- `injectDeserializerRegistry(DeserializerRegistry)` - injects deserializer registry into deserializer. Useful when custom deserializer needs access to other deserializers.

NOTE: DeserializerRegistryInjector is not OFGeneralDeserializer descendand. It is a standalone interface.

MessageCodeKey and its descendants These keys are used as for deserializer lookup in DeserializerRegistry. MessageCodeKey should is used in general, while its descendants are used in more special cases. For Example ActionDeserializerKey is used for Action deserializer lookup and (de)registration. Vendor is provided with special keys, which contain only the most necessary fields. These keys usually start with "Experimenter" prefix (MatchEntryDeserializerKey is an exception).

MessageCodeKey has these fields:

- short version - Openflow wire version number
- int value - value read from byte message
- Class<?> clazz - class of object being creating
- [1] The scenario starts in a custom bundle which wants to extend library's functionality. The custom bundle creates deserializers which implement exposed `OFDeserializer` / `HeaderDeserializer` interfaces (wrapped under `OFGeneralDeserializer` unifying super interface).
- [2] Created deserializers are paired with corresponding ExperimenterKeys, which are used for deserializer lookup. If you don't know what key should be used with your (de)serializer implementation, please visit *Registration keys* page.
- [3] Paired deserializers are passed to the OF Library via **SwitchConnectionProvider**.*registerCustomDeserializer(key, impl)*. Library registers the deserializer.
  - While registering, Library checks if the deserializer is an instance of **DeserializerRegistryInjector** interface. If yes, the DeserializerRegistry (which stores all deserializer references) is injected into the deserializer.

This is particularly useful when the deserializer needs access to other deserializers. For example `IntructionsDeserializer` needs access to `ActionsDeserializer` in order to be able to process OFPIT_WRITE_ACTIONS/OFPIT_APPLY_ACTIONS instructions.

Detailed walkthrough: Serialization extensibility

External interface & class description.

OFGeneralSerializer:

- OFSerializer<E extends DataObject>
  - *serialize(E,ByteBuf)* - serializes E into given ByteBuf
- `HeaderSerializer<E extends DataObject>`
  - *serializeHeaders(E,ByteBuf)* - serializes E headers (used in Multipart TableFeatures messages)

SerializerRegistryInjector * `injectSerializerRegistry(SerializerRegistry)` - injects serializer registry into serializer. Useful when custom serializer needs access to other serializers.

NOTE: SerializerRegistryInjector is not OFGeneralSerializer descendand.

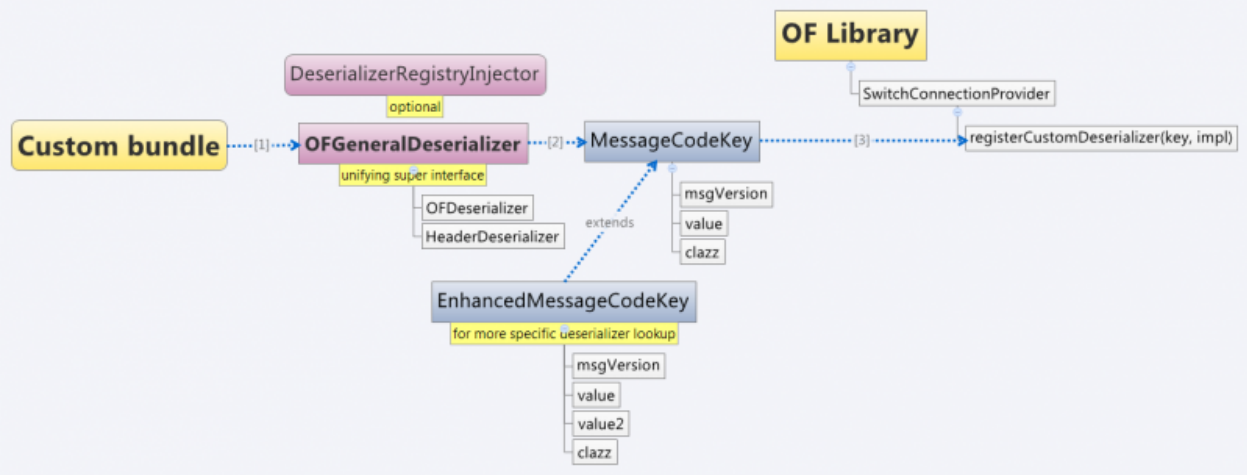

Fig. 29: Deserialization scenario walkthrough

MessageTypeKey and its descendants These keys are used as for serializer lookup in SerializerRegistry. MessageTypeKey should is used in general, while its descendants are used in more special cases. For Example ActionSerializerKey is used for Action serializer lookup and (de)registration. Vendor is provided with special keys, which contain only the most necessary fields. These keys usually start with "Experimenter" prefix (MatchEntrySerializerKey is an exception).

MessageTypeKey has these fields:

- *short version* - Openflow wire version number
- *Class<E> msgType* - DTO class

Scenario walkthrough

- [1] Serialization extensbility principles are similar to the deserialization principles. The scenario starts in a custom bundle. The custom bundle creates serializers which implement exposed OFSerializer / HeaderSerializer interfaces (wrapped under OFGeneralSerializer unifying super interface).
- [2] Created serializers are paired with their ExperimenterKeys, which are used for serializer lookup. If you don't know what key should be used with your serializer implementation, please visit *Registration keys* page.
- [3] Paired serializers are passed to the OF Library via **SwitchConnectionProvider**.*registerCustomSerializer(key, impl)*. Library registers the serializer.
- While registering, Library checks if the serializer is an instance of **SerializerRegistryInjector** interface. If yes, the SerializerRegistry (which stores all serializer references) is injected into the serializer.

This is particularly useful when the serializer needs access to other deserializers. For example IntructionsSerializer needs access to ActionsSerializer in order to be able to process OFPIT_WRITE_ACTIONS/OFPIT_APPLY_ACTIONS instructions.

Internal description

SwitchConnectionProvider `SwitchConnectionProvider` constructs and initializes both deserializer and serializer registries with default (de)serializers. It also injects the `DeserializerRegistry` into the `DeserializationFactory`, the `SerializerRegistry` into the `SerializationFactory`. When call to register custom (de)serializer is made, `SwitchConnectionProvider` calls register method on appropriate registry.

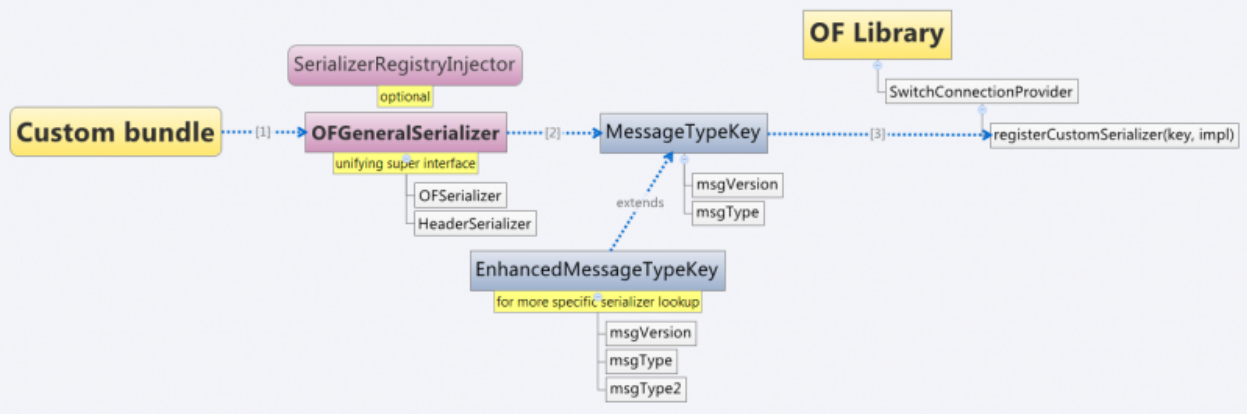

Fig. 30: Serialization scenario walkthrough

DeserializerRegistry / SerializerRegistry Both registries contain init() method to initialize default (de)serializers. Registration checks if key or (de)serializer implementation are not `null`. If at least one of the is `null`, `NullPointerException` is thrown. Else the (de)serializer implementation is checked if it is `(De)SerializerRegistryInjector` instance. If it is an instance of this interface, the registry is injected into this (de)serializer implementation.

`GetSerializer(key)` or `GetDeserializer(key)` performs registry lookup. Because there are two separate interfaces that might be put into the registry, the registry uses their unifying super interface. Get(De)Serializer(key) method casts the super interface to desired type. There is also a null check for the (de)serializer received from the registry. If the deserializer wasn't found, `NullPointerException` with key description is thrown.

Registration keys

Deserialization.

Possible openflow extensions and their keys

There are three vendor specific extensions in Openflow v1.0 and eight in Openflow v1.3. These extensions are registered under registration keys, that are shown in table below:

Extension type	Open-Flow	Registration key	Utility class
Vendor message	1.0	ExperimenterIdDeserializerKey(1, experimenterId, ExperimenterMessage.class)	ExperimenterDeserializerKeyFactory
Action	1.0	ExperimenterActionDeserializerKey(1, experimenter ID)	.
Stats message	1.0	ExperimenterMultipartReplyMessageDeserializerKey(1, experimenter ID)	ExperimenterDeserializerKeyFactory
Experimenter message	1.3	ExperimenterIdDeserializerKey(4, experimenterId, ExperimenterMessage.class)	ExperimenterDeserializerKeyFactory
Match entry	1.3	MatchEntryDeserializerKey(4, (number) ${oxm_class}, (number) ${oxm_field});	.
		key.setExperimenterId(experi menter ID);	.
Action	1.3	ExperimenterActionDeserializerKey(4, experimenter ID)	.
Instruction	1.3	ExperimenterInstructionDeserializerKey(4, experimenter ID)	.
Multipart	1.3	ExperimenterIdDeserializerKey(4, experimenterId, MultipartReplyMessage.class)	ExperimenterDeserializerKeyFactory
Multipart - Table features	1.3	ExperimenterIdDeserializerKey(4, experimenterId, TableFeatureProperties.class)	ExperimenterDeserializerKeyFactory
Error	1.3	ExperimenterIdDeserializerKey(4, experimenterId, ErrorMessage.class)	ExperimenterDeserializerKeyFactory
Queue property	1.3	ExperimenterIdDeserializerKey(4, experimenterId, QueueProperty.class)	ExperimenterDeserializerKeyFactory
Meter band type	1.3	ExperimenterIdDeserializerKey(4, experimenterId, MeterBandExperimenterCase.class)	ExperimenterDeserializerKeyFactory

Table: **Deserialization**

Serialization.

Possible openflow extensions and their keys

There are three vendor specific extensions in Openflow v1.0 and seven Openflow v1.3. These extensions are registered under registration keys, that are shown in table below:

Extension type	Open-Flow	Registration key	Utility class
Vendor message	1.0	ExperimenterIdSerializerKey< >(1, experimenterId, ExperimenterInput.class)	ExperimenterSerialize rKeyFactory
Action	1.0	ExperimenterActionSerializer Key(1, experimenterId, subtype)	.
Stats message	1.0	ExperimenterMultipartRequest SerializerKey(1, experimenter ID)	ExperimenterSerialize rKeyFactory
Experimenter message	1.3	ExperimenterIdSerializerKey< >(4, experimenterId, ExperimenterInput.class)	ExperimenterSerialize rKeyFactory
Match entry	1.3	MatchEntrySerializerKey<>(4, (class) ${oxm_class}, (class) ${oxm_field});	.
		key.setExperimenterId(experi menter ID)	.
Action	1.3	ExperimenterActionSerializer Key(4, experimenterId, subtype)	.
Instruction	1.3	ExperimenterInstructionSeria lizerKey(4, experimenter ID)	.
Multipart	1.3	ExperimenterIdSerializerKey< >(4, experimenterId, MultipartRequestExperimenter Case.class)	ExperimenterSerialize rKeyFactory
Multipart - Table features	1.3	ExperimenterIdSerializerKey< >(4, experimenterId, TableFeatureProperties.class)	ExperimenterSerialize rKeyFactory
Meter band type	1.3	ExperimenterIdSerializerKey< >(4, experimenterId, MeterBandExperimenterCase.cl ass)	ExperimenterSerialize rKeyFactory

Table: **Serialization**

OpenFlow Plugin Project Developer Guide

This section covers topics which are developer specific and which have not been covered in the user guide. Please see the OpenFlow plugin user guide first.

It can be found on the OpenDaylight software download page.

Event Sequences

Session Establishment

The *OpenFlow Protocol Library* provides interface **SwitchConnectionHandler** which contains method *onSwitchConnected* (step 1). This event is raised in the OpenFlow Protocol Library when an OpenFlow device connects to OpenDaylight and caught in the **ConnectionManagerImpl** class in the OpenFlow plugin.

There the plugin creates a new instance of the **ConnectionContextImpl** class (step 1.1) and also instances of **HandshakeManagerImpl** (which uses **HandshakeListenerImpl**) and **ConnectionReadyListenerImpl**. **ConnectionReadyListenerImpl** contains method *onConnectionReady()* which is called when connection is prepared. This method starts the handshake with the OpenFlow device (switch) from the OpenFlow plugin side. Then handshake can be also started from device side. In this case method *shake()* from **HandshakeManagerImpl** is called (steps 1.1.1 and 2).

The handshake consists of an exchange of HELLO messages in addition to an exchange of device features (steps 2.1. and 3). The handshake is completed by **HandshakeManagerImpl**. After receiving device features, the **HandshakeListenerImpl** is notifed via the *onHanshakeSuccessfull()* method. After this, the device features, node id and

connection state are stored in a **ConnectionContext** and the method *deviceConnected()* of **DeviceManagerImpl** is called.

When *deviceConnected()* is called, it does the following:

1. creates a new transaction chain (step 4.1)
2. creates a new instance of **DeviceContext** (step 4.2.2)
3. initializes the device context: the static context of device is populated by calling *createDeviceFeaturesForOF<version>()* to populate table, group, meter features and port descriptions (step 4.2.1 and 4.2.1.1)
4. creates an instance of **RequestContext** for each type of feature

When the OpenFlow device responds to these requests (step 4.2.1.1) with multipart replies (step 5) they are processed and stored to MD-SAL operational datastore. The *createDeviceFeaturesForOF<version>()* method returns a **Future** which is processed in the callback (step 5.1) (part of *initializeDeviceContext()* in the *deviceConnected()* method) by calling the method *onDeviceCtxLevelUp()* from **StatisticsManager** (step 5.1.1).

The call to *createDeviceFeaturesForOF<version>()*: . creates a new instance of **StatisticsContextImpl** (step 5.1.1.1).

1. calls *gatherDynamicStatistics()* on that instance which returns a **Future** which will produce a value when done
   (a) this method calls methods to get dynamic data (flows, tables, groups) from the device (step 5.1.1.2, 5.1.1.2.1, 5.1.1.2.1.1)
   (b) if everything works, this data is also stored in the MD-SAL operational datastore

If the **Future** is successful, it is processed (step 6.1.1) in a callback in **StatisticsManagerImpl** which:

1. schedules the next time to poll the device for statistics
2. sets the device state to synchronized (step 6.1.1.2)
3. calls *onDeviceContextLevelUp()* in **RpcManagerImpl**

The *onDeviceContextLevelUp()* call:

1. creates a new instance of **RequestContextImpl**
2. registers implementation for supported services
3. calls *onDeviceContextLevelUp()* in **DeviceManagerImpl** (step 6.1.1.2.1.2) which causes the information about the new device be be written to the MD-SAL operational datastore (step 6.1.1.2.2)

Handshake

The first thing that happens when an OpenFlow device connects to OpenDaylight is that the OpenFlow plugin gathers basic information about the device and establishes agreement on key facts like the version of OpenFlow which will be used. This process is called the handshake.

The handshake starts with HELLO message which can be sent either by the OpenFlow device or the OpenFlow plugin. After this, there are several scenarios which can happen:

1. if the first HELLO message contains a *version bitmap*, it is possible to determine if there is a common version of OpenFlow or not:
   (a) if there is a single common version use it and the **VERSION IS SETTLED**
   (b) if there are more than one common versions, use the highest (newest) protocol and the **VERSION IS SETTLED**
   (c) if there are no common versions, the device is **DISCONNECTED**
2. if the first HELLO message does not contain a *version bitmap*, then STEB-BY-STEP negotiation is used

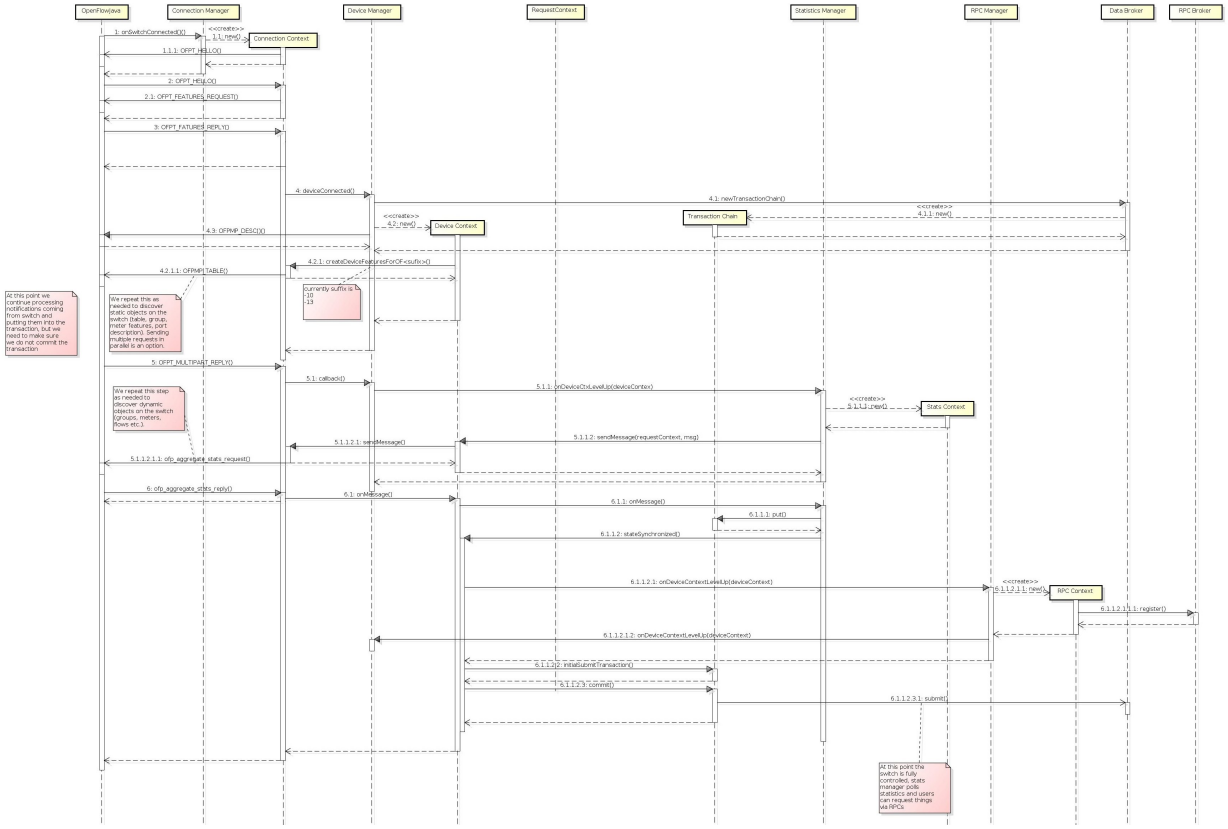

Fig. 31: Session establishment

3. if second (or more) HELLO message is received, then STEP-BY-STEP negotiation is used

STEP-BY-STEP negotiation:

- if last version proposed by the OpenFlow plugin is the same as the version received from the OpenFlow device, then the **VERSION IS SETTLED**
- if the version received in the current HELLO message from the device is the same as from previous then negotiation has failed and the device is **DISCONNECTED**
- if the last version from the device is greater than the last version proposed from the plugin, wait for the next HELLO message in the hope that it will advertise support for a lower version
- if the last version from the device is is less than the last version proposed from the plugin:
  - propose the highest version the plugin supports that is less than or equal to the version received from the device and wait for the next HELLO message
  - if if the plugin doesn't support a lower version, the device is **DISCONNECTED**

After selecting of version we can say that the **VERSION IS SETTLED** and the OpenFlow plugin can ask device for its features. At this point handshake ends.

Adding a Flow

There are two ways to add a flow in in the OpenFlow plugin: adding it to the MD-SAL config datastore or calling an RPC. Both of these can either be done using the native MD-SAL interfaces or using RESTCONF. This discussion focuses on calling the RPC.

If user send flow via REST interface (step 1) it will cause that *invokeRpc()* is called on **RpcBroker**. The **RpcBroker** then looks for an appropriate implementation of the interface. In the case of the OpenFlow plugin, this is the *addFlow()* method of **SalFlowServiceImpl** (step 1.1). The same thing happens if the RPC is called directly from the native MD-SAL interfaces.

The *addFlow()* method then

1. calls the *commitEntry()* method (step 2) from the OpenFlow Protocol Library which is responsible for sending the flow to the device
2. creates a new **RequestContext** by calling *createRequestContext()* (step 3)
3. creates a callback to handle any events that happen because of sending the flow to the device

The callback method is triggered when a barrier reply message (step 2.1) is received from the device indicating that the flow was either installed or an appropriate error message was sent. If the flow was successfully sent to the device, the RPC result is set to success (step 5). // **SalFlowService** contains inside method *addFlow()* other callback which caught notification from callback for barrier message.

At this point, no information pertaining to the flow has been added to the MD-SAL operational datastore. That is accomplished by the periodic gathering of statistics from OpenFlow devices.

The **StatisticsContext** for each given OpenFlow device periodically polls it using *gatherStatistics()* of **StatisticsGatheringUtil** which issues an OpenFlow OFPT_MULTIPART_REQUEST - OFPMP_FLOW. The response to this request (step 7) is processed in **StatisticsGatheringUtil** class where flow data is written to the MD-SAL operational datastore via the *writeToTransaction()* method of **DeviceContext**.

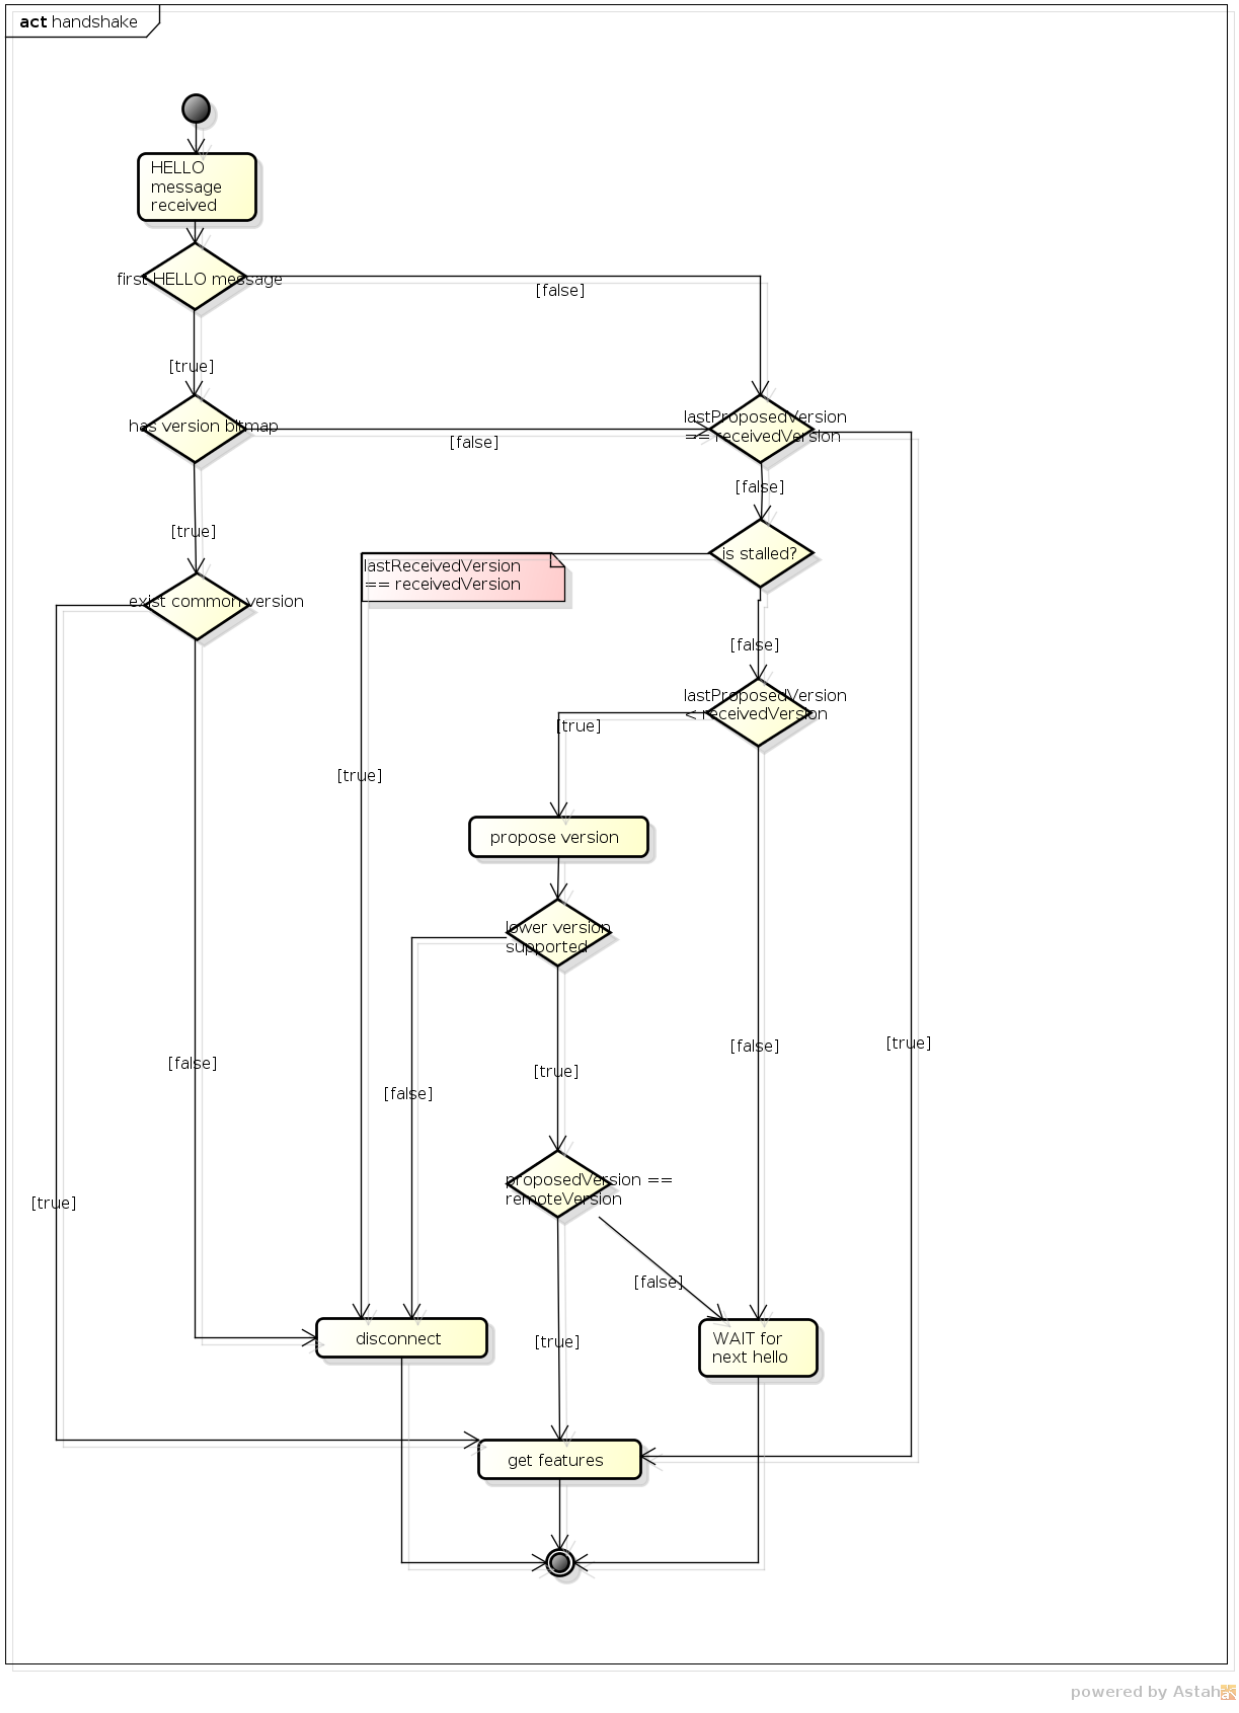

Fig. 32: Handshake process

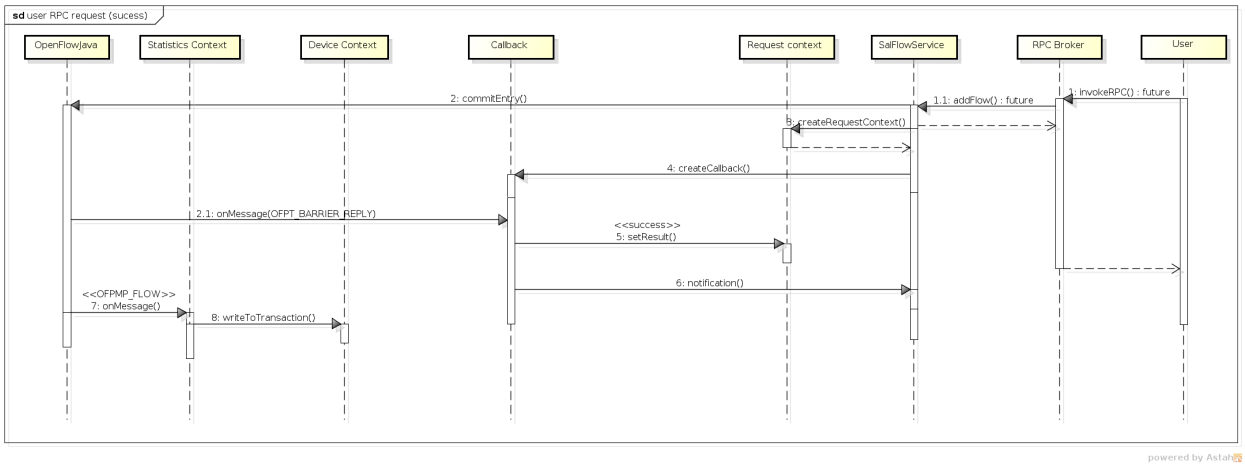

Fig. 33: Add flow

Description of OpenFlow Plugin Modules

The OpenFlow plugin project contains a variety of OpenDaylight modules, which are loaded using the configuration subsystem. This section describes the YANG files used to model each module.

General model (interfaces) - openflow-plugin-cfg.yang.

- the provided module is defined (`identity openflow-provider`)
- and target implementation is assigned (`...OpenflowPluginProvider`)

Implementation model - openflow-plugin-cfg-impl.yang

- the implementation of module is defined (`identity openflow-provider-impl`)
  - class name of generated implementation is defined (ConfigurableOpenFlowProvider)
- via augmentation the configuration of module is defined:
  - this module requires instance of binding-aware-broker (`container binding-aware-broker`)
  - and list of openflow-switch-connection-provider (those are provided by openflowjava, one plugin instance will orchestrate multiple openflowjava modules)

Generating config and sal classes out of yangs

In order to involve suitable code generators, this is needed in pom:

```
<build> ...
  <plugins>
    <plugin>
      <groupId>org.opendaylight.yangtools</groupId>
      <artifactId>yang-maven-plugin</artifactId>
      <executions>
        <execution>
          <goals>
            <goal>generate-sources</goal>
          </goals>
          <configuration>
```

(continues on next page)

(continued from previous page)

```
            <codeGenerators>
              <generator>
                <codeGeneratorClass>
                  org.opendaylight.controller.config.yangjmxgenerator.plugin.
↪JMXGenerator
                </codeGeneratorClass>
                <outputBaseDir>${project.build.directory}/generated-sources/config</
↪outputBaseDir>
                <additionalConfiguration>
                  <namespaceToPackage1>
                    urn:opendaylight:params:xml:ns:yang:controller==org.opendaylight.
↪controller.config.yang
                  </namespaceToPackage1>
                </additionalConfiguration>
              </generator>
              <generator>
                <codeGeneratorClass>
                  org.opendaylight.yangtools.maven.sal.api.gen.plugin.
↪CodeGeneratorImpl
                </codeGeneratorClass>
                <outputBaseDir>${project.build.directory}/generated-sources/sal</
↪outputBaseDir>
              </generator>
              <generator>
                <codeGeneratorClass>org.opendaylight.yangtools.yang.unified.doc.
↪generator.maven.DocumentationGeneratorImpl</codeGeneratorClass>
                <outputBaseDir>${project.build.directory}/site/models</outputBaseDir>
              </generator>
            </codeGenerators>
            <inspectDependencies>true</inspectDependencies>
          </configuration>
        </execution>
      </executions>
      <dependencies>
        <dependency>
          <groupId>org.opendaylight.controller</groupId>
          <artifactId>yang-jmx-generator-plugin</artifactId>
          <version>0.2.5-SNAPSHOT</version>
        </dependency>
        <dependency>
          <groupId>org.opendaylight.yangtools</groupId>
          <artifactId>maven-sal-api-gen-plugin</artifactId>
          <version>${yangtools.version}</version>
          <type>jar</type>
        </dependency>
      </dependencies>
    </plugin>
    ...
```

- JMX generator (target/generated-sources/config)
- sal CodeGeneratorImpl (target/generated-sources/sal)

Altering generated files

Those files were generated under src/main/java in package as referred in yangs (if exist, generator will not overwrite them):

- ConfigurableOpenFlowProviderModuleFactory

  here the **instantiateModule** methods are extended in order to capture and inject osgi BundleContext into module, so it can be injected into final implementation - **OpenflowPluginProvider** + `module.setBundleContext(bundleContext);`

- ConfigurableOpenFlowProviderModule

  here the **createInstance** method is extended in order to inject osgi BundleContext into module implementation + `pluginProvider.setContext(bundleContext);`

Configuration xml file

Configuration file contains

- required capabilities
  - modules definitions from openflowjava
  - modules definitions from openflowplugin
- modules definition
  - openflow:switch:connection:provider:impl (listening on port 6633, name=openflow-switch-connection-provider-legacy-impl)
  - openflow:switch:connection:provider:impl (listening on port 6653, name=openflow-switch-connection-provider-default-impl)
  - openflow:common:config:impl (having 2 services (wrapping those 2 previous modules) and binding-broker-osgi-registry injected)
- provided services
  - openflow-switch-connection-provider-default
  - openflow-switch-connection-provider-legacy
  - openflow-provider

```
<snapshot>
 <required-capabilities>
   <capability>
↪urn:opendaylight:params:xml:ns:yang:openflow:switch:connection:provider:impl?
↪module=openflow-switch-connection-provider-impl&revision=2014-03-28</capability>
   <capability>
↪urn:opendaylight:params:xml:ns:yang:openflow:switch:connection:provider?
↪module=openflow-switch-connection-provider&revision=2014-03-28</capability>
   <capability>urn:opendaylight:params:xml:ns:yang:openflow:common:config:impl?
↪module=openflow-provider-impl&revision=2014-03-26</capability>
   <capability>urn:opendaylight:params:xml:ns:yang:openflow:common:config?
↪module=openflow-provider&revision=2014-03-26</capability>
 </required-capabilities>

 <configuration>
```

(continues on next page)

(continued from previous page)

```
    <modules xmlns="urn:opendaylight:params:xml:ns:yang:controller:config">
      <module>
        <type xmlns:prefix=
↪"urn:opendaylight:params:xml:ns:yang:openflow:switch:connection:provider:impl">
↪prefix:openflow-switch-connection-provider-impl</type>
        <name>openflow-switch-connection-provider-default-impl</name>
        <port>6633</port>
        <switch-idle-timeout>15000</switch-idle-timeout>
      </module>
      <module>
        <type xmlns:prefix=
↪"urn:opendaylight:params:xml:ns:yang:openflow:switch:connection:provider:impl">
↪prefix:openflow-switch-connection-provider-impl</type>
        <name>openflow-switch-connection-provider-legacy-impl</name>
        <port>6653</port>
        <switch-idle-timeout>15000</switch-idle-timeout>
      </module>


      <module>
        <type xmlns:prefix=
↪"urn:opendaylight:params:xml:ns:yang:openflow:common:config:impl">prefix:openflow-
↪provider-impl</type>
        <name>openflow-provider-impl</name>

        <openflow-switch-connection-provider>
          <type xmlns:ofSwitch=
↪"urn:opendaylight:params:xml:ns:yang:openflow:switch:connection:provider">
↪ofSwitch:openflow-switch-connection-provider</type>
          <name>openflow-switch-connection-provider-default</name>
        </openflow-switch-connection-provider>
        <openflow-switch-connection-provider>
          <type xmlns:ofSwitch=
↪"urn:opendaylight:params:xml:ns:yang:openflow:switch:connection:provider">
↪ofSwitch:openflow-switch-connection-provider</type>
          <name>openflow-switch-connection-provider-legacy</name>
        </openflow-switch-connection-provider>


        <binding-aware-broker>
          <type xmlns:binding=
↪"urn:opendaylight:params:xml:ns:yang:controller:md:sal:binding">binding:binding-
↪broker-osgi-registry</type>
          <name>binding-osgi-broker</name>
        </binding-aware-broker>
      </module>
    </modules>

    <services xmlns="urn:opendaylight:params:xml:ns:yang:controller:config">
      <service>
        <type xmlns:prefix=
↪"urn:opendaylight:params:xml:ns:yang:openflow:switch:connection:provider">
↪prefix:openflow-switch-connection-provider</type>
        <instance>
          <name>openflow-switch-connection-provider-default</name>
          <provider>/modules/module[type='openflow-switch-connection-provider-impl
↪'][name='openflow-switch-connection-provider-default-impl']</provider>
```

(continues on next page)

(continued from previous page)

```
          </instance>
          <instance>
            <name>openflow-switch-connection-provider-legacy</name>
            <provider>/modules/module[type='openflow-switch-connection-provider-impl
↪'][name='openflow-switch-connection-provider-legacy-impl']</provider>
          </instance>
        </service>

        <service>
          <type xmlns:prefix=
↪"urn:opendaylight:params:xml:ns:yang:openflow:common:config">prefix:openflow-
↪provider</type>
          <instance>
            <name>openflow-provider</name>
            <provider>/modules/module[type='openflow-provider-impl'][name='openflow-
↪provider-impl']</provider>
          </instance>
        </service>
      </services>


  </configuration>
</snapshot>
```

API changes

In order to provide multiple instances of modules from openflowjava there is an API change. Previously OFPlugin got access to SwitchConnectionProvider exposed by OFJava and injected collection of configurations so that for each configuration new instance of tcp listening server was created. Now those configurations are provided by configSubsystem and configured modules (wrapping the original SwitchConnectionProvider) are injected into OFPlugin (wrapping SwitchConnectionHandler).

Providing config file (IT, local distribution/base, integration/distributions/base)

openflowplugin-it

Here the whole configuration is contained in one file (controller.xml). Required entries needed in order to startup and wire OEPlugin + OFJava are simply added there.

OFPlugin/distribution/base

Here new config file has been added (src/main/resources/configuration/initial/42-openflow-protocol-impl.xml) and is being copied to config/initial subfolder of build.

integration/distributions/build

In order to push the actual config into config/initial subfolder of distributions/base in integration project there was a new artifact in OFPlugin created - **openflowplugin-controller-config**, containing only the config xml file under src/main/resources. Another change was committed into integration project. During build this config xml is being extracted and copied to the final folder in order to be accessible during controller run.

Internal message statistics API

To aid in testing and diagnosis, the OpenFlow plugin provides information about the number and rate of different internal events.

The implementation does two things: collects event counts and exposes counts. Event counts are grouped by message type, e.g., **PacketInMessage**, and checkpoint, e.g., *TO_SWITCH_ENQUEUED_SUCCESS*. Once gathered, the results are logged as well as being exposed using OSGi command line (deprecated) and JMX.

Collect

Each message is counted as it passes through various processing checkpoints. The following checkpoints are defined as a Java enum and tracked:

```
/**
  * statistic groups overall in OFPlugin
  */
enum STATISTIC_GROUP {
     /** message from switch, enqueued for processing */
     FROM_SWITCH_ENQUEUED,
     /** message from switch translated successfully - source */
     FROM_SWITCH_TRANSLATE_IN_SUCCESS,
     /** message from switch translated successfully - target */
     FROM_SWITCH_TRANSLATE_OUT_SUCCESS,
     /** message from switch where translation failed - source */
     FROM_SWITCH_TRANSLATE_SRC_FAILURE,
     /** message from switch finally published into MD-SAL */
     FROM_SWITCH_PUBLISHED_SUCCESS,
     /** message from switch - publishing into MD-SAL failed */
     FROM_SWITCH_PUBLISHED_FAILURE,

     /** message from MD-SAL to switch via RPC enqueued */
     TO_SWITCH_ENQUEUED_SUCCESS,
     /** message from MD-SAL to switch via RPC NOT enqueued */
     TO_SWITCH_ENQUEUED_FAILED,
     /** message from MD-SAL to switch - sent to OFJava successfully */
     TO_SWITCH_SUBMITTED_SUCCESS,
     /** message from MD-SAL to switch - sent to OFJava but failed*/
     TO_SWITCH_SUBMITTED_FAILURE
}
```

When a message passes through any of those checkpoints then counter assigned to corresponding checkpoint and message is incremented by 1.

Expose statistics

As described above, there are three ways to access the statistics:

- OSGi command line (this is considered deprecated)

  ```
  osgi> dumpMsgCount
  ```

- OpenDaylight logging console (statistics are logged here every 10 seconds)

  required logback settings : `<logger name="org.opendaylight.openflowplugin.openflow.md.queue.MessageSpyCounterImpl" level="DEBUG"\/>`

- JMX (via JConsole)

  start OpenFlow plugin with the `-jmx` parameter

  start JConsole by running `jconsole`

  the JConsole MBeans tab should contain org.opendaylight.controller

  RuntimeBean has a msg-spy-service-impl

  Operations provides makeMsgStatistics report functionality

Example results

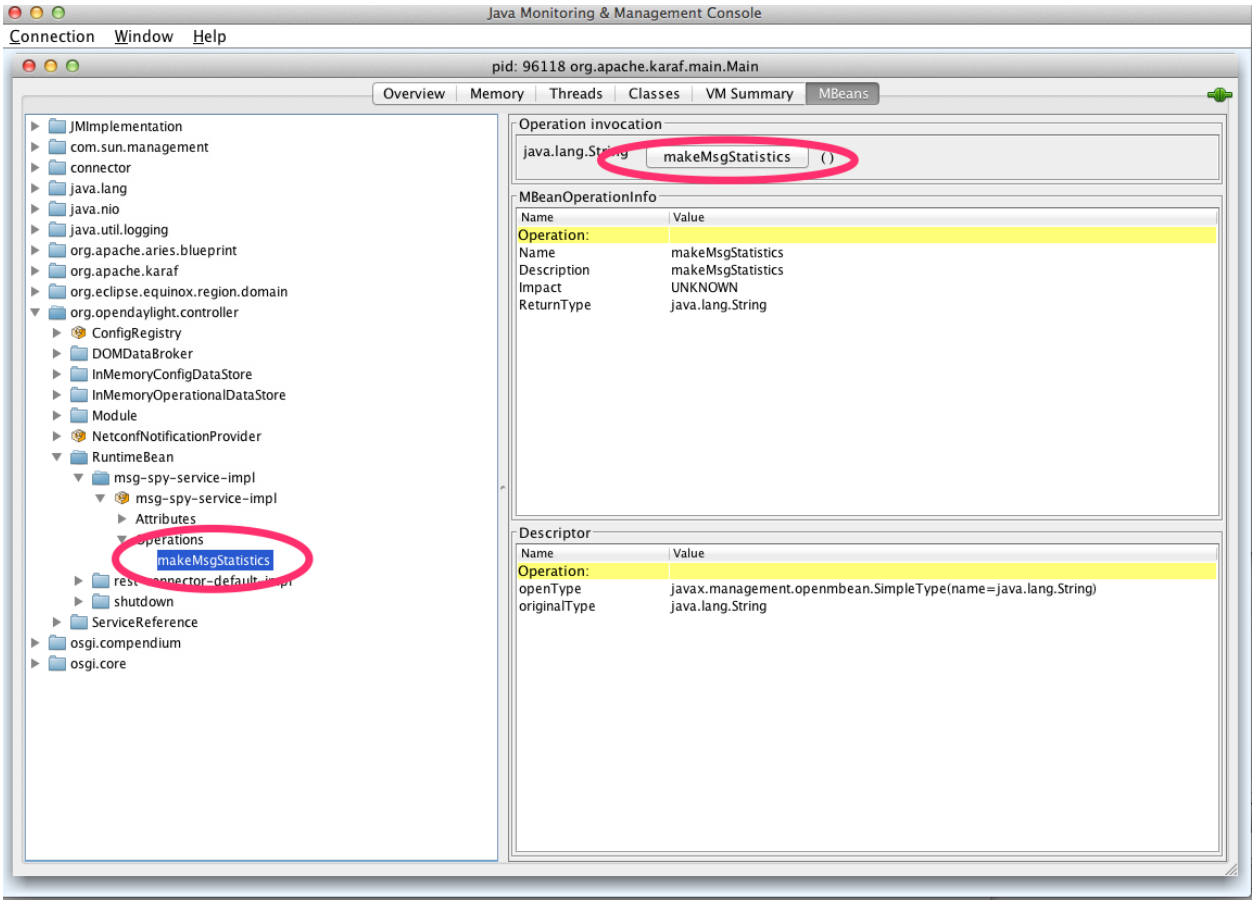

Fig. 34: OFplugin Debug stats.png

```
DEBUG o.o.o.s.MessageSpyCounterImpl - FROM_SWITCH_ENQUEUED: MSG[PortStatusMessage] ->␣
↪+0 | 1
DEBUG o.o.o.s.MessageSpyCounterImpl - FROM_SWITCH_ENQUEUED:␣
↪MSG[MultipartReplyMessage] -> +24 | 81
DEBUG o.o.o.s.MessageSpyCounterImpl - FROM_SWITCH_ENQUEUED: MSG[PacketInMessage] ->␣
↪+8 | 111
DEBUG o.o.o.s.MessageSpyCounterImpl - FROM_SWITCH_TRANSLATE_IN_SUCCESS:␣
↪MSG[PortStatusMessage] -> +0 | 1
DEBUG o.o.o.s.MessageSpyCounterImpl - FROM_SWITCH_TRANSLATE_IN_SUCCESS:␣
↪MSG[MultipartReplyMessage] -> +24 | 81
```

(continues on next page)

(continued from previous page)

```
DEBUG o.o.o.s.MessageSpyCounterImpl - FROM_SWITCH_TRANSLATE_IN_SUCCESS:␣
↪MSG[PacketInMessage] -> +8 | 111
DEBUG o.o.o.s.MessageSpyCounterImpl - FROM_SWITCH_TRANSLATE_OUT_SUCCESS:␣
↪MSG[QueueStatisticsUpdate] -> +3 | 7
DEBUG o.o.o.s.MessageSpyCounterImpl - FROM_SWITCH_TRANSLATE_OUT_SUCCESS:␣
↪MSG[NodeUpdated] -> +0 | 3
DEBUG o.o.o.s.MessageSpyCounterImpl - FROM_SWITCH_TRANSLATE_OUT_SUCCESS:␣
↪MSG[NodeConnectorStatisticsUpdate] -> +3 | 7
DEBUG o.o.o.s.MessageSpyCounterImpl - FROM_SWITCH_TRANSLATE_OUT_SUCCESS:␣
↪MSG[GroupDescStatsUpdated] -> +3 | 7
DEBUG o.o.o.s.MessageSpyCounterImpl - FROM_SWITCH_TRANSLATE_OUT_SUCCESS:␣
↪MSG[FlowsStatisticsUpdate] -> +3 | 19
DEBUG o.o.o.s.MessageSpyCounterImpl - FROM_SWITCH_TRANSLATE_OUT_SUCCESS:␣
↪MSG[PacketReceived] -> +8 | 111
DEBUG o.o.o.s.MessageSpyCounterImpl - FROM_SWITCH_TRANSLATE_OUT_SUCCESS:␣
↪MSG[MeterFeaturesUpdated] -> +0 | 3
DEBUG o.o.o.s.MessageSpyCounterImpl - FROM_SWITCH_TRANSLATE_OUT_SUCCESS:␣
↪MSG[GroupStatisticsUpdated] -> +3 | 7
DEBUG o.o.o.s.MessageSpyCounterImpl - FROM_SWITCH_TRANSLATE_OUT_SUCCESS:␣
↪MSG[GroupFeaturesUpdated] -> +0 | 3
DEBUG o.o.o.s.MessageSpyCounterImpl - FROM_SWITCH_TRANSLATE_OUT_SUCCESS:␣
↪MSG[MeterConfigStatsUpdated] -> +3 | 7
DEBUG o.o.o.s.MessageSpyCounterImpl - FROM_SWITCH_TRANSLATE_OUT_SUCCESS:␣
↪MSG[MeterStatisticsUpdated] -> +3 | 7
DEBUG o.o.o.s.MessageSpyCounterImpl - FROM_SWITCH_TRANSLATE_OUT_SUCCESS:␣
↪MSG[NodeConnectorUpdated] -> +0 | 12
DEBUG o.o.o.s.MessageSpyCounterImpl - FROM_SWITCH_TRANSLATE_OUT_SUCCESS:␣
↪MSG[FlowTableStatisticsUpdate] -> +3 | 8
DEBUG o.o.o.s.MessageSpyCounterImpl - FROM_SWITCH_TRANSLATE_SRC_FAILURE: no activity␣
↪detected
DEBUG o.o.o.s.MessageSpyCounterImpl - FROM_SWITCH_PUBLISHED_SUCCESS:␣
↪MSG[QueueStatisticsUpdate] -> +3 | 7
DEBUG o.o.o.s.MessageSpyCounterImpl - FROM_SWITCH_PUBLISHED_SUCCESS: MSG[NodeUpdated]␣
↪-> +0 | 3
DEBUG o.o.o.s.MessageSpyCounterImpl - FROM_SWITCH_PUBLISHED_SUCCESS:␣
↪MSG[NodeConnectorStatisticsUpdate] -> +3 | 7
DEBUG o.o.o.s.MessageSpyCounterImpl - FROM_SWITCH_PUBLISHED_SUCCESS:␣
↪MSG[GroupDescStatsUpdated] -> +3 | 7
DEBUG o.o.o.s.MessageSpyCounterImpl - FROM_SWITCH_PUBLISHED_SUCCESS:␣
↪MSG[FlowsStatisticsUpdate] -> +3 | 19
DEBUG o.o.o.s.MessageSpyCounterImpl - FROM_SWITCH_PUBLISHED_SUCCESS:␣
↪MSG[PacketReceived] -> +8 | 111
DEBUG o.o.o.s.MessageSpyCounterImpl - FROM_SWITCH_PUBLISHED_SUCCESS:␣
↪MSG[MeterFeaturesUpdated] -> +0 | 3
DEBUG o.o.o.s.MessageSpyCounterImpl - FROM_SWITCH_PUBLISHED_SUCCESS:␣
↪MSG[GroupStatisticsUpdated] -> +3 | 7
DEBUG o.o.o.s.MessageSpyCounterImpl - FROM_SWITCH_PUBLISHED_SUCCESS:␣
↪MSG[GroupFeaturesUpdated] -> +0 | 3
DEBUG o.o.o.s.MessageSpyCounterImpl - FROM_SWITCH_PUBLISHED_SUCCESS:␣
↪MSG[MeterConfigStatsUpdated] -> +3 | 7
DEBUG o.o.o.s.MessageSpyCounterImpl - FROM_SWITCH_PUBLISHED_SUCCESS:␣
↪MSG[MeterStatisticsUpdated] -> +3 | 7
DEBUG o.o.o.s.MessageSpyCounterImpl - FROM_SWITCH_PUBLISHED_SUCCESS:␣
↪MSG[NodeConnectorUpdated] -> +0 | 12
DEBUG o.o.o.s.MessageSpyCounterImpl - FROM_SWITCH_PUBLISHED_SUCCESS:␣
↪MSG[FlowTableStatisticsUpdate] -> +3 | 8
DEBUG o.o.o.s.MessageSpyCounterImpl - FROM_SWITCH_PUBLISHED_FAILURE: no activity␣
↪detected
```

(continues on next page)

(continued from previous page)

```
DEBUG o.o.o.s.MessageSpyCounterImpl - TO_SWITCH_ENQUEUED_SUCCESS: MSG[AddFlowInput] ->
↪ +0 | 12
DEBUG o.o.o.s.MessageSpyCounterImpl - TO_SWITCH_ENQUEUED_FAILED: no activity detected
DEBUG o.o.o.s.MessageSpyCounterImpl - TO_SWITCH_SUBMITTED_SUCCESS: MSG[AddFlowInput] -
↪> +0 | 12
DEBUG o.o.o.s.MessageSpyCounterImpl - TO_SWITCH_SUBMITTED_FAILURE: no activity␣
↪detected
```

Application: Forwarding Rules Synchronizer

Basics

Description

Forwarding Rules Synchronizer (FRS) is a newer version of Forwarding Rules Manager (FRM). It was created to solve most shortcomings of FRM. FRS solving errors with retry mechanism. Sending barrier if needed. Using one service for flows, groups and meters. And it has less changes requests send to device since calculating difference and using compression queue.

It is located in the Java package:

```
package org.opendaylight.openflowplugin.applications.frsync;
```

Listeners

- 1x config - FlowCapableNode
- 1x operational - Node

System of work

- one listener in config datastore waiting for changes
  - update cache
  - skip event if operational not present for node
  - send syncup entry to reactor for synchronization
    - node added: after part of modification and whole operational snapshot
    - node updated: after and before part of modification
    - node deleted: null and before part of modification
- one listener in operational datastore waiting for changes
  - update cache
  - on device connected
    - register for cluster services
  - on device disconnected remove from cache
    - remove from cache

        * unregister for cluster services
    - if registered for reconciliation
        * do reconciliation through syncup (only when config present)
- reactor *(provides syncup w/decorators assembled in this order)*
    - Cluster decorator - skip action if not master for device
    - FutureZip decorator (FutureZip extends Future decorator)
        * Future - run delegate syncup in future - submit task to executor service
        * FutureZip - provides state compression - compress optimized config delta if waiting for execution with new one
    - Guard decorator - per device level locking
    - Retry decorator - register for reconciliation if syncup failed
    - Reactor impl - calculate diff from after/before parts of syncup entry and execute

Strategy

In the *old* FRM uses an incremental strategy with all changes made one by one, where FRS uses a flat batch system with changes made in bulk. It uses one service SalFlatBatchService instead of three (flow, group, meter).

Boron release

FRS is used in Boron as separate feature and it is not loaded by any other feature. It has to be run separately.

```
odl-openflowplugin-app-forwardingrules-sync
```

FRS additions

Retry mechanism

- is started when change request to device return as failed (register for reconcile)
- wait for next consistent operational and do reconciliation with actual config (not only diff)

ZipQueue

- only the diff (before/after) between last config changes is sent to device
- when there are more config changes for device in a row waiting to be processed they are compressed into one entry (after is still replaced with the latest)

Cluster-aware

- FRS is cluster aware using ClusteringSingletonServiceProvider from the MD-SAL
- on mastership change reconciliation is done (register for reconcile)

SalFlatBatchService

FRS uses service with implemented barrier waiting logic between dependent objects

Service: SalFlatBatchService

Basics

SalFlatBatchService was created along forwardingrules-sync application as the service that should application used by default. This service uses only one input with bag of flow/group/meter objects and their common add/update/remove action. So you practically send only one input (of specific bags) to this service.

- interface: *org.opendaylight.yang.gen.v1.urn.opendaylight.flat.batch.service.rev160321.SalFlatBatchService*
- implementation: *org.opendaylight.openflowplugin.impl.services.SalFlatBatchServiceImpl*
- method: *processFlatBatch(input)*
- input: *org.opendaylight.yang.gen.v1.urn.opendaylight.flat.batch.service.rev160321.ProcessFlatBatchInput*

Usage benefits

- possibility to use only one input bag with particular failure analysis preserved
- automatic barrier decision (chain+wait)
- less RPC routing in cluster environment (since one call encapsulates all others)

ProcessFlatBatchInput

Input for SalFlatBatchService (ProcessFlatBatchInput object) consists of:

- node - NodeRef
- batch steps - List<Batch> - defined action + bag of objects + order for failures analysis
  - BatchChoice - yang-modeled action choice (e.g. FlatBatchAddFlowCase) containing batch bag of objects (e.g. flows to be added)
  - BatchOrder - (integer) order of batch step (should be incremented by single action)
- exitOnFirstError - boolean flag

Workflow

1. prepare **list of steps** based on input
2. **mark barriers** in steps where needed
3. prepare particular **F/G/M-batch** service calls from **Flat-batch** steps
   - F/G/M-batch services encapsulate bulk of single service calls
   - they actually chain barrier after processing all single calls if actual step is marked as barrier-needed
4. **chain** futures and **start** executing
   - start all actions that can be run simultaneously (chain all on one starting point)

- in case there is a step marked as barrier-needed
  - wait for all fired jobs up to one with barrier
  - merge rpc results (status, errors, batch failures) into single one
  - the latest job with barrier is new starting point for chaining

Services encapsulation

- SalFlatBatchService
  - SalFlowBatchService
    - SalFlowService
  - SalGroupBatchService
    - SalGroupService
  - SalMeterBatchService
    - SalMeterService

Barrier decision

- decide on actual step and all previous steps since the latest barrier
- if condition in table below is satisfied the latest step before actual is marked as barrier-needed

actual step	previous steps contain
FLOW_ADD *or* FLOW_UPDATE	GROUP_ADD *or* METER_ADD
GROUP_ADD	GROUP_ADD *or* GROUP_UPDATE
GROUP_REMOVE	FLOW_UPDATE *or* FLOW_REMOVE *or* GROUP_UPDATE *or* GROUP_REMOVE
METER_REMOVE	FLOW_UPDATE *or* FLOW_REMOVE

Error handling

There is flag in ProcessFlatBatchInput to stop process on the first error.

- *true* - if partial step is not successful stop whole processing
- *false* (default) - try to process all steps regardless partial results

If error occurs in any of partial steps upper FlatBatchService call will return as unsuccessful in both cases. However every partial error is attached to general flat batch result along with BatchFailure (contains BatchOrder and BatchItemIdChoice to identify failed step).

Cluster singleton approach in plugin

Basics

Description

The existing OpenDaylight service deployment model assumes symmetric clusters, where all services are activated on all nodes in the cluster. However, many services require that there is a single active service instance per cluster. We call such services *singleton services*. The Entity Ownership Service (EOS) represents the base Leadership choice for one Entity instance. Every Cluster Singleton service **type** must have its own Entity and every Cluster Singleton service **instance** must have its own Entity Candidate. Every registered Entity Candidate should be notified about its actual role. All this "work" is done by MD-SAL so the Openflowplugin need "only" to register as service in **SingletonClusteringServiceProvider** given by MD-SAL.

Change against using EOS service listener

In this new clustering singleton approach plugin uses API from the MD-SAL project: SingletonClusteringService which comes with three methods.

```
instantiateServiceInstance()
closeServiceInstance()
getIdentifier()
```

This service has to be registered to a SingletonClusteringServiceProvider from MD-SAL which take care if mastership is changed in cluster environment.

First method in SingletonClusteringService is being called when the cluster node becomes a MASTER. Second is being called when status changes to SLAVE or device is disconnected from cluster. Last method plugins returns NodeId as ServiceGroupIdentifier Startup after device is connected

On the start up the plugin we need to initialize first four managers for each working area providing information and services

- Device manager
- RPC manager
- Role manager
- Statistics manager

After connection the device the listener Device manager get the event and start up to creating the context for this connection. Startup after device connection

Services are managed by SinlgetonClusteringServiceProvider from MD-SAL project. So in startup we simply create a instance of LifecycleService and register all contexts into it.

Role change

Plugin is no longer registered as Entity Ownership Service (EOS) listener therefore does not need to and cannot respond on EOS ownership changes.

Service start

Services start asynchronously but the start is managed by LifecycleService. If something goes wrong LifecycleService stop starting services in context and this speeds up the reconnect process. But the services haven't changed and plugin need to start all this:

- Activating transaction chain manager
- Initial gathering of device statistics
- Initial submit to DS
- Sending role MASTER to device
- RPC services registration
- Statistics gathering start

Service stop

If closeServiceInstance occurred plugin just simply try to store all unsubmitted transactions and close the transaction chain manager, stop RPC services, stop Statistics gathering and after that all unregister txEntity from EOS.

Yang models and API

Model
Openflow basic types
opendaylight-table-types.yang
opendaylight-action-types.yang
opendaylight-flow-types.yang
opendaylight-meter-types.yang
opendaylight-group-types.yang
opendaylight-match-types.yang
opendaylight-port-types.yang
opendaylight-queue-types.yang
Openflow services
sal-table.yang
sal-group.yang
sal-queue.yang
flow-errors.yang
flow-capable-transaction.yang
sal-flow.yang
sal-meter.yang
flow-topology-discovery.yang
node-errors.yang
node-config.yang
sal-echo.yang
sal-port.yang
packet-processing.yang
flow-node-inventory.yang
Openflow statistics
opendaylight-queue-statistics.yang

Continued on next page

Table 1 – continued from previous page

Model
opendaylight-flow-table-statistics.yang
opendaylight-port-statistics.yang
opendaylight-statistics-types.yang
opendaylight-group-statistics.yang
opendaylight-flow-statistics.yang
opendaylight-meter-statistics.yang

Karaf feature tree

Fig. 35: Openflow plugin karaf feature tree

Short HOWTO create such a tree.

Wiring up notifications

Introduction

We need to translate OpenFlow messages coming up from the *OpenFlow Protocol Library* into MD-SAL Notification objects and then publish them to the MD-SAL.

Mechanics

1. Create a Translator class
2. Register the Translator
3. Register the notificationPopListener to handle your Notification Objects

Create a Translator class

You can see an example in PacketInTranslator.java.

First, simply create the class

```
public class PacketInTranslator implements IMDMessageTranslator<OfHeader, List
↪<DataObject>> {
```

Then implement the translate function:

```
public class PacketInTranslator implements IMDMessageTranslator<OfHeader, List
↪<DataObject>> {

    protected static final Logger LOG = LoggerFactory
            .getLogger(PacketInTranslator.class);
    @Override
    public PacketReceived translate(SwitchConnectionDistinguisher cookie,
            SessionContext sc, OfHeader msg) {
            ...
    }
```

Make sure to check that you are dealing with the expected type and cast it:

```
if(msg instanceof PacketInMessage) {
    PacketInMessage message = (PacketInMessage)msg;
    List<DataObject> list = new CopyOnWriteArrayList<DataObject>();
```

Do your transation work and return

```
PacketReceived pktInEvent = pktInBuilder.build();
list.add(pktInEvent);
return list;
```

Register your Translator Class

Next you need to go to MDController.java and in init() add register your Translator:

```
public void init() {
        LOG.debug("Initializing!");
        messageTranslators = new ConcurrentHashMap<>();
        popListeners = new ConcurrentHashMap<>();
        //TODO: move registration to factory
        addMessageTranslator(ErrorMessage.class, OF10, new ErrorTranslator());
        addMessageTranslator(ErrorMessage.class, OF13, new ErrorTranslator());
        addMessageTranslator(PacketInMessage.class,OF10, new PacketInTranslator());
        addMessageTranslator(PacketInMessage.class,OF13, new PacketInTranslator());
```

Notice that there is a separate registration for each of OpenFlow 1.0 and OpenFlow 1.3. Basically, you indicate the type of OpenFlow Protocol Library message you wish to translate for, the OpenFlow version, and an instance of your Translator.

Register your MD-SAL Message for Notification to the MD-SAL

Now, also in MDController.init() register to have the notificationPopListener handle your MD-SAL Message:

```
addMessagePopListener(PacketReceived.class, new NotificationPopListener<DataObject>
↪());
```

You are done

That's all there is to it. Now when a message comes up from the OpenFlow Protocol Library, it will be translated and published to the MD-SAL.

Message Order Preservation

While the Helium release of OpenFlow Plugin relied on queues to ensure messages were delivered in order, subsequent releases instead ensure that all the messages from a given device are delivered using the same thread and thus message order is guaranteed without queues. The OpenFlow plugin allocates a number of threads equal to twice the number of processor cores on machine it is run, e.g., 8 threads if the machine has 4 cores.

Note: While each device is assigned to one thread, multiple devices can be assigned to the same thread.

OpFlex agent-ovs Developer Guide

Overview

agent-ovs is a policy agent that works with OVS to enforce a group-based policy networking model with locally attached virtual machines or containers. The policy agent is designed to work well with orchestration tools like OpenStack.

agent-ovs Architecture

agent-ovs uses libopflex to communicate with an OpFlex-based policy repository to enforce policy on network endpoints attached to OVS by an orchestration system.

The key components are:

- Agent - coordinates startup and configuration
- Renderers - Renderers are responsible for rendering policy. This is a very general mechanism but the currently-implemented renderer is the stitched-mode renderer that can work along with with hardware fabrics such as ACI that support policy enforcement.
- EndpointManager - Keep track of network endpoints and declare them to the endpoint repository
- PolicyManager - Keep track of and index policies
- FlowManager - render policies to OVS

API Reference Documentation

Internal API documentation can be found here: https://jenkins.opendaylight.org/opflex/job/opflex-merge/ws/agent-ovs/doc/html/index.html

OpFlex genie Developer Guide

Overview

Genie is a tool for code generation from a model. It supports generating C++ and Java code. C++ can be generated suitable for use with libopflex. C++ and Java can be generated as a plain set of objects.

Group-based Policy Model

The group-based policy model is included with the genie tool and can be found under the MODEL directory. By running mvn exec:java, libmodelgbp will be generated as a library project that, when built and installed, will work with libopflex. This model is used by the OVS agent.

API Reference Documentation

Complete API documentation for the generated libmodelgbp can be found here: https://jenkins.opendaylight.org/opflex/job/opflex-merge/ws/libopflex/doc/html/index.html

OpFlex libopflex Developer Guide

Overview

The OpFlex framework allows you to develop agents that can communicate using the OpFlex protocol and act as a policy element in an OpFlex-based distributed control system. The OpFlex architecture provides a distributed control system based on a declarative policy information model. The policies are defined at a logically centralized policy repository and enforced within a set of distributed policy elements. The policy repository communicates with the subordinate policy elements using the OpFlex control protocol. This protocol allows for bidirectional communication of policy, events, statistics, and faults.

Rather than simply providing access to the OpFlex protocol, this framework allows you to directly manipulate a management information tree containing a hierarchy of managed objects. This tree is kept in sync as needed with other policy elements in the system, and you are automatically notified when important changes to the model occur. Additionally, we can ensure that only those managed objects that are important to the local policy element are synchronized locally.

Object Model

Interactions with the OpFlex framework happen through the management information tree. This is a tree of managed objects defined by an object model specific to your application. There are a few important major category of objects that can appear in the model.

- First, there is the policy object. A policy object represents some data related to a policy that describes a user intent for how the system should behave. A policy object is stored in the policy repository which is the source of "truth" for this object.
- Second, there is an endpoint object. A endpoint represents an entity in the system to which we want to apply policy, which could be a network interface, a storage array, or other relevent policy endpoint. Endpoints are discovered and reported by policy elements locally, and are synchronized into the endpoint repository. The originating policy element is the source of truth for the endpoints it discovers. Policy elements can retrieve information about endpoints discovered by other policy elements by resolving endpoints from the endpoint repository.
- Third, there is the observable object. An observable object represents some state related to the operational status or health of the policy element. Observable objects will be reported to the observer.
- Finally, there is the local-only object. This is the simplest object because it exists only local to a particular policy element. These objects can be used to store state specific to that policy element, or as helpers to resolve other objects. Read on to learn more.

You can use the genie tool that is included with the framework to produce your application model along with a set of generated accessor classes that can work with this framework library. You should refer to the documentation that accompanies the genie tool for information on how to use to to generate your object model. Later in this guide, we'll go through examples of how to use the generated managed object accessor classes.

Programming by Side Effect

When developing software on the OpFlex framework, you'll need to think in a slightly different way. Rather than calling an API function that would perform some specific action, you'll need to write a managed object to the managed object database. Writing that object to the store will trigger the side effect of performing the action that you want.

For example, a policy element will need to have a component responsible for discovering policy endpoints. When it discovers a policy endpoint, it would write an endpoint object into the managed object database. That endpoint object will contain a reference to policy that is relevant to the endpoint object. This will trigger a whole cascade of events. First, the framework will notice that an endpoint object has been created and it will write it to the endpoint repository. Second, the framework to will attempt to resolve the unresolved reference to the relevent policy object. There might be a whole chain of policy resolutions that will be automatically performed to download all the relevent policy until there are no longer any dangling references.

As long as there is a locally-created object in the system with a reference to that policy, the framework will continually ensure that the policy and any transitive policies are kept up to date. The policy element can subscribe to updates to these policy classes that will be invoked either the first time the policy is resolved or any time the policy changes.

A consequence of this design is that the managed object database can be temporarily in an inconsistent state with unresolved dangling references. Eventually, however, the inconsistency will be fully resolved. The policy element must be able to cleanly handle partially-resolved or inconsistent state and eventually reach the correct state as it receives these update notifications. Note that, in the OpFlex architecture, when there is no policy that specifically allows a particular action, that action must be prevented.

Let's cover one slightly more complex example. If a policy element needs to discover information about an endpoint that is not local to that policy element, it will need to retrieve that information from the endpoint repository. However, just as there is no API call to retrieve a policy object from the policy repository, there is no API call to retrieve an endpoint from the endpoint repository.

To get this information, the policy element needs to create a local-only object that references the endpoint. Once it creates this local-only object, if the endpoint is not already resolved, the framework will notice the dangling reference and automatically resolve the endpoint from the endpoint respository. When the endpoint resolution completes, the framework deliver an update notification to the policy element. The policy element will continue to receive any updates related to that endpoint until the policy element remove the local-only reference to the object. Once this occurs, the framework can garbage-collect any unreferenced objects.

Threading and Ownership

The OpFlex framework uses a somewhat unique threading model. Each managed object in the system belongs to a particular owner. An owner would typically be a single thread that is reponsible for all updates to objects with that owner. Though anything can read the state of a managed object, only the owner of a managed object is permitted to write to it. Though this is not strictly required for correctness, the performance of the system wil be best if you ensure that only one thread at a time is writing to objects with a particular owner.

Change notifications are delivered in a serialized fashion by a single thread. Blocking this thread from a notification callback will stall delivery of all notifications. It is therefore best practice to ensure that you do not block or perform long-running operations from a notification callback.

Key APIs and Interfaces

Basic Usage and Initialization

The primary interface point into the framework is opflex::ofcore::OFFramework. You can choose to instantiate your own copy of the framework, or you can use the static default instance.

Before you can use the framework, you must initialize it by installing your model metadata. The model metadata is accessible through the generated model library. In this case, it assumes your model is called "mymodel":

```
#include <opflex/ofcore/OFFramework.h>
#include <mymodel/metadata/metadata.hpp>
// ...
using opflex::ofcore::OFFramework;
OFFramework::defaultInstance().setModel(mymodel::getMetadata());
```

The other critical piece of information required for initialization is the OpFlex identity information. The identity information is required in order to successfully connect to OpFlex peers. In OpFlex, each component has a unique name within its policy domain, and each policy domain is identified by a globally unique domain name. You can set this identity information by calling:

```
OFFramework::defaultInstance()
    .setOpflexIdentity("[component name]", "[unique domain]");
```

You can then start the framework simply by calling:

```
OFFramework::defaultInstance().start();
```

Finally, you can add peers after the framework is started by calling the `opflex::ofcore::OFFramework::addPeer` method:

```
OFFramework::defaultInstance().addPeer("192.168.1.5", 1234);
```

When connecting to the peer, that peer may provide an additional list of peers to connect to, which will be automatically added as peers. If the peer does not include itself in the list, then the framework will disconnect from that peer and add the peers in the list. In this way, it is possible to automatically bootstrap the correct set of peers using a known hostname or IP address or a known, fixed anycast IP address.

To cleanly shut down, you can call:

```
OFFramework::defaultInstance().stop();
```

Working with Data in the Tree

Reading from the Tree

You can access data in the managed tree using the generated accessor classes. The details of exactly which classes you'll use will depend on the model you're using, but let's assume that we have a simple model called "simple" with the following classes:

- root - The root node. The URI for the root node is "/"
- foo - A policy object, and a child of root, with a scalar string property called "bar", and a unsigned 64-bit integer property called baz. The bar property is the naming property for foo. The URI for a foo object would be "/foo/[value of bar]/"

- fooref - A local-only child of root, with a reference to a foo, and a scalar string property called "bar". The bar property is the naming property for foo. The URI for a fooref object would be "/fooref/[value of bar]/"

In this example, we'll have a generated class for each of the objects. There are two main ways to get access to an object in the tree.

First, we can get instantiate an accessor class to any node in the tree by calling one of its static resolve functions. The resolve functions can take either an already-built URI that represents the object, or you can call the version that will locate the object by its naming properties.

Second, we can access the object also from its parent object using the appropriate child resolver member functions.

However we read it, the object we get back is an immutable view into the object it references. The properties set locally on that object will not change even though the underlying object may have been updated in the store. Note, however, that its children can change between when you first retrieve the object and when you resolve any children.

Another thing that is critical to note again is that when you attempt to resolve an object, you can get back nothing, even if the object actually does exist on another OpFlex node. You must ensure that some object in the managed object database references the remote managed object you want before it will be visible to you.

To get access to the root node using the default framework instance, we can simply call:

```
using boost::shared_ptr;
using boost::optional;
optional<shared_ptr<simple::root> > r(simple::root::resolve());
```

Note that whenever we can a resolve function, we get back our data in the form of an optional shared pointer to the object instance. If the node does not exist, then the optional will be set to boost::none. Note that if you dereference an optional that has not been set, you'll trigger an assert, so you must check the return as follows:

```
if (!r) {
   // handle missing object
}
```

Now let's get a child node of the root in three different ways:

```
// Get foo1 by constructing its URI from the root
optional<shared_ptr<simple::foo> > foo1(simple::foo::resolve("test"));
// get foo1 by constructing its URI relative to its parent
foo1 = r.get()->resolveFoo("test");
// get foo1 by manually building its URI
foo1 = simple::foo::resolve(opflex::modb::URIBuilder()
                            .addElement("foo")
                            .addElement("test")
                            .build());
```

All three of these calls will give us the same object, which is the "foo" object located at "/foo/test/".

The foo class has a single string property called "bar". We can easily access it as follows:

```
const std::string& barv = foo1.getBar();
```

Writing to the Tree

Writing to the tree is nearly as easy as reading from it. The key concept to understand is the mutator object. If you want to make changes to the tree, you must allocate a mutator object. The mutator will register itself in some thread-local storage in the framework instance you're using. The mutator is specific to a single "owner" for the data, so you can only make changes to data associated with that owner.

Whenever you modify one of the accessor classes, the change is actually forwarded to the currently-active mutator. You won't see any of the changes you make until you call the commit member function on the mutator. When you do that, all the changes you made are written into the store.

Once the changes are written into the store, you will need to call the appropriate resolve function again to see the changes.

Allocating a mutator is simple. To create a mutator for the default framework instance associated with the owner "owner1", just allocate the mutator on the stack. Be sure to call commit() before it goes out of scope or you'll lose your changes.

```
{
    opflex::modb::Mutator mutator("owner1");
    // make changes here
    mutator.commit();
}
```

Note that if an exception is thrown while making changes but before committing, the mutator will go out of scope and the changes will be discarded.

To create a new node, you must call the appropriate add[Child] member function on its parent. This function takes parameters for each of the naming properties for the object:

```
shared_ptr<simple::foo> newfoo(root->addFoo("test"));
```

This will return a shared pointer to a new foo object that has been registered in the active mutator but not yet committed. The "bar" naming property will be set automatically, but if you want to set the "baz" property now, you can do so by calling:

```
newfoo->setBaz(42);
```

Note that creating the root node requires a call to the special static class method createRootElement:

```
shared_ptr<simple::root> newroot(simple::root::createRootElement());
```

Here's a complete example that ties this all together:

```
{
    opflex::modb::Mutator mutator("owner1");
    shared_ptr<simple::root> newroot(simple::root::createRootElement());
    shared_ptr<simple::root> newfoo(newroot->addFoo("test"));
    newfoo->setBaz(42);

    mutator.commit();
}
```

Update Notifications

When using the OpFlex framework, you're likely to find that most of your time is spend responding to changes in the managed object database. To get these notifications, you're going to need to register some number of listeners.

You can register an object listener to see all changes related to a particular class by calling a static function for that class. You'll then get notifications whenever any object in that class is added, updated, or deleted. The listener should queue a task to read the new state and perform appropriate processing. If this function blocks or peforms a long-running operation, then the dispatching of update notifications will be stalled, but there will not be any other deleterious effects.

If multiple changes happen to the same URI, then at least one notification will be delivered but some events may be consolidated.

The update you get will tell you the URI and the Class ID of the changed object. The class ID is a unique ID for each class. When you get the update, you'll need to call the appropriate resolve function to retrieve the new value.

You'll need to create your own object listener derived from opflex::modb::ObjectListener:

```
class MyListener : public ObjectListener {
public:
    MyListener() { }
    virtual void objectUpdated(class_id_t class_id, const URI& uri) {
        // Your handler here
    }
};
```

To register your listener with the default framework instance, just call the appropriate class static method:

```
MyListener listener;
simple::foo::registerListener(&listener);
// main loop
simple::foo::unregisterListener(&listener);
```

The listener will now recieve notifications whenever any foo or any children of any foo object changes.

Note that you must ensure that you unregister your listeners before deallocating them.

API Reference Documentation

Complete API documentation can be found through doxygen here: https://jenkins.opendaylight.org/opflex/job/opflex-merge/ws/libopflex/doc/html/index.html

OVSDB Developer Guide

OVSDB Integration

The Open vSwitch database (OVSDB) Southbound Plugin component for OpenDaylight implements the OVSDB RFC 7047 management protocol that allows the southbound configuration of switches that support OVSDB. The component comprises a library and a plugin. The OVSDB protocol uses JSON-RPC calls to manipulate a physical or virtual switch that supports OVSDB. Many vendors support OVSDB on various hardware platforms. The OpenDaylight controller uses the library project to interact with an OVS instance.

Note: Read the OVSDB User Guide before you begin development.

OpenDaylight OVSDB southbound plugin architecture and design

OpenVSwitch (OVS) is generally accepted as the unofficial standard for Virtual Switching in the Open hypervisor based solutions. Every other Virtual Switch implementation, properietery or otherwise, uses OVS in some form. For information on OVS, see Open vSwitch.

In Software Defined Networking (SDN), controllers and applications interact using two channels: OpenFlow and OVSDB. OpenFlow addresses the forwarding-side of the OVS functionality. OVSDB, on the other hand, addresses

the management-plane. A simple and concise overview of Open Virtual Switch Database(OVSDB) is available at: http://networkstatic.net/getting-started-ovsdb/

Overview of OpenDaylight Controller architecture

The OpenDaylight controller platform is designed as a highly modular and plugin based middleware that serves various network applications in a variety of use-cases. The modularity is achieved through the Java OSGi framework. The controller consists of many Java OSGi bundles that work together to provide the required controller functionalities.

The bundles can be placed in the following broad categories:

- Network Service Functional Modules (Examples: Topology Manager, Inventory Manager, Forwarding Rules Manager,and others)
- NorthBound API Modules (Examples: Topology APIs, Bridge Domain APIs, Neutron APIs, Connection Manager APIs, and others)
- Service Abstraction Layer(SAL)- (Inventory Services, DataPath Services, Topology Services, Network Config, and others)
- SouthBound Plugins (OpenFlow Plugin, OVSDB Plugin, OpenDove Plugin, and others)
- Application Modules (Simple Forwarding, Load Balancer)

Each layer of the Controller architecture performs specified tasks, and hence aids in modularity. While the Northbound API layer addresses all the REST-Based application needs, the SAL layer takes care of abstracting the SouthBound plugin protocol specifics from the Network Service functions.

Each of the SouthBound Plugins serves a different purpose, with some overlapping. For example, the OpenFlow plugin might serve the Data-Plane needs of an OVS element, while the OVSDB plugin can serve the management plane needs of the same OVS element. As the OpenFlow Plugin talks OpenFlow protocol with the OVS element, the OVSDB plugin will use OVSDB schema over JSON-RPC transport.

OVSDB southbound plugin

The Open vSwitch Database Management Protocol-draft-02 and Open vSwitch Manual provide theoretical information about OVSDB. The OVSDB protocol draft is generic enough to lay the groundwork on Wire Protocol and Database Operations, and the OVS Manual currently covers 13 tables leaving space for future OVS expansion, and vendor expansions on proprietary implementations. The OVSDB Protocol is a database records transport protocol using JSON RPC1.0. For information on the protocol structure, see Getting Started with OVSDB. The OpenDaylight OVSDB southbound plugin consists of one or more OSGi bundles addressing the following services or functionalities:

- Connection Service - Based on Netty
- Network Configuration Service
- Bidirectional JSON-RPC Library
- OVSDB Schema definitions and Object mappers
- Overlay Tunnel management
- OVSDB to OpenFlow plugin mapping service
- Inventory Service

Connection service

One of the primary services that most southbound plugins provide in OpenDaylight a Connection Service. The service provides protocol specific connectivity to network elements, and supports the connectivity management services as specified by the OpenDaylight Connection Manager. The connectivity services include:

- Connection to a specified element given IP-address, L4-port, and other connectivity options (such as authentication,. . .)
- Disconnection from an element
- Handling Cluster Mode change notifications to support the OpenDaylight Clustering/High-Availability feature

Network Configuration Service

The goal of the OpenDaylight Network Configuration services is to provide complete management plane solutions needed to successfully install, configure, and deploy the various SDN based network services. These are generic services which can be implemented in part or full by any south-bound protocol plugin. The south-bound plugins can be either of the following:

- The new network virtualization protocol plugins such as OVSDB JSON-RPC
- The traditional management protocols such as SNMP or any others in the middle.

The above definition, and more information on Network Configuration Services, is available at : https://wiki.opendaylight.org/view/OpenDaylight_Controller:NetworkConfigurationServices

Bidirectional JSON-RPC library

The OVSDB plugin implements a Bidirectional JSON-RPC library. It is easy to design the library as a module that manages the Netty connection towards the Element.

The main responsibilities of this Library are:

- Demarshal and marshal JSON Strings to JSON objects
- Demarshal and marshal JSON Strings from and to the Network Element.

OVSDB Schema definitions and Object mappers

The OVSDB Schema definitions and Object Mapping layer sits above the JSON-RPC library. It maps the generic JSON objects to OVSDB schema POJOs (Plain Old Java Object) and vice-versa. This layer mostly provides the Java Object definition for the corresponding OVSDB schema (13 of them) and also will provide much more friendly API abstractions on top of these object data. This helps in hiding the JSON semantics from the functional modules such as Configuration Service and Tunnel management.

On the demarshaling side the mapping logic differentiates the Request and Response messages as follows :

- Request messages are mapped by its “method”

- Response messages are mapped by their IDs which were originally populated by the Request message. The JSON semantics of these OVSDB schema is quite complex. The following figures summarize two of the end-to-end scenarios:

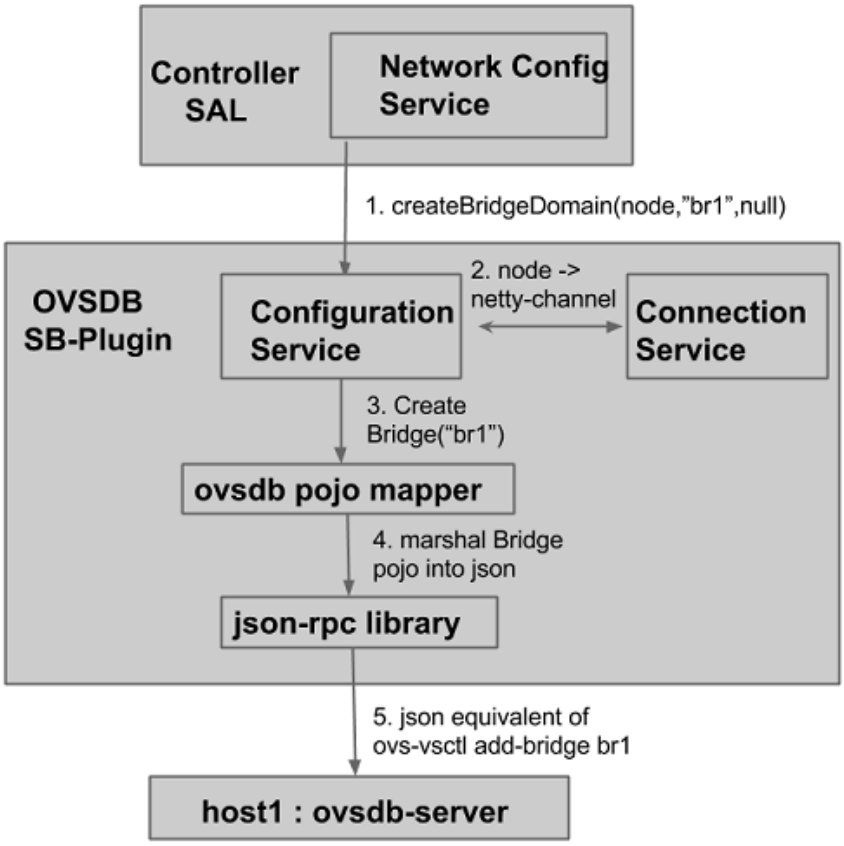

Fig. 36: End-to-end handling of a Create Bridge request

Overlay tunnel management

Network Virtualization using OVS is achieved through Overlay Tunnels. The actual Type of the Tunnel may be GRE, VXLAN, or STT. The differences in the encapsulation and configuration decide the tunnel types. Establishing a tunnel using configuration service requires just the sending of OVSDB messages towards the ovsdb-server. However, the scaling issues that would arise on the state management at the data-plane (using OpenFlow) can get challenging. Also, this module can assist in various optimizations in the presence of Gateways. It can also help in providing Service guarantees for the VMs using these overlays with the help of underlay orchestration.

OVSDB to OpenFlow plugin mapping service

The connect() of the ConnectionService would result in a Node that represents an ovsdb-server. The CreateBridgeDomain() Configuration on the above Node would result in creating an OVS bridge. This OVS Bridge is an OpenFlow Agent for the OpenDaylight OpenFlow plugin with its own Node represented as (example)

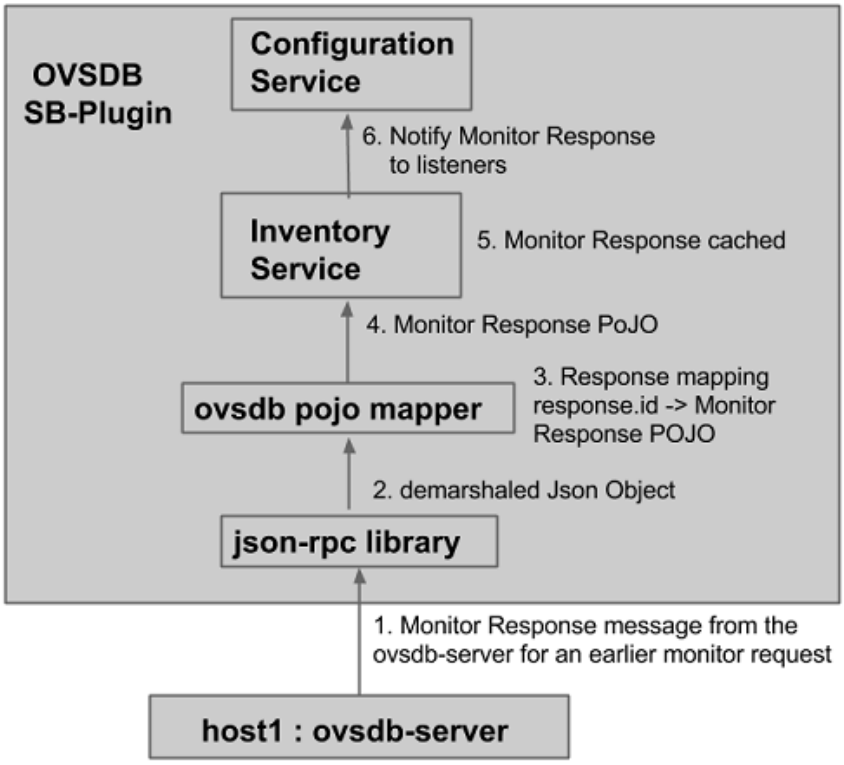

Fig. 37: End-to-end handling of a monitor response

OF|xxxx.yyyy.zzzz. Without any help from the OVSDB plugin, the Node Mapping Service of the Controller platform would not be able to map the following:

```
{OVSDB_NODE + BRIDGE_IDENTFIER} <---> {OF_NODE}.
```

Without such mapping, it would be extremely difficult for the applications to manage and maintain such nodes. This Mapping Service provided by the OVSDB plugin would essentially help in providing more value added services to the orchestration layers that sit atop the Northbound APIs (such as OpenStack).

OVSDB: New features

Schema independent library

The OVS connection is a node which can have multiple databases. Each database is represented by a schema. A single connection can have multiple schemas. OSVDB supports multiple schemas. Currently, these are two schemas available in the OVSDB, but there is no restriction on the number of schemas. Owing to the Northbound v3 API, no code changes in ODL are needed for supporting additional schemas.

Schemas:

- openvswitch : Schema wrapper that represents http://openvswitch.org/ovs-vswitchd.conf.db.5.pdf
- hardwarevtep: Schema wrapper that represents http://openvswitch.org/docs/vtep.5.pdf

OVSDB Library Developer Guide

Overview

The OVSDB library manages the Netty connections to network nodes and handles bidirectional JSON-RPC messages. It not only provides OVSDB protocol functionality to OpenDaylight OVSDB plugin but also can be used as standalone JAVA library for OVSDB protocol.

The main responsibilities of OVSDB library include:

- Manage connections to peers
- Marshal and unmarshal JSON Strings to JSON objects.
- Marshal and unmarshal JSON Strings from and to the Network Element.

Connection Service

The OVSDB library provides connection management through the OvsdbConnection interface. The OvsdbConnection interface provides OVSDB connection management APIs which include both active and passive connections. From the library perspective, active OVSDB connections are initiated from the controller to OVS nodes while passive OVSDB connections are initiated from OVS nodes to the controller. In the active connection scenario an application needs to provide the IP address and listening port of OVS nodes to the library management API. On the other hand, the library management API only requires the info of the controller listening port in the passive connection scenario.

For a passive connection scenario, the library also provides a connection event listener through the OvsdbConnectionListener interface. The listener interface has connected() and disconnected() methods to notify an application when a new passive connection is established or an existing connection is terminated.

SSL Connection

In addition to a regular TCP connection, the OvsdbConnection interface also provides a connection management API for an SSL connection. To start an OVSDB connection with SSL, an application will need to provide a Java SSLContext object to the management API. There are different ways to create a Java SSLContext, but in most cases a Java KeyStore with certificate and private key provided by the application is required. Detailed steps about how to create a Java SSLContext is out of the scope of this document and can be found in the Java documentation for JAVA Class SSlContext.

In the active connection scenario, the library uses the given SSLContext to create a Java SSLEngine and configures the SSL engine with the client mode for SSL handshaking. Normally clients are not required to authenticate themselves.

In the passive connection scenario, the library uses the given SSLContext to create a Java SSLEngine which will operate in server mode for SSL handshaking. For security reasons, the SSLv3 protocol and some cipher suites are disabled. Currently the OVSDB server only supports the TLS_RSA_WITH_AES_128_CBC_SHA cipher suite and the following protocols: SSLv2Hello, TLSv1, TLSv1.1, TLSv1.2.

The SSL engine is also configured to operate on two-way authentication mode for passive connection scenarios, i.e., the OVSDB server (controller) will authenticate clients (OVS nodes) and clients (OVS nodes) are also required to authenticate the server (controller). In the two-way authentication mode, an application should keep a trust manager to store the certificates of trusted clients and initialize a Java SSLContext with this trust manager. Thus during the SSL handshaking process the OVSDB server (controller) can use the trust manager to verify clients and only accept connection requests from trusted clients. On the other hand, users should also configure OVS nodes to authenticate the controller. Open vSwitch already supports this functionality in the ovsdb-server command with option `--ca-cert=cacert.pem` and `--bootstrap-ca-cert=cacert.pem`. On the OVS node, a user can use the option `--ca-cert=cacert.pem` to specify a controller certificate directly and the node will only allow connections to the controller with the specified certificate. If the OVS node runs ovsdb-server with option `--bootstrap-ca-cert=cacert.pem`, it will authenticate the controller with the specified certificate cacert.pem. If the certificate file doesn't exist, it will attempt to obtain a certificate from the peer (controller) on its first SSL connection and save it to the named PEM file `cacert.pem`. Here is an example of ovsdb-server with `--bootstrap-ca-cert=cacert.pem` option:

```
ovsdb-server --pidfile --detach --log-file --remote punix:/var/run/
openvswitch/db.sock --remote=db:hardware_vtep,Global,managers --private-key=/
etc/openvswitch/ovsclient-privkey.pem -- certificate=/etc/openvswitch/
ovsclient-cert.pem --bootstrap-ca-cert=/etc/openvswitch/vswitchd.cacert
```

OVSDB protocol transactions

The OVSDB protocol defines the RPC transaction methods in RFC 7047. The following RPC methods are supported in OVSDB protocol:

- List databases
- Get schema
- Transact
- Cancel
- Monitor
- Update notification

- Monitor cancellation
- Lock operations
- Locked notification
- Stolen notification
- Echo

According to RFC 7047, an OVSDB server must implement all methods, and an OVSDB client is only required to implement the "Echo" method and otherwise free to implement whichever methods suit its needs. However, the OVSDB library currently doesn't support all RPC methods. For the "Echo" method, the library can handle "Echo" messages from a peer and send a JSON response message back, but the library doesn't support actively sending an "Echo" JSON request to a peer. Other unsupported RPC methods are listed below:

- Cancel
- Lock operations
- Locked notification
- Stolen notification

In the OVSDB library the RPC methods are defined in the Java interface OvsdbRPC. The library also provides a high-level interface OvsdbClient as the main interface to interact with peers through the OVSDB protocol. In the passive connection scenario, each connection will have a corresponding OvsdbClient object, and the application can obtain the OvsdbClient object through connection listener callback methods. In other words, if the application implements the OvsdbConnectionListener interface, it will get notifications of connection status changes with the corresponding OvsdbClient object of that connection.

OVSDB database operations

RFC 7047 also defines database operations, such as insert, delete, and update, to be performed as part of a "transact" RPC request. The OVSDB library defines the data operations in Operations.java and provides the TransactionBuilder class to help build "transact" RPC requests. To build a JSON-RPC transact request message, the application can obtain the TransactionBuilder object through a transactBuilder() method in the OvsdbClient interface.

The TransactionBuilder class provides the following methods to help build transactions:

- getOperations(): Get the list of operations in this transaction.
- add(): Add data operation to this transaction.
- build(): Return the list of operations in this transaction. This is the same as the getOperations() method.
- execute(): Send the JSON RPC transaction to peer.
- getDatabaseSchema(): Get the database schema of this transaction.

If the application wants to build and send a "transact" RPC request to modify OVSDB tables on a peer, it can take the following steps:

1. Statically import parameter "op" in Operations.java

   ```
   import static org.opendaylight.ovsdb.lib.operations.Operations.op;
   ```

2. Obtain transaction builder through transacBuilder() method in OvsdbClient:

   ```
   TransactionBuilder transactionBuilder = ovsdbClient.transactionBuilder(dbSchema);
   ```

3. Add operations to transaction builder:

   ```
   transactionBuilder.add(op.insert(schema, row));
   ```

4. Send transaction to peer and get JSON RPC response:

   ```
   operationResults = transactionBuilder.execute().get();
   ```

Note: Although the "select" operation is supported in the OVSDB library, the library implementation is a little different from RFC 7047. In RFC 7047, section 5.2.2 describes the "select" operation as follows:

"The "rows" member of the result is an array of objects. Each object corresponds to a matching row, with each column specified in "columns" as a member, the column's name as the member name, and its value as the member value. If "columns" is not specified, all the table's columns are included (including the internally generated "_uuid" and "_version" columns)."

The OVSDB library implementation always requires the column's name in the "columns" field of a JSON message. If the "columns" field is not specified, none of the table's columns are included. If the application wants to get the table entry with all columns, it needs to specify all the columns' names in the "columns" field.

Reference Documentation

RFC 7047 The Open vSwitch Databse Management Protocol https://tools.ietf.org/html/rfc7047

OVSDB MD-SAL Southbound Plugin Developer Guide

Overview

The Open vSwitch Database (OVSDB) Model Driven Service Abstraction Layer (MD-SAL) Southbound Plugin provides an MD-SAL based interface to Open vSwitch systems. This is done by augmenting the MD-SAL topology node with a YANG model which replicates some (but not all) of the Open vSwitch schema.

OVSDB MD-SAL Southbound Plugin Architecture and Operation

The architecture and operation of the OVSDB MD-SAL Southbound plugin is illustrated in the following set of diagrams.

Connecting to an OVSDB Node

An OVSDB node is a system which is running the OVS software and is capable of being managed by an OVSDB manager. The OVSDB MD-SAL Southbound plugin in OpenDaylight is capable of operating as an OVSDB manager. Depending on the configuration of the OVSDB node, the connection of the OVSDB manager can be active or passive.

Active OVSDB Node Manager Workflow

An active OVSDB node manager connection is made when OpenDaylight initiates the connection to the OVSDB node. In order for this to work, you must configure the OVSDB node to listen on a TCP port for the connection (i.e. OpenDaylight is active and the OVSDB node is passive). This option can be configured on the OVSDB node using the following command:

```
ovs-vsctl set-manager ptcp:6640
```

The following diagram illustrates the sequence of events which occur when OpenDaylight initiates an active OVSDB manager connection to an OVSDB node.

Step 1 Create an OVSDB node by using RESTCONF or an OpenDaylight plugin. The OVSDB node is listed under the OVSDB topology node.

Step 2 Add the OVSDB node to the OVSDB MD-SAL southbound configuration datastore. The OVSDB southbound provider is registered to listen for data change events on the portion of the MD-SAL topology data store which contains the OVSDB southbound topology node augmentations. The addition of an OVSDB node causes an event which is received by the OVSDB Southbound provider.

Step 3 The OVSDB Southbound provider initiates a connection to the OVSDB node using the connection information provided in the configuration OVSDB node (i.e. IP address and TCP port number).

Step 4 The OVSDB Southbound provider adds the OVSDB node to the OVSDB MD-SAL operational data store. The operational data store contains OVSDB node objects which represent active connections to OVSDB nodes.

Step 5 The OVSDB Southbound provider requests the schema and databases which are supported by the OVSDB node.

Step 6 The OVSDB Southbound provider uses the database and schema information to construct a monitor request which causes the OVSDB node to send the controller any updates made to the OVSDB databases on the OVSDB node.

Passive OVSDB Node Manager Workflow

A passive OVSDB node connection to OpenDaylight is made when the OVSDB node initiates the connection to OpenDaylight. In order for this to work, you must configure the OVSDB node to connect to the IP address and OVSDB port on which OpenDaylight is listening. This option can be configured on the OVSDB node using the following command:

```
ovs-vsctl set-manager tcp:<IP address>:6640
```

The following diagram illustrates the sequence of events which occur when an OVSDB node connects to OpenDaylight.

Step 1 The OVSDB node initiates a connection to OpenDaylight.

Step 2 The OVSDB Southbound provider adds the OVSDB node to the OVSDB MD-SAL operational data store. The operational data store contains OVSDB node objects which represent active connections to OVSDB nodes.

Step 3 The OVSDB Southbound provider requests the schema and databases which are supported by the OVSDB node.

Step 4 The OVSDB Southbound provider uses the database and schema information to construct a monitor request which causes the OVSDB node to send back any updates which have been made to the OVSDB databases on the OVSDB node.

OVSDB Node ID in the Southbound Operational MD-SAL

When OpenDaylight initiates an active connection to an OVSDB node, it writes an external-id to the Open_vSwitch table on the OVSDB node. The external-id is an OpenDaylight instance identifier which identifies the OVSDB topology node which has just been created. Here is an example showing the value of the *opendaylight-iid* entry in the external-ids column of the Open_vSwitch table where the node-id of the OVSDB node is *ovsdb:HOST1*.

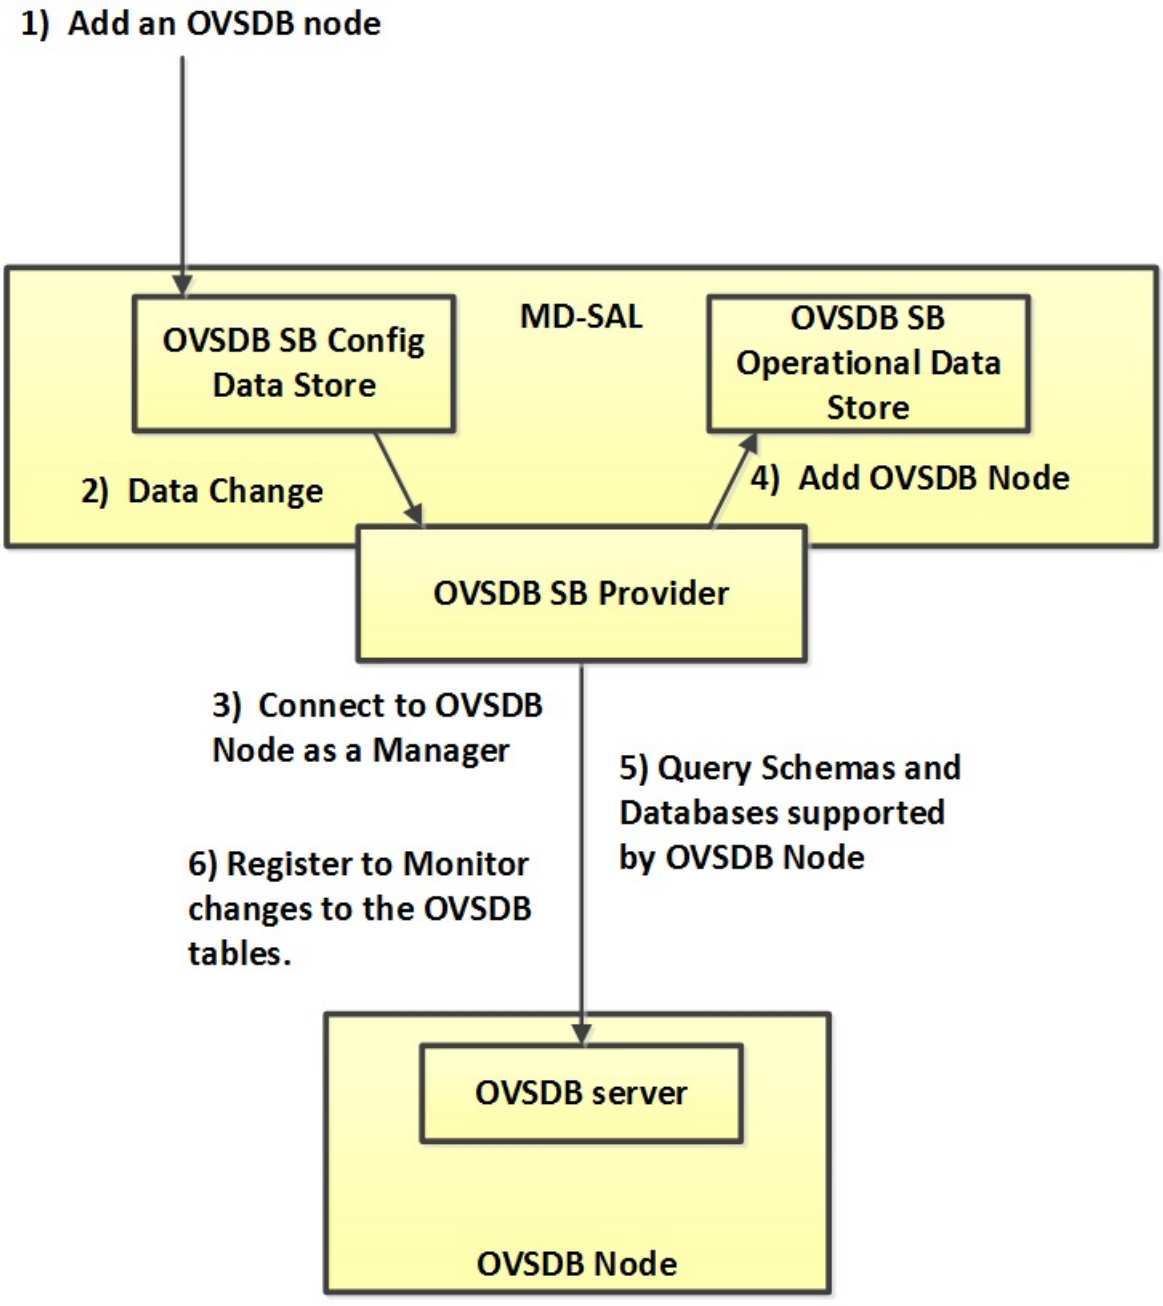

Fig. 38: Active OVSDB Manager Connection

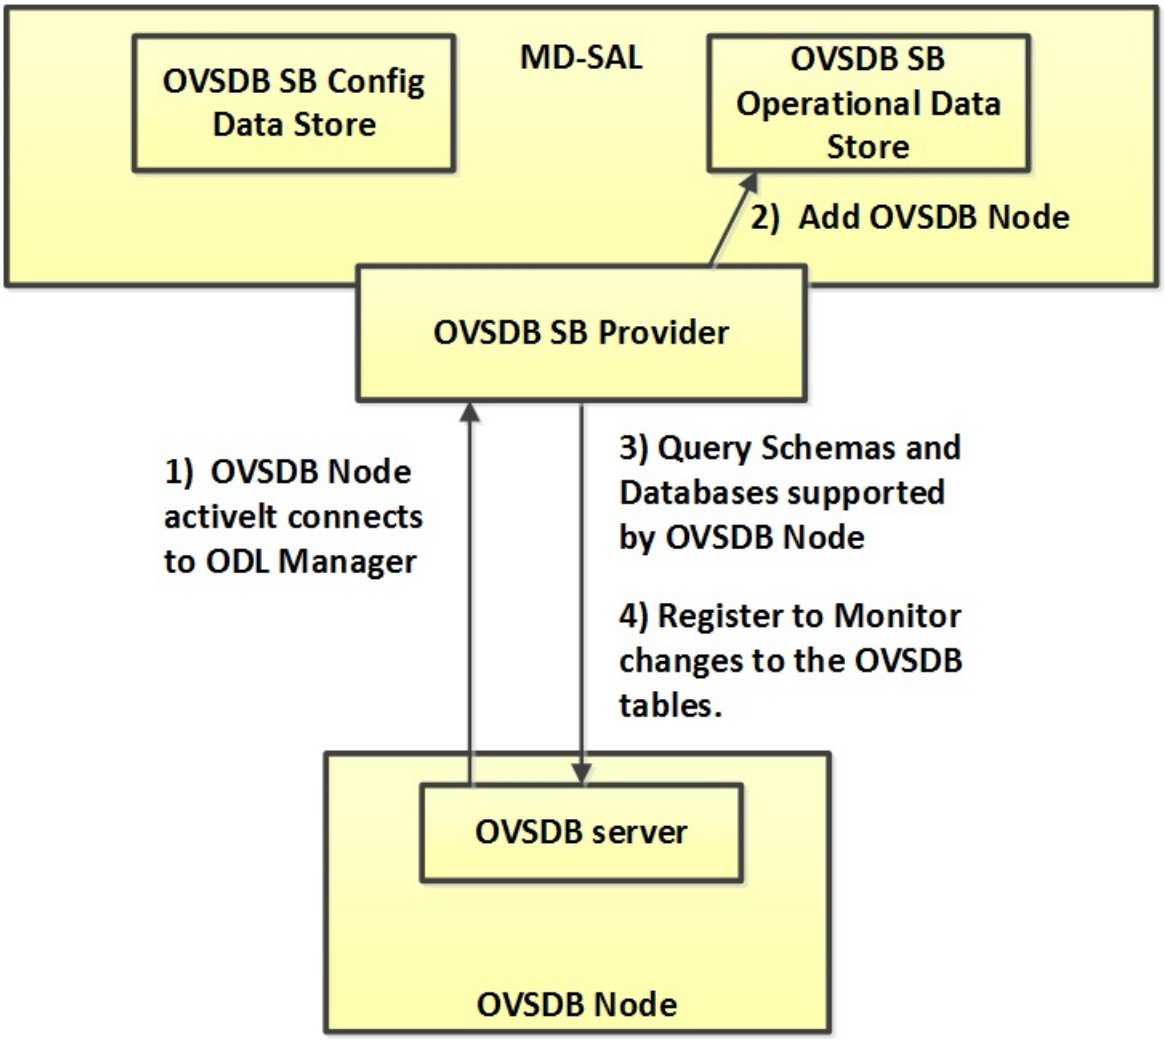

Fig. 39: Passive OVSDB Manager Connection

```
$ ovs-vsctl list open_vswitch
...
external_ids        : {opendaylight-iid="/network-topology:network-topology/network-
↪topology:topology[network-topology:topology-id='ovsdb:1']/network-
↪topology:node[network-topology:node-id='ovsdb:HOST1']"}
...
```

The *opendaylight-iid* entry in the external-ids column of the Open_vSwitch table causes the OVSDB node to have same node-id in the operational MD-SAL datastore as in the configuration MD-SAL datastore. This holds true if the OVSDB node manager settings are subsequently changed so that a passive OVSDB manager connection is made.

If there is no *opendaylight-iid* entry in the external-ids column and a passive OVSDB manager connection is made, then the node-id of the OVSDB node in the operational MD-SAL datastore will be constructed using the UUID of the Open_vSwitch table as follows.

```
"node-id": "ovsdb://uuid/b8dc0bfb-d22b-4938-a2e8-b0084d7bd8c1"
```

The *opendaylight-iid* entry can be removed from the Open_vSwitch table using the following command.

```
$ sudo ovs-vsctl remove open_vswitch . external-id "opendaylight-iid"
```

OVSDB Changes by using OVSDB Southbound Config MD-SAL

After the connection has been made to an OVSDB node, you can make changes to the OVSDB node by using the OVSDB Southbound Config MD-SAL. You can make CRUD operations by using the RESTCONF interface or by a plugin using the MD-SAL APIs. The following diagram illustrates the high-level flow of events.

Step 1 A change to the OVSDB Southbound Config MD-SAL is made. Changes include adding or deleting bridges and ports, or setting attributes of OVSDB nodes, bridges or ports.

Step 2 The OVSDB Southbound provider receives notification of the changes made to the OVSDB Southbound Config MD-SAL data store.

Step 3 As appropriate, OVSDB transactions are constructed and transmitted to the OVSDB node to update the OVSDB database on the OVSDB node.

Step 4 The OVSDB node sends update messages to the OVSDB Southbound provider to indicate the changes made to the OVSDB nodes database.

Step 5 The OVSDB Southbound provider maps the changes received from the OVSDB node into corresponding changes made to the OVSDB Southbound Operational MD-SAL data store.

Detecting changes in OVSDB coming from outside OpenDaylight

Changes to the OVSDB nodes database may also occur independently of OpenDaylight. OpenDaylight also receives notifications for these events and updates the Southbound operational MD-SAL. The following diagram illustrates the sequence of events.

Step 1 Changes are made to the OVSDB node outside of OpenDaylight (e.g. ovs-vsctl).

Step 2 The OVSDB node constructs update messages to inform OpenDaylight of the changes made to its databases.

Step 3 The OVSDB Southbound provider maps the OVSDB database changes to corresponding changes in the OVSDB Southbound operational MD-SAL data store.

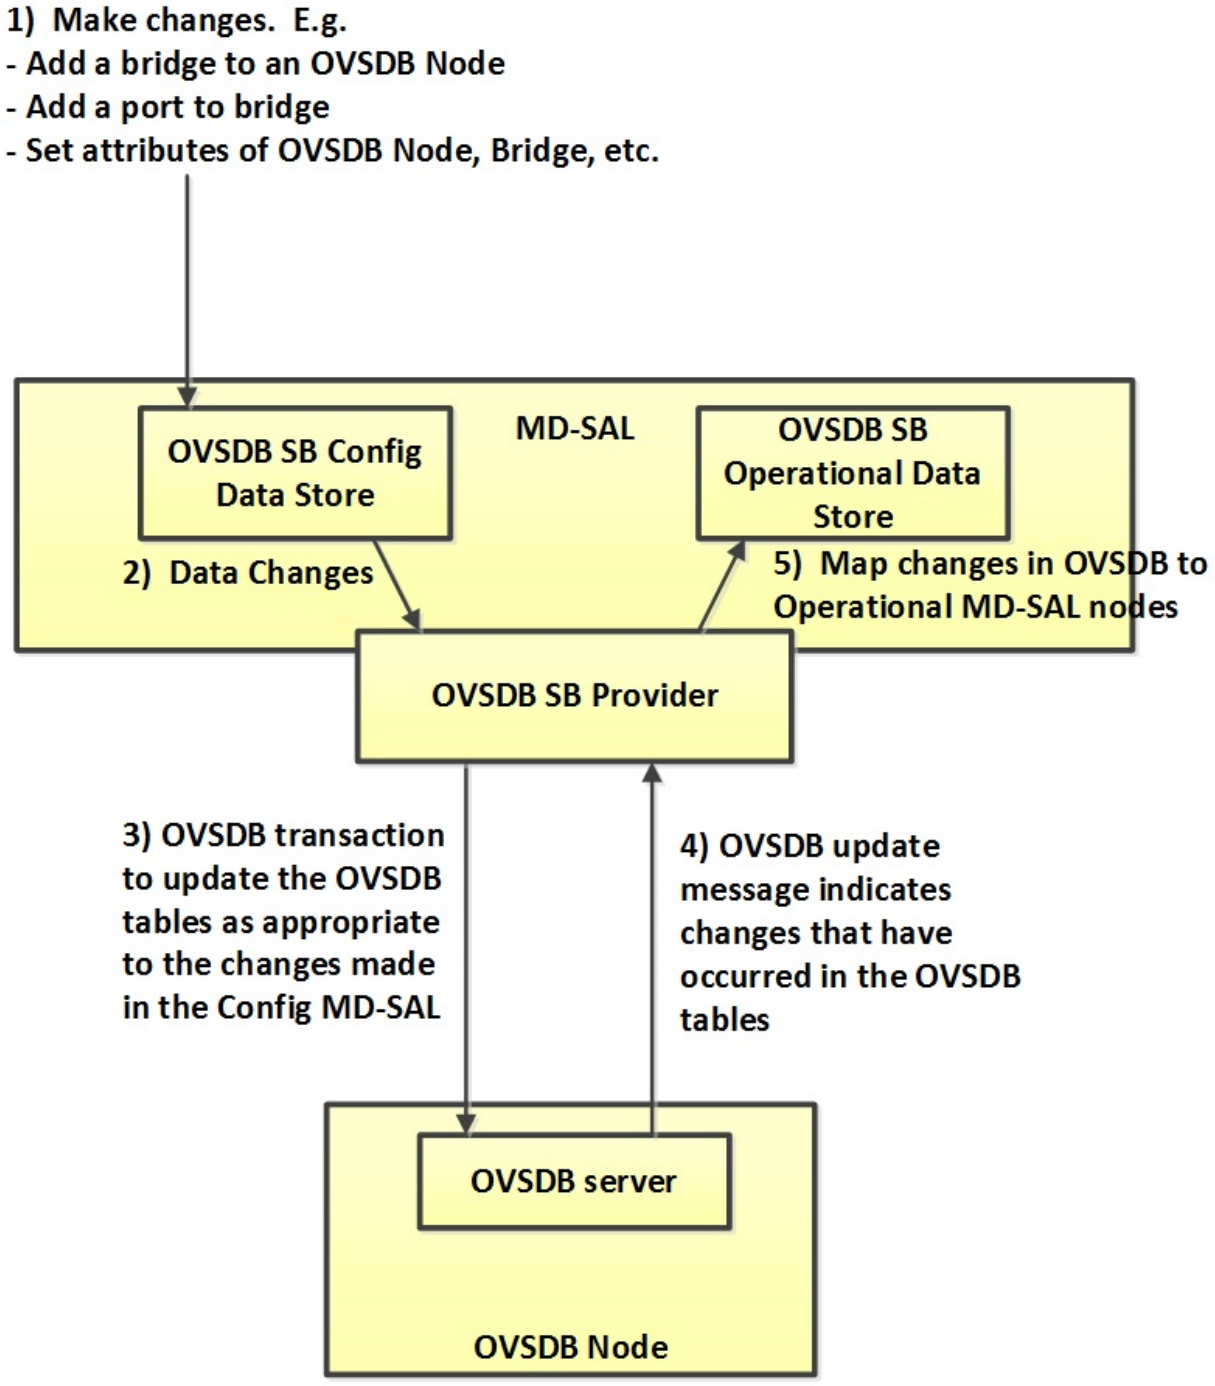

Fig. 40: OVSDB Changes by using the Southbound Config MD-SAL

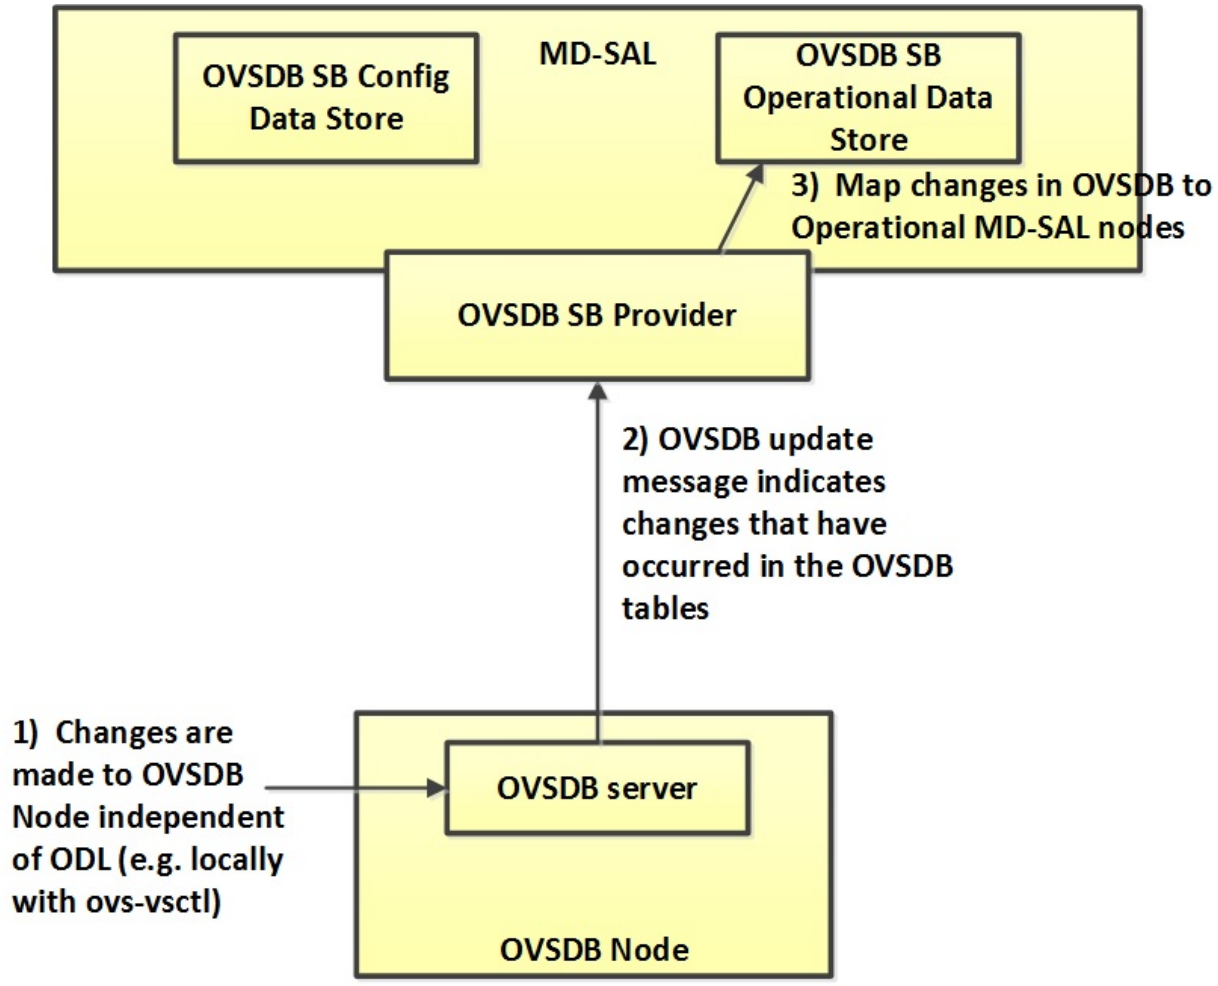

Fig. 41: OVSDB Changes made directly on the OVSDB node

OVSDB Model

The OVSDB Southbound MD-SAL operates using a YANG model which is based on the abstract topology node model found in the network topology model.

The augmentations for the OVSDB Southbound MD-SAL are defined in the ovsdb.yang file.

There are three augmentations:

ovsdb-node-augmentation This augments the topology node and maps primarily to the Open_vSwitch table of the OVSDB schema. It contains the following attributes.

- **connection-info** - holds the local and remote IP address and TCP port numbers for the OpenDaylight to OVSDB node connections
- **db-version** - version of the OVSDB database
- **ovs-version** - version of OVS
- **list managed-node-entry** - a list of references to ovsdb-bridge-augmentation nodes, which are the OVS bridges managed by this OVSDB node
- **list datapath-type-entry** - a list of the datapath types supported by the OVSDB node (e.g. *system*, *netdev*) - depends on newer OVS versions
- **list interface-type-entry** - a list of the interface types supported by the OVSDB node (e.g. *internal*, *vxlan*, *gre*, *dpdk*, etc.) - depends on newer OVS verions
- **list openvswitch-external-ids** - a list of the key/value pairs in the Open_vSwitch table external_ids column
- **list openvswitch-other-config** - a list of the key/value pairs in the Open_vSwitch table other_config column

ovsdb-bridge-augmentation This augments the topology node and maps to an specific bridge in the OVSDB bridge table of the associated OVSDB node. It contains the following attributes.

- **bridge-uuid** - UUID of the OVSDB bridge
- **bridge-name** - name of the OVSDB bridge
- **bridge-openflow-node-ref** - a reference (instance-identifier) of the OpenFlow node associated with this bridge
- **list protocol-entry** - the version of OpenFlow protocol to use with the OpenFlow controller
- **list controller-entry** - a list of controller-uuid and is-connected status of the OpenFlow controllers associated with this bridge
- **datapath-id** - the datapath ID associated with this bridge on the OVSDB node
- **datapath-type** - the datapath type of this bridge
- **fail-mode** - the OVSDB fail mode setting of this bridge
- **flow-node** - a reference to the flow node corresponding to this bridge
- **managed-by** - a reference to the ovsdb-node-augmentation (OVSDB node) that is managing this bridge
- **list bridge-external-ids** - a list of the key/value pairs in the bridge table external_ids column for this bridge
- **list bridge-other-configs** - a list of the key/value pairs in the bridge table other_config column for this bridge

ovsdb-termination-point-augmentation This augments the topology termination point model. The OVSDB Southbound MD-SAL uses this model to represent both the OVSDB port and OVSDB interface for a given port/interface in the OVSDB schema. It contains the following attributes.

- **port-uuid** - UUID of an OVSDB port row
- **interface-uuid** - UUID of an OVSDB interface row
- **name** - name of the port
- **interface-type** - the interface type
- **list options** - a list of port options
- **ofport** - the OpenFlow port number of the interface
- **ofport_request** - the requested OpenFlow port number for the interface
- **vlan-tag** - the VLAN tag value
- **list trunks** - list of VLAN tag values for trunk mode
- **vlan-mode** - the VLAN mode (e.g. access, native-tagged, native-untagged, trunk)
- **list port-external-ids** - a list of the key/value pairs in the port table external_ids column for this port
- **list interface-external-ids** - a list of the key/value pairs in the interface table external_ids interface for this interface
- **list port-other-configs** - a list of the key/value pairs in the port table other_config column for this port
- **list interface-other-configs** - a list of the key/value pairs in the interface table other_config column for this interface

Examples of OVSDB Southbound MD-SAL API

Connect to an OVSDB Node

This example RESTCONF command adds an OVSDB node object to the OVSDB Southbound configuration data store and attempts to connect to the OVSDB host located at the IP address 10.11.12.1 on TCP port 6640.

```
POST http://<host>:8181/restconf/config/network-topology:network-topology/topology/
↪ovsdb:1/
Content-Type: application/json
{
  "node": [
     {
       "node-id": "ovsdb:HOST1",
       "connection-info": {
         "ovsdb:remote-ip": "10.11.12.1",
         "ovsdb:remote-port": 6640
       }
     }
  ]
}
```

Query the OVSDB Southbound Configuration MD-SAL

Following on from the previous example, if the OVSDB Southbound configuration MD-SAL is queried, the RESTCONF command and the resulting reply is similar to the following example.

```
GET http://<host>:8080/restconf/config/network-topology:network-topology/topology/
↪ovsdb:1/
Application/json data in the reply
{
  "topology": [
    {
      "topology-id": "ovsdb:1",
      "node": [
        {
          "node-id": "ovsdb:HOST1",
          "ovsdb:connection-info": {
            "remote-port": 6640,
            "remote-ip": "10.11.12.1"
          }
        }
      ]
    }
  ]
}
```

Reference Documentation

Openvswitch schema

OVSDB Hardware VTEP Developer Guide

Overview

TBD

OVSDB Hardware VTEP Architecture

TBD

PCEP Developer Guide

Overview

This section provides an overview of **feature odl-bgpcep-pcep-all** . This feature will install everything needed for PCEP (Path Computation Element Protocol) including establishing the connection, storing information about LSPs (Label Switched Paths) and displaying data in network-topology overview.

PCEP Architecture

Each feature represents a module in the BGPCEP codebase. The following diagram illustrates how the features are related.

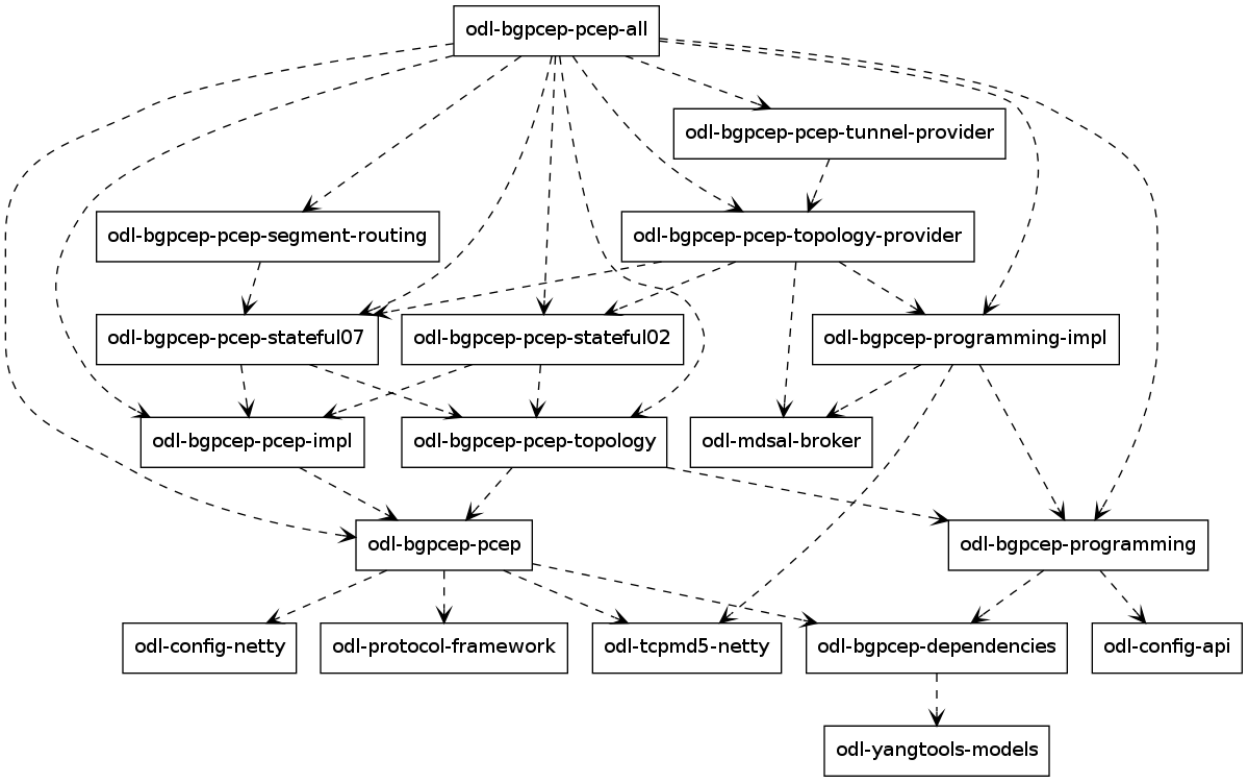

Fig. 42: PCEP Dependency Tree

Key APIs and Interfaces

PCEP

Session handling

32-pcep.xml defines only pcep-dispatcher the parser should be using (global-pcep-extensions), factory for creating session proposals (you can create different proposals for different PCCs (Path Computation Clients)).

```
<module>
 <type xmlns:prefix="urn:opendaylight:params:xml:ns:yang:controller:pcep:impl">
↪prefix:pcep-dispatcher-impl</type>
 <name>global-pcep-dispatcher</name>
 <pcep-extensions>
  <type xmlns:pcepspi="urn:opendaylight:params:xml:ns:yang:controller:pcep:spi">
↪pcepspi:extensions</type>
  <name>global-pcep-extensions</name>
 </pcep-extensions>
 <pcep-session-proposal-factory>
  <type xmlns:pcep="urn:opendaylight:params:xml:ns:yang:controller:pcep">pcep:pcep-
↪session-proposal-factory</type>
  <name>global-pcep-session-proposal-factory</name>
 </pcep-session-proposal-factory>
 <boss-group>
  <type xmlns:netty="urn:opendaylight:params:xml:ns:yang:controller:netty">
↪netty:netty-threadgroup</type>
```

(continues on next page)

(continued from previous page)

```
  <name>global-boss-group</name>
 </boss-group>
 <worker-group>
  <type xmlns:netty="urn:opendaylight:params:xml:ns:yang:controller:netty">
↪netty:netty-threadgroup</type>
  <name>global-worker-group</name>
 </worker-group>
</module>
```

For user configuration of PCEP, check User Guide.

Parser

The base PCEP parser includes messages and attributes from RFC5441, RFC5541, RFC5455, RFC5557 and RFC5521.

Registration

All parsers and serializers need to be registered into *Extension provider*. This *Extension provider* is configured in initial configuration of the parser-spi module (*32-pcep.xml*).

```
<module>
 <type xmlns:prefix="urn:opendaylight:params:xml:ns:yang:controller:pcep:spi">
↪prefix:pcep-extensions-impl</type>
 <name>global-pcep-extensions</name>
 <extension>
  <type xmlns:pcepspi="urn:opendaylight:params:xml:ns:yang:controller:pcep:spi">
↪pcepspi:extension</type>
  <name>pcep-parser-base</name>
 </extension>
 <extension>
  <type xmlns:pcepspi="urn:opendaylight:params:xml:ns:yang:controller:pcep:spi">
↪pcepspi:extension</type>
  <name>pcep-parser-ietf-stateful07</name>
 </extension>
 <extension>
  <type xmlns:pcepspi="urn:opendaylight:params:xml:ns:yang:controller:pcep:spi">
↪pcepspi:extension</type>
  <name>pcep-parser-ietf-initiated00</name>
 </extension>
 <extension>
  <type xmlns:pcepspi="urn:opendaylight:params:xml:ns:yang:controller:pcep:spi">
↪pcepspi:extension</type>
  <name>pcep-parser-sync-optimizations</name>
 </extension>
</module>
```

- *pcep-parser-base* - will register parsers and serializers implemented in pcep-impl module
- *pcep-parser-ietf-stateful07* - will register parsers and serializers of draft-ietf-pce-stateful-pce-07 implementation
- *pcep-parser-ietf-initiated00* - will register parser and serializer of draft-ietf-pce-pce-initiated-lsp-00 implementation
- *pcep-parser-sync-optimizations* - will register parser and serializers of draft-ietf-pce-stateful-sync-optimizations-03 implementation

Stateful07 module is a good example of a PCEP parser extension.

Configuration of PCEP parsers specifies one implementation of *Extension provider* that will take care of registering mentioned parser extensions: SimplePCEPExtensionProviderContext. All registries are implemented in package pcep-spi.

Parsing

Parsing of PCEP elements is mostly done equally to BGP, the only exception is message parsing, that is described here.

In BGP messages, parsing of first-level elements (path-attributes) can be validated in a simple way, as the attributes should be ordered chronologically. PCEP, on the other hand, has a strict object order policy, that is described in RBNF (Routing Backus-Naur Form) in each RFC. Therefore the algorithm for parsing here is to parse all objects in order as they appear in the message. The result of parsing is a list of *PCEPObjects*, that is put through validation. *validate()* methods are present in each message parser. Depending on the complexity of the message, it can contain either a simple condition (checking the presence of a mandatory object) or a full state machine.

In addition to that, PCEP requires sending error message for each documented parsing error. This is handled by creating an empty list of messages *errors* which is then passed as argument throughout whole parsing process. If some parser encounters *PCEPDocumentedException*, it has the duty to create appropriate PCEP error message and add it to this list. In the end, when the parsing is finished, this list is examined and all messages are sent to peer.

Better understanding provides this sequence diagram:

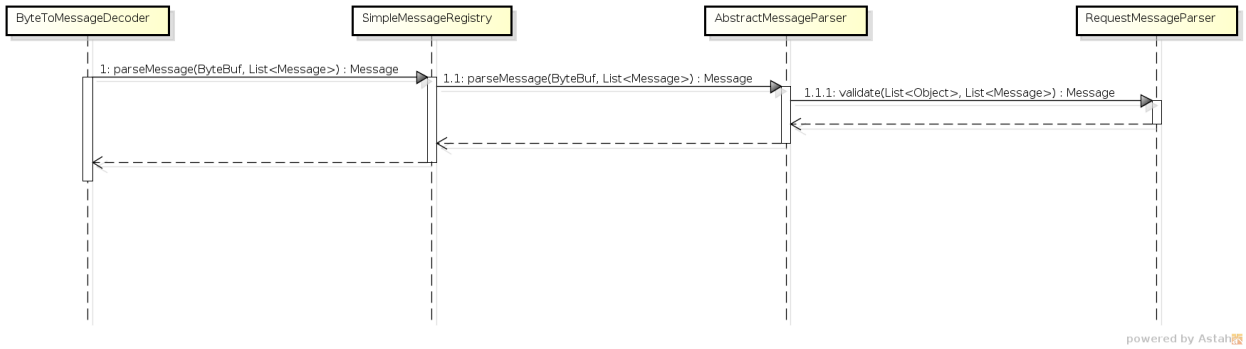

Fig. 43: Parsing

PCEP IETF stateful

This section summarizes module pcep-ietf-stateful07. The term *stateful* refers to draft-ietf-pce-stateful-pce and draft-ietf-pce-pce-initiated-lsp in versions draft-ietf-pce-stateful-pce-07 with draft-ietf-pce-pce-initiated-lsp-00.

We will upgrade our implementation, when the stateful draft gets promoted to RFC.

The stateful module is implemented as extensions to pcep-base-parser. The stateful draft declared new elements as well as additional fields or TLVs (type,length,value) to known objects. All new elements are defined in yang models, that contain augmentations to elements defined in pcep-types.yang. In the case of extending known elements, the *Parser* class merely extends the base class and overrides necessary methods as shown in following diagram:

All parsers (including those for newly defined PCEP elements) have to be registered via the *Activator* class. This class is present in both modules.

In addition to parsers, the stateful module also introduces additional session proposal. This proposal includes new fields defined in stateful drafts for Open object.

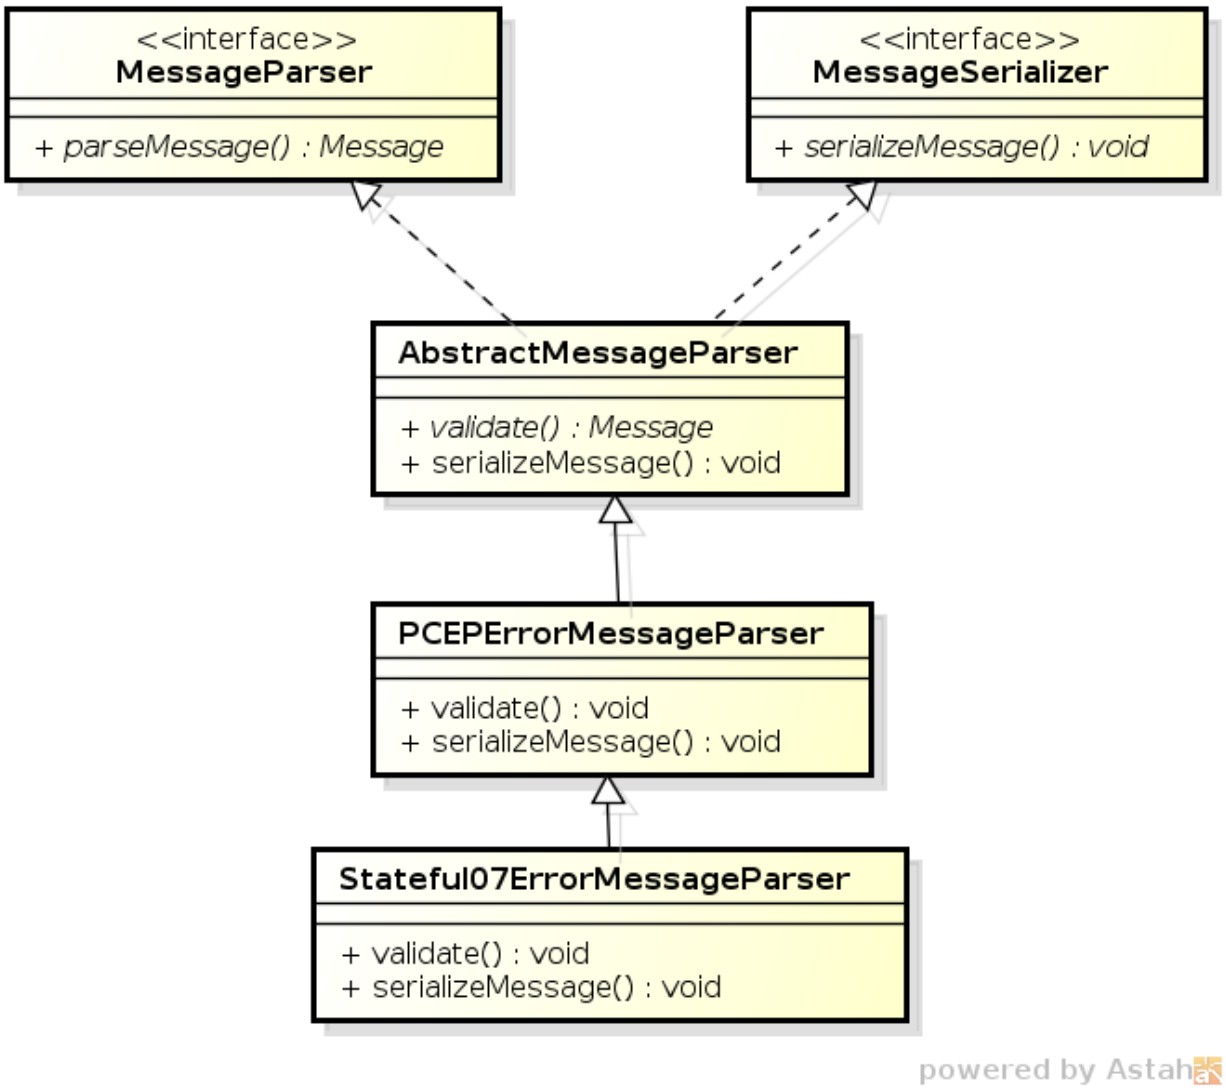

Fig. 44: Extending existing parsers

PCEP segment routing (SR)

PCEP Segment Routing is an extension of base PCEP and pcep-ietf-stateful-07 extension. The pcep-segment-routing module implements draft-ietf-pce-segment-routing-01.

The extension brings new SR-ERO (Explicit Route Object) and SR-RRO (Reported Route Object) subobject composed of SID (Segment Identifier) and/or NAI (Node or Adjacency Identifier). The segment Routing path is carried in the ERO and RRO object, as a list of SR-ERO/SR-RRO subobjects in an order specified by the user. The draft defines new TLV - SR-PCE-CAPABILITY TLV, carried in PCEP Open object, used to negotiate Segment Routing ability.

The yang models of subobject, SR-PCE-CAPABILITY TLV and appropriate augmentations are defined in odl-pcep-segment-routing.yang.
The pcep-segment-routing module includes parsers/serializers for new subobject (SrEroSubobjectParser) and TLV (SrPceCapabilityTlvParser).

The pcep-segment-routing module implements draft-ietf-pce-lsp-setup-type-01, too. The draft defines new TLV - Path Setup Type TLV, which value indicate path setup signaling technique. The TLV may be included in RP(Request Parameters)/SRP(Stateful PCE Request Parameters) object. For the default RSVP-TE (Resource Reservation Protocol), the TLV is omitted. For Segment Routing, PST = 1 is defined.

The Path Setup Type TLV is modeled with yang in module pcep-types.yang. A parser/serializer is implemented in PathSetupTypeTlvParser and it is overriden in segment-routing module to provide the aditional PST.

PCEP Synchronization Procedures Optimization

Optimizations of Label Switched Path State Synchronization Procedures for a Stateful PCE draft-ietf-pce-stateful-sync-optimizations-03 specifies following optimizations for state synchronization and the corresponding PCEP procedures and extensions:

- **State Synchronization Avoidance:** To skip state synchronization if the state has survived and not changed during session restart.
- **Incremental State Synchronization:** To do incremental (delta) state synchronization when possible.
- **PCE-triggered Initial Synchronization:** To let PCE control the timing of the initial state synchronization. The capability can be applied to both full and incremental state synchronization.
- **PCE-triggered Re-synchronization:** To let PCE re-synchronize the state for sanity check.

PCEP Topology

PCEP data is displayed only through one URL that is accessible from the base network-topology URL:

http://localhost:8181/restconf/operational/network-topology:network-topology/topology/pcep-topology

Each PCC will be displayed as a node:

```
<node>
 <path-computation-client>
  <ip-address>42.42.42.42</ip-address>
  <state-sync>synchronized</state-sync>
  <stateful-tlv>
   <stateful>
```

(continues on next page)

(continued from previous page)

```
    <initiation>true</initiation>
    <lsp-update-capability>true</lsp-update-capability>
   </stateful>
  </stateful-tlv>
 </path-computation-client>
 <node-id>pcc://42.42.42.42</node-id>
</node>
</source>
```

If some tunnels are configured on the network, they would be displayed on the same page, within a node that initiated the tunnel:

```
<node>
 <path-computation-client>
  <state-sync>synchronized</state-sync>
  <stateful-tlv>
   <stateful>
    <initiation>true</initiation>
    <lsp-update-capability>true</lsp-update-capability>
   </stateful>
  </stateful-tlv>
  <reported-lsp>
   <name>foo</name>
   <lsp>
    <operational>down</operational>
    <sync>false</sync>
    <ignore>false</ignore>
    <plsp-id>1</plsp-id>
    <create>false</create>
    <administrative>true</administrative>
    <remove>false</remove>
    <delegate>true</delegate>
    <processing-rule>false</processing-rule>
    <tlvs>
    <lsp-identifiers>
      <ipv4>
       <ipv4-tunnel-sender-address>43.43.43.43</ipv4-tunnel-sender-address>
       <ipv4-tunnel-endpoint-address>0.0.0.0</ipv4-tunnel-endpoint-address>
       <ipv4-extended-tunnel-id>0.0.0.0</ipv4-extended-tunnel-id>
      </ipv4>
      <tunnel-id>0</tunnel-id>
      <lsp-id>0</lsp-id>
     </lsp-identifiers>
     <symbolic-path-name>
      <path-name>Zm9v</path-name>
     </symbolic-path-name>
    </tlvs>
   </lsp>
  </reported-lsp>
  <ip-address>43.43.43.43</ip-address>
 </path-computation-client>
 <node-id>pcc://43.43.43.43</node-id>
</node>
```

Note that, the <*path-name*> tag displays tunnel name in Base64 encoding.

API Reference Documentation

Javadocs are generated while creating mvn:site and they are located in target/ directory in each module.

PacketCable Developer Guide

PCMM Specification

PacketCable™ Multimedia Specification

System Overview

These components introduce a DOCSIS QoS Service Flow management using the PCMM protocol. The driver component is responsible for the PCMM/COPS/PDP functionality required to service requests from PacketCable Provider and FlowManager. Requests are transposed into PCMM Gate Control messages and transmitted via COPS to the CCAP/CMTS. This plugin adheres to the PCMM/COPS/PDP functionality defined in the CableLabs specification. PacketCable solution is an MDSAL compliant component.

PacketCable Components

The packetcable maven project is comprised of several modules.

Bundle	Description
packetcable-driver	A common module that containts the COPS stack and manages all connections to CCAPS/CMTSes.
packetcable-emulator	A basic CCAP emulator to facilitate testing the the plugin when no physical CCAP is avaible.
packetcable-policy-karaf	Generates a Karaf distribution with a config that loads all the packetcable features at runtime.
packetcable-policy-model	Contains the YANG information model.
packetcable-policy-server	Provider hosts the model processing, RESTCONF, and API implementation.

Setting Logging Levels

From the Karaf console

```
log:set <LEVEL> (<PACKAGE>|<BUNDLE>)
Example
log:set DEBUG org.opendaylight.packetcable.packetcable-policy-server
```

Tools for Testing

Postman REST client for Chrome

Install the Chrome extension

Download and import sample packetcable collection

View Rest API

1. Install the `odl-mdsal-apidocs` feature from the karaf console.
2. Open http://localhost:8181/apidoc/explorer/index.html default dev build user/pass is admin/admin
3. Navigate to the PacketCable section.

Yang-IDE

Editing yang can be done in any text editor but Yang-IDE will help prevent mistakes.

Setup and Build Yang-IDE for Eclipse

Using Wireshark to Trace PCMM

1. To start wireshark with privileges issue the following command:

```
sudo wireshark &
```

2. Select the interface to monitor.
3. Use the Filter to only display COPS messages by applying "cops" in the filter field.

Debugging and Verifying DQoS Gate (Flows) on the CCAP/CMTS

Below are some of the most useful CCAP/CMTS commands to verify flows have been enabled on the CMTS.

Cisco

Cisco CMTS Cable Command Reference

Find the Cable Modem

```
10k2-DSG#show cable modem
                                                                                D
MAC Address    IP Address      I/F           MAC             Prim RxPwr  Timing Num I
                                             State           Sid  (dBmv) Offset CPE P
0010.188a.faf6 0.0.0.0         C8/0/0/U0     offline         1    0.00   1482   0   N
74ae.7600.01f3 10.32.115.150   C8/0/10/U0    online          1    -0.50  1431   0   Y
0010.188a.fad8 10.32.115.142   C8/0/10/UB    w-online        2    -0.50  1507   1   Y
000e.0900.00dd 10.32.115.143   C8/0/10/UB    w-online        3    1.00   1677   0   Y
e86d.5271.304f 10.32.115.168   C8/0/10/UB    w-online        6    -0.50  1419   1   Y
```

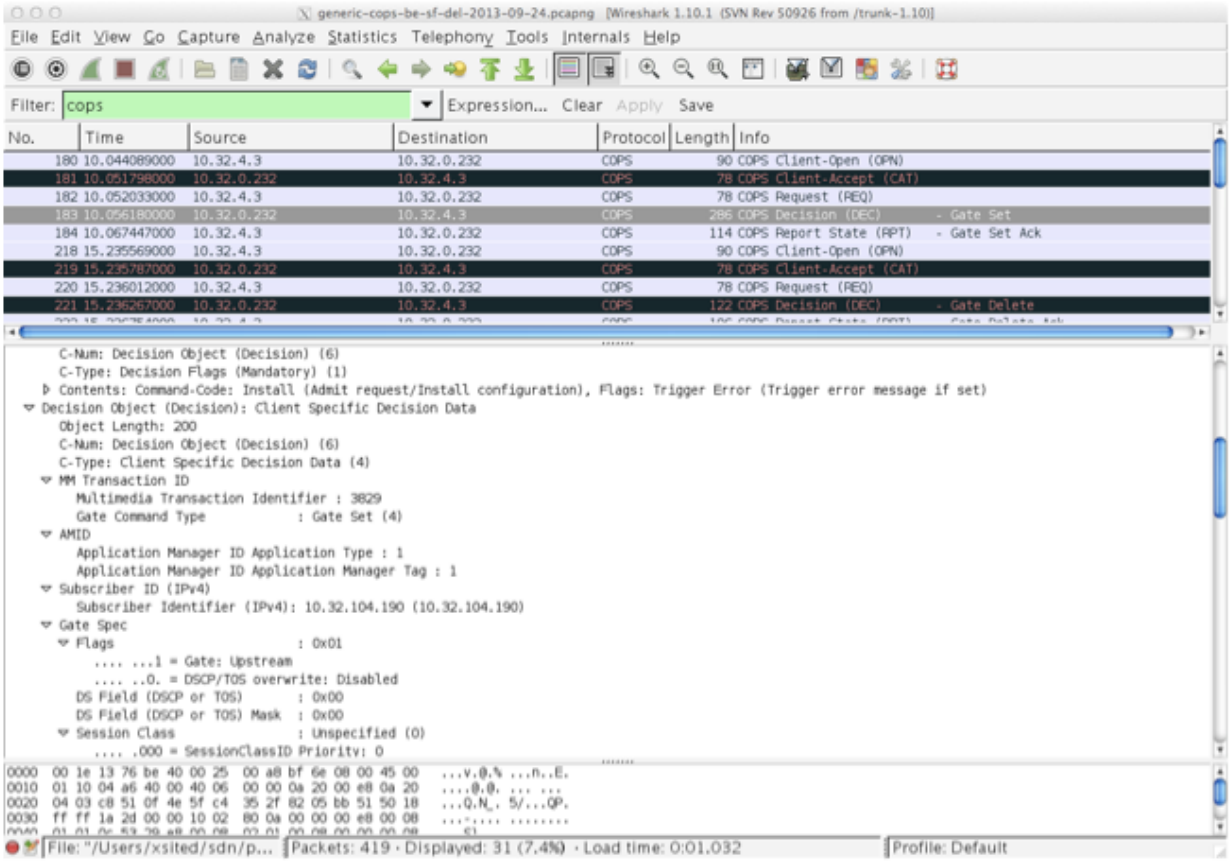

Fig. 45: Wireshark looking for COPS messages.

Show PCMM Plugin Connection

```
10k2-DSG#show packetcabl ?
  cms     Gate Controllers connected to this PacketCable client
  event   Event message server information
  gate    PacketCable gate information
  global  PacketCable global information

10k2-DSG#show packetcable cms
GC-Addr         GC-Port  Client-Addr    COPS-handle  Version PSID Key PDD-Cfg


10k2-DSG#show packetcable cms
GC-Addr         GC-Port  Client-Addr    COPS-handle  Version PSID Key PDD-Cfg
10.32.0.240     54238    10.32.15.3     0x4B9C8150/1    4.0   0    0   0
```

Show COPS Messages

```
debug cops details
```

Use CM Mac Address to List Service Flows

```
10k2-DSG#show cable modem
                                                                          D
MAC Address    IP Address      I/F           MAC           Prim RxPwr  Timing Num I
                                             State         Sid  (dBmv) Offset CPE P
0010.188a.faf6 ---             C8/0/0/UB     w-online      1    0.50   1480   1   N
74ae.7600.01f3 10.32.115.150   C8/0/10/U0    online        1    -0.50  1431   0   Y
0010.188a.fad8 10.32.115.142   C8/0/10/UB    w-online      2    -0.50  1507   1   Y
000e.0900.00dd 10.32.115.143   C8/0/10/UB    w-online      3    0.00   1677   0   Y
e86d.5271.304f 10.32.115.168   C8/0/10/UB    w-online      6    -0.50  1419   1   Y


10k2-DSG#show cable modem 000e.0900.00dd service-flow


SUMMARY:
MAC Address    IP Address      Host          MAC           Prim  Num Primary    DS
                               Interface     State         Sid   CPE Downstream RfId
000e.0900.00dd 10.32.115.143   C8/0/10/UB    w-online      3     0   Mo8/0/2:1  2353


Sfid  Dir Curr  Sid   Sched  Prio MaxSusRate  MaxBrst     MinRsvRate  Throughput
          State       Type
23    US  act   3     BE     0    0           3044        0           39
30    US  act   16    BE     0    500000      3044        0           0
24    DS  act   N/A   N/A    0    0           3044        0           17



UPSTREAM SERVICE FLOW DETAIL:

SFID  SID   Requests   Polls      Grants     Delayed    Dropped    Packets
```

(continues on next page)

(continued from previous page)

```
                                                Grants     Grants
23    3     784         0          784         0          0          784
30    16    0           0          0           0          0          0


DOWNSTREAM SERVICE FLOW DETAIL:

SFID  RP_SFID QID    Flg Policer               Scheduler              FrwdIF
                         Xmits      Drops      Xmits      Drops
24    33019   131550     0          0          777        0          Wi8/0/2:2

Flags Legend:
$: Low Latency Queue (aggregated)
~: CIR Queue
```

Deleting a PCMM Gate Message from the CMTS

```
10k2-DSG#test cable dsd  000e.0900.00dd 30
```

Find service flows

All gate controllers currently connected to the PacketCable client are displayed

```
show cable modem 00:11:22:33:44:55 service flow   ????
show cable modem
```

Debug and display PCMM Gate messages

```
debug packetcable gate control
debug packetcable gate events
show packetcable gate summary
show packetcable global
show packetcable cms
```

Debug COPS messages

```
debug cops detail
debug packetcable cops
debug cable dynamic_qos trace
```

Integration Verification

Checkout the integration project and perform regression tests.

```
git clone ssh://${ODL_USERNAME}@git.opendaylight.org:29418/integration.git
git clone https:/git.opendaylight.org/gerrit/integration.git
```

1. Check and edit the integration/features/src/main/resources/features.xml and follow the directions there.
2. Check and edit the integration/features/pom.xml and add a dependency for your feature file
3. Build integration/features and debug

```
mvn clean install
```

Test your feature in the integration/distributions/extra/karaf/ distribution

```
cd integration/distributions/extra/karaf/
mvn clean install
cd target/assembly/bin
./karaf
```

service-wrapper

Install http://karaf.apache.org/manual/latest/users-guide/wrapper.html

```
opendaylight-user@root>feature:install service-wrapper
opendaylight-user@root>wrapper:install --help
DESCRIPTION
        wrapper:install

Install the container as a system service in the OS.

SYNTAX
        wrapper:install [options]

OPTIONS
        -d, --display
                The display name of the service.
                (defaults to karaf)
        --help
                Display this help message
        -s, --start-type
                Mode in which the service is installed. AUTO_START or DEMAND_START␣
↪(Default: AUTO_START)
                (defaults to AUTO_START)
        -n, --name
                The service name that will be used when installing the service.␣
↪(Default: karaf)
                (defaults to karaf)
        -D, --description
                The description of the service.
                (defaults to )

opendaylight-user@root> wrapper:install
Creating file: /home/user/odl/distribution-karaf-0.5.0-Boron/bin/karaf-wrapper
Creating file: /home/user/odl/distribution-karaf-0.5.0-Boron/bin/karaf-service
Creating file: /home/user/odl/distribution-karaf-0.5.0-Boron/etc/karaf-wrapper.conf
Creating file: /home/user/odl/distribution-karaf-0.5.0-Boron/lib/libwrapper.so
Creating file: /home/user/odl/distribution-karaf-0.5.0-Boron/lib/karaf-wrapper.jar
Creating file: /home/user/odl/distribution-karaf-0.5.0-Boron/lib/karaf-wrapper-main.
↪jar

Setup complete.  You may wish to tweak the JVM properties in the wrapper␣
↪configuration file:
```

(continues on next page)

(continued from previous page)

```
/home/user/odl/distribution-karaf-0.5.0-Boron/etc/karaf-wrapper.conf
before installing and starting the service.


Ubuntu/Debian Linux system detected:
  To install the service:
    $ ln -s /home/user/odl/distribution-karaf-0.5.0-Boron/bin/karaf-service /etc/init.
↪d/

  To start the service when the machine is rebooted:
    $ update-rc.d karaf-service defaults

  To disable starting the service when the machine is rebooted:
    $ update-rc.d -f karaf-service remove

  To start the service:
    $ /etc/init.d/karaf-service start

  To stop the service:
    $ /etc/init.d/karaf-service stop

  To uninstall the service :
    $ rm /etc/init.d/karaf-service
```

Service Function Chaining

OpenDaylight Service Function Chaining (SFC) Overview

OpenDaylight Service Function Chaining (SFC) provides the ability to define an ordered list of a network services (e.g. firewalls, load balancers). These service are then "stitched" together in the network to create a service chain. This project provides the infrastructure (chaining logic, APIs) needed for ODL to provision a service chain in the network and an end-user application for defining such chains.

- ACE - Access Control Entry
- ACL - Access Control List
- SCF - Service Classifier Function
- SF - Service Function
- SFC - Service Function Chain
- SFF - Service Function Forwarder
- SFG - Service Function Group
- SFP - Service Function Path
- RSP - Rendered Service Path
- NSH - Network Service Header

SFC Classifier Control and Date plane Developer guide

Overview

Description of classifier can be found in: https://datatracker.ietf.org/doc/draft-ietf-sfc-architecture/

Classifier manages everything from starting the packet listener to creation (and removal) of appropriate ip(6)tables rules and marking received packets accordingly. Its functionality is **available only on Linux** as it leverdges **NetfilterQueue**, which provides access to packets matched by an **iptables** rule. Classifier requires **root privileges** to be able to operate.

So far it is capable of processing ACL for MAC addresses, ports, IPv4 and IPv6. Supported protocols are TCP and UDP.

Classifier Architecture

Python code located in the project repository sfc-py/common/classifier.py.

Note: classifier assumes that Rendered Service Path (RSP) **already exists** in ODL when an ACL referencing it is obtained

---

1. sfc_agent receives an ACL and passes it for processing to the classifier
2. the RSP (its SFF locator) referenced by ACL is requested from ODL
3. if the RSP exists in the ODL then ACL based iptables rules for it are applied

After this process is over, every packet successfully matched to an iptables rule (i.e. successfully classified) will be NSH encapsulated and forwarded to a related SFF, which knows how to traverse the RSP.

Rules are created using appropriate iptables command. If the Access Control Entry (ACE) rule is MAC address related both iptables and ip6tabeles rules re issued. If ACE rule is IPv4 address related, only iptables rules are issued, same for IPv6.

Note: iptables **raw** table contains all created rules

Information regarding already registered RSP(s) are stored in an internal data-store, which is represented as a dictionary:

```
{rsp_id: {'name': <rsp_name>,
          'chains': {'chain_name': (<ipv>,),
                     ...
                     },
          'sff': {'ip': <ip>,
                  'port': <port>,
                  'starting-index': <starting-index>,
                  'transport-type': <transport-type>
                  },
          },
...
}
```

- `name`: name of the RSP

- `chains`: dictionary of iptables chains related to the RSP with information about IP version for which the chain exists
- `SFF`: SFF forwarding parameters
  - `ip`: SFF IP address
  - `port`: SFF port
  - `starting-index`: index given to packet at first RSP hop
  - `transport-type`: encapsulation protocol

Key APIs and Interfaces

This features exposes API to configure classifier (corresponds to service-function-classifier.yang)

API Reference Documentation

See: sfc-model/src/main/yang/service-function-classifier.yang

SFC-OVS Plugin

Overview

SFC-OVS provides integration of SFC with Open vSwitch (OVS) devices. Integration is realized through mapping of SFC objects (like SF, SFF, Classifier, etc.) to OVS objects (like Bridge, TerminationPoint=Port/Interface). The mapping takes care of automatic instantiation (setup) of corresponding object whenever its counterpart is created. For example, when a new SFF is created, the SFC-OVS plugin will create a new OVS bridge and when a new OVS Bridge is created, the SFC-OVS plugin will create a new SFF.

SFC-OVS Architecture

SFC-OVS uses the OVSDB MD-SAL Southbound API for getting/writing information from/to OVS devices. The core functionality consists of two types of mapping:

1. mapping from OVS to SFC
   - OVS Bridge is mapped to SFF
   - OVS TerminationPoints are mapped to SFF DataPlane locators
2. mapping from SFC to OVS
   - SFF is mapped to OVS Bridge
   - SFF DataPlane locators are mapped to OVS TerminationPoints

Key APIs and Interfaces

- SFF to OVS mapping API (methods to convert SFF object to OVS Bridge and OVS TerminationPoints)
- OVS to SFF mapping API (methods to convert OVS Bridge and OVS TerminationPoints to SFF object)

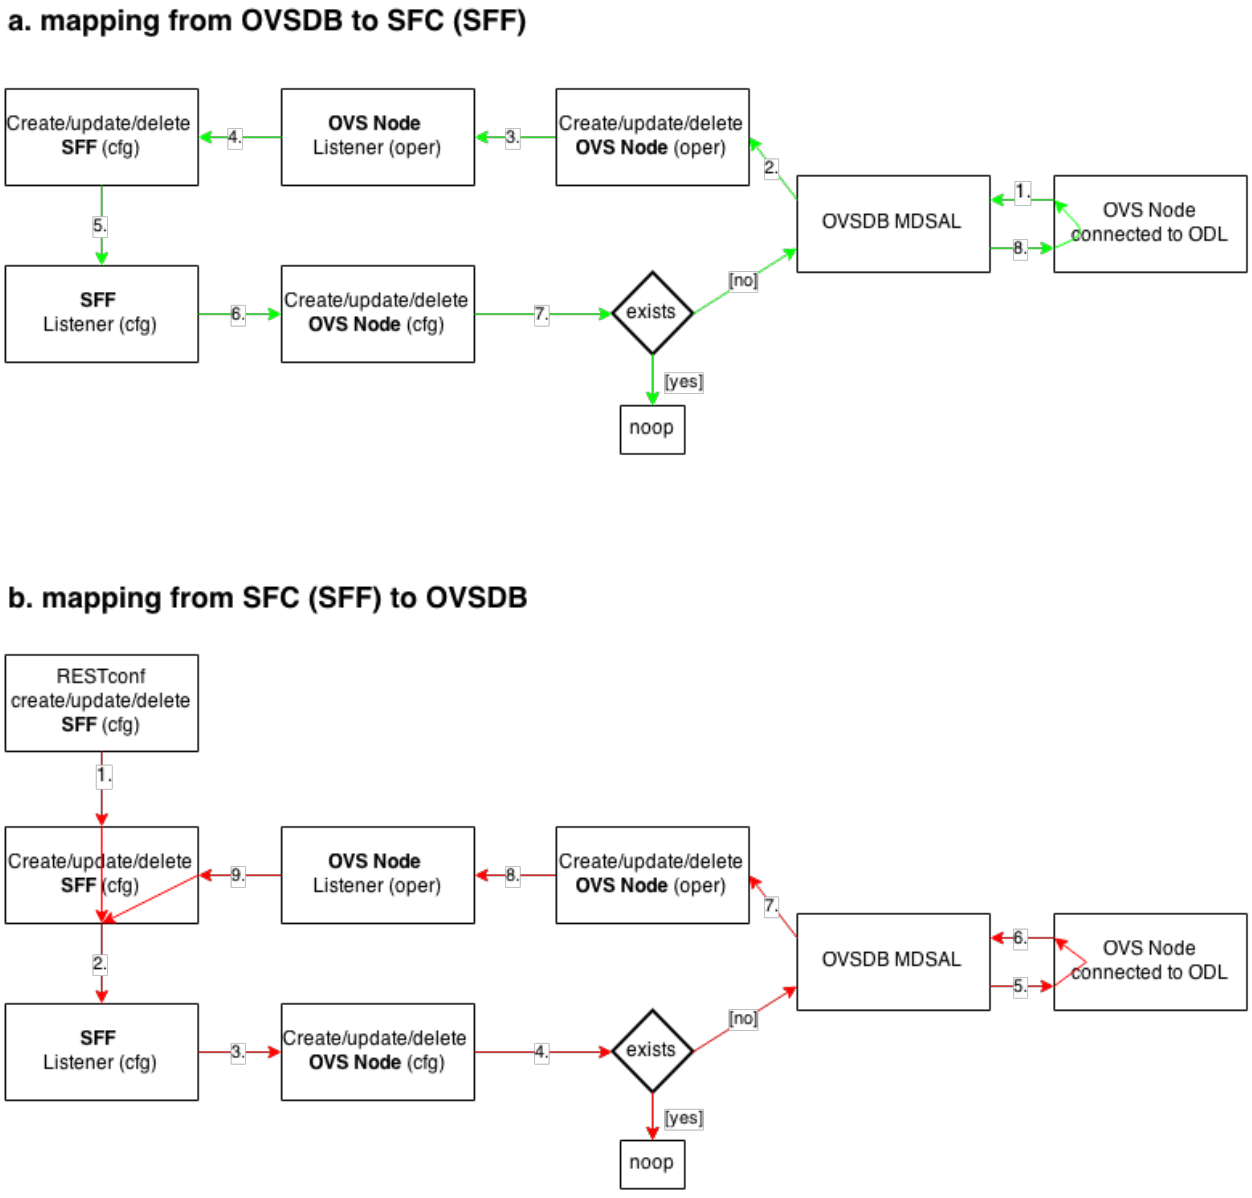

Fig. 46: SFC <—> OVS mapping flow diagram

SFC Southbound REST Plugin

Overview

The Southbound REST Plugin is used to send configuration from DataStore down to network devices supporting a REST API (i.e. they have a configured REST URI). It supports POST/PUT/DELETE operations, which are triggered accordingly by changes in the SFC data stores.

- Access Control List (ACL)
- Service Classifier Function (SCF)
- Service Function (SF)
- Service Function Group (SFG)
- Service Function Schedule Type (SFST)
- Service Function Forwader (SFF)
- Rendered Service Path (RSP)

Southbound REST Plugin Architecture

1. **listeners** - used to listen on changes in the SFC data stores
2. **JSON exporters** - used to export JSON-encoded data from binding-aware data store objects
3. **tasks** - used to collect REST URIs of network devices and to send JSON-encoded data down to these devices

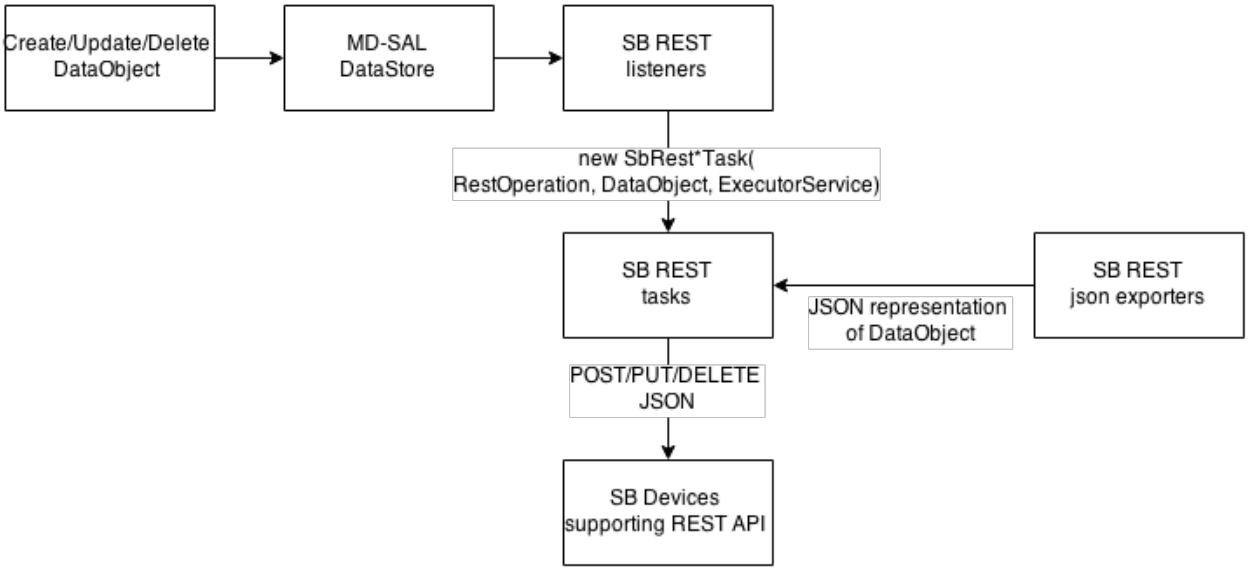

Fig. 47: Southbound REST Plugin Architecture diagram

Key APIs and Interfaces

The plugin provides Southbound REST API designated to listening REST devices. It supports POST/PUT/DELETE operations. The operation (with corresponding JSON-encoded data) is sent to unique REST URL belonging to certain datatype.

- Access Control List (ACL): `http://<host>:<port>/config/ietf-acl:access-lists/access-list/`
- Service Function (SF): `http://<host>:<port>/config/service-function:service-functions/service-function/`
- Service Function Group (SFG): `http://<host>:<port>/config/service-function:service-function-groups/service-function-group/`
- Service Function Schedule Type (SFST): `http://<host>:<port>/config/service-function-scheduler-type:service-function-scheduler-types/service-function-scheduler-type/`
- Service Function Forwarder (SFF): `http://<host>:<port>/config/service-function-forwarder:service-function-forwarders/service-function-forwarder/`
- Rendered Service Path (RSP): `http://<host>:<port>/operational/rendered-service-path:rendered-service-paths/rendered-service-path/`

Therefore, network devices willing to receive REST messages must listen on these REST URLs.

Note: Service Classifier Function (SCF) URL does not exist, because SCF is considered as one of the network devices willing to receive REST messages. However, there is a listener hooked on the SCF data store, which is triggering POST/PUT/DELETE operations of ACL object, because ACL is referenced in `service-function-classifier.yang`

Service Function Load Balancing Developer Guide

Overview

SFC Load-Balancing feature implements load balancing of Service Functions, rather than a one-to-one mapping between Service Function Forwarder and Service Function.

Load Balancing Architecture

Service Function Groups (SFG) can replace Service Functions (SF) in the Rendered Path model. A Service Path can only be defined using SFGs or SFs, but not a combination of both.

Relevant objects in the YANG model are as follows:

1. Service-Function-Group-Algorithm:

```
Service-Function-Group-Algorithms {
    Service-Function-Group-Algorithm {
        String name
        String type
    }
}
```

```
Available types: ALL, SELECT, INDIRECT, FAST_FAILURE
```

2. Service-Function-Group:

```
Service-Function-Groups {
    Service-Function-Group {
        String name
        String serviceFunctionGroupAlgorithmName
        String type
        String groupId
        Service-Function-Group-Element {
            String service-function-name
            int index
        }
    }
}
```

3. ServiceFunctionHop: holds a reference to a name of SFG (or SF)

Key APIs and Interfaces

This feature enhances the existing SFC API.

REST API commands include: * For Service Function Group (SFG): read existing SFG, write new SFG, delete existing SFG, add Service Function (SF) to SFG, and delete SF from SFG * For Service Function Group Algorithm (SFG-Alg): read, write, delete

Bundle providing the REST API: sfc-sb-rest * Service Function Groups and Algorithms are defined in: sfc-sfg and sfc-sfg-alg * Relevant JAVA API: SfcProviderServiceFunctionGroupAPI, SfcProviderServiceFunctionGroupAlgAPI

Service Function Scheduling Algorithms

Overview

When creating the Rendered Service Path (RSP), the earlier release of SFC chose the first available service function from a list of service function names. Now a new API is introduced to allow developers to develop their own schedule algorithms when creating the RSP. There are four scheduling algorithms (Random, Round Robin, Load Balance and Shortest Path) are provided as examples for the API definition. This guide gives a simple introduction of how to develop service function scheduling algorithms based on the current extensible framework.

Architecture

The following figure illustrates the service function selection framework and algorithms.

The YANG Model defines the Service Function Scheduling Algorithm type identities and how they are stored in the MD-SAL data store for the scheduling algorithms.

The MD-SAL data store stores all informations for the scheduling algorithms, including their types, names, and status.

The API provides some basic APIs to manage the informations stored in the MD-SAL data store, like putting new items into it, getting all scheduling algorithms, etc.

The RESTCONF API provides APIs to manage the informations stored in the MD-SAL data store through RESTful calls.

The Service Function Chain Renderer gets the enabled scheduling algorithm type, and schedules the service functions with scheduling algorithm implementation.

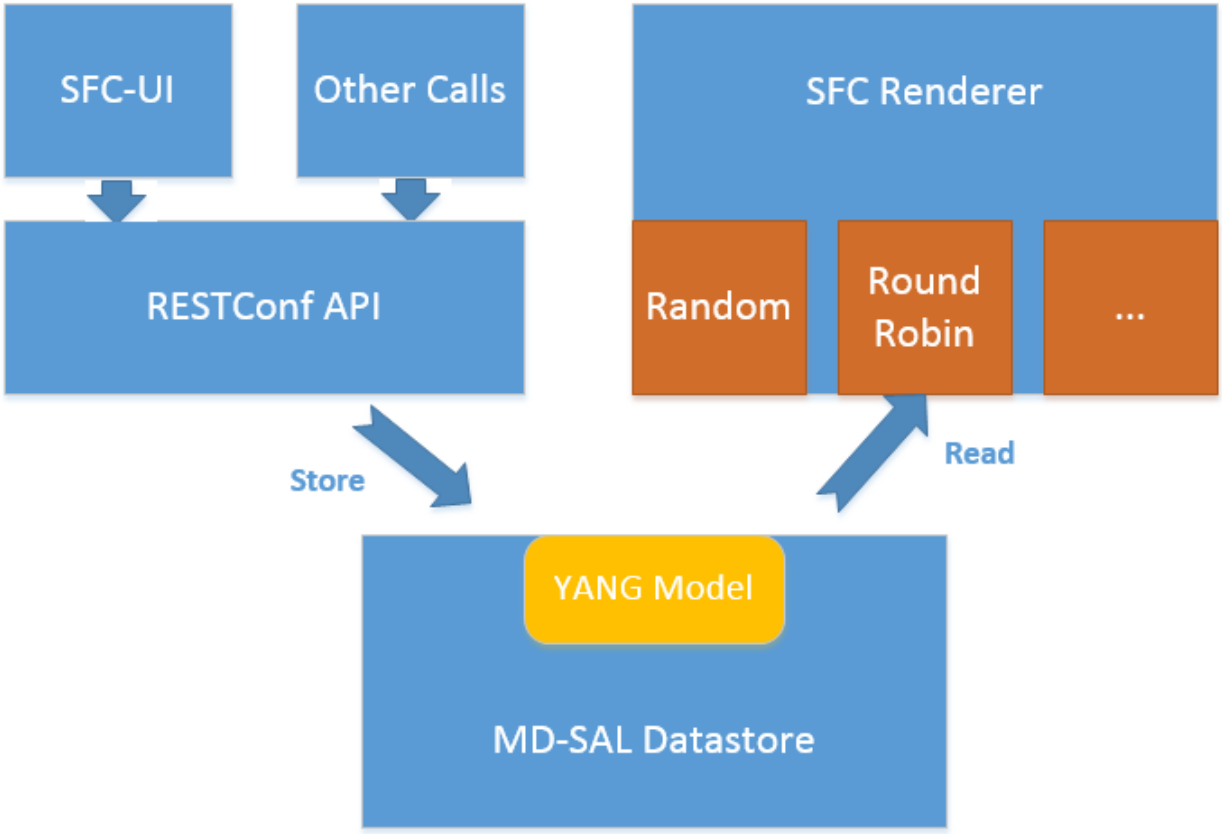

Fig. 48: SF Scheduling Algorithm framework Architecture

Key APIs and Interfaces

While developing a new Service Function Scheduling Algorithm, a new class should be added and it should extend the base schedule class SfcServiceFunctionSchedulerAPI. And the new class should implement the abstract function:

```
public List<String> scheduleServiceFuntions(ServiceFunctionChain chain, int
serviceIndex).
```

- **"ServiceFunctionChain chain"**: the chain which will be rendered
- **"int serviceIndex"**: the initial service index for this rendered service path
- **"List<String>"**: a list of service funtion names which scheduled by the Service Function Scheduling Algorithm.

API Reference Documentation

Please refer the API docs generated in the mdsal-apidocs.

SNBI Developer Guide

Overview

Key distribution in a scaled network has always been a challenge. Typically, operators must perform some manual key distribution process before secure communication is possible between a set of network devices. The Secure Network Bootstrapping Infrastructure (SNBI) project securely and automatically brings up an integrated set of network devices and controllers, simplifying the process of bootstrapping network devices with the keys required for secure communication. SNBI enables connectivity to the network devices by assigning unique IPv6 addresses and bootstrapping devices with the required keys. Admission control of devices into a specific domain is achieved using whitelist of authorized devices.

SNBI Architecture

At a high level, SNBI architecture consists of the following components:

- SNBI Registrar
- SNBI Forwarding Element (FE)

SNBI Registrar

Registrar is a device in a network that validates device against a whitelist and delivers device domain certificate. Registrar includes the following:

- RESTCONF API for Domain Whitelist Configuration
- Certificate Authority
- SNBI Southbound Plugin

RESTCONF API for Domain Whitelist Configuration:.

RESTCONF APIs are used to configure the whitelist set device in the registrar in the controller. The registrar interacts with the MD-SAL to obtain the whitelist set of devices and validate the device trying to join a domain. Furthermore it is possible to run multiple registrar instances pertaining to each domain.

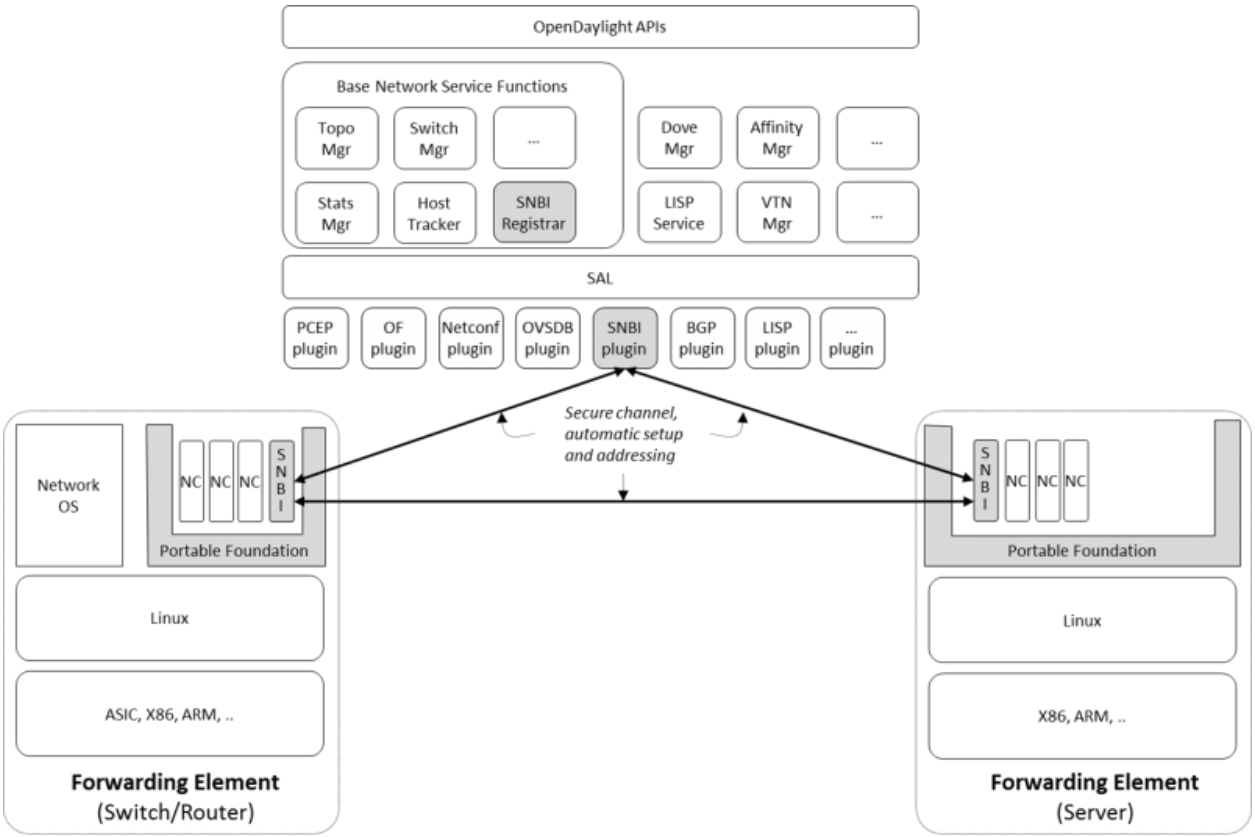

Fig. 49: SNBI Architecture Diagram

SNBI Southbound Plugin:.

The Southbound Plugin implements the protocol state machine necessary to exchange device identifiers, and deliver certificates. The southbound plugin interacts with MD-SAL and the certificate authority to validate and create device domain certificates. The device domain certificate thus generated could be used to prove the validity of the devices within the domain.

Certificate Authority:.

A simple certificate authority is implemented using the Bouncy Castle package. The Certificate Authority creates the certificates from the device CSR requests received from the devices. The certificates thus generated are delievered to the devices using the Southbound Plugin as discussed earlier.

SNBI Forwarding Element (FE)

The SNBI Forwarding Element runs on Linux machines which have to join the domain. The Device UDI(Universal Device Identifier) or the device identifier could be derived from a multitude of parameters in the host machine, but most of the parameters derived from the host are known ahead or doesn't remain constant across reloads. Therefore, each of the SNBI FE should be configured explicitly with a UDI that is already present in the device white list. The registrar service IP address must be provided to the first host (Forwarding Element) to be bootstrapped. As mentioned in the *section_title* section, the registrar service IP address is **fd08::aaaa:bbbb:1**. The First Forwarding Element must be configured with this IPv6 address.

The forwarding element must be installed or unpacked on a Linux host whose network layer traffic must be secured. The FE performs the following functions:

- Neighour Discovery
- Bootstrapping with device domain certificates
- Host Configuration

Neighbour Discovery

Neighbour Discovery (ND) is the first step in accommodating devices in a secure network. SNBI performs periodic neighbour discovery of SNBI agents by transmitting ND hello packets. The discovered devices are populated in an ND table. Neighbour Discovery is periodic and bidirectional. ND hello packets are transmitted every 10 seconds. A 40 second refresh timer is set for each discovered neighbour. On expiry of the refresh timer, the Neighbour Adjacency is removed from the ND table as the Neighbour Adjacency is no longer valid. It is possible that the same SNBI neighbour is discovered on multiple links, the expiry of a device on one link does not automatically remove the device entry from the ND table. In the exchange of ND keepalives, the device UDI is exchanged.

Bootstrapping with Device Domain Certificates

Bootstrapping a device involves the following sequential steps:

- Authenticate a device using device identifier (UDI-Universal Device Identifier or SUDI-Secure Universal Device Identifier) - The device identifier is exchanged in the hello messages.
- Allocate the appropriate device ID and IPv6 address to uniquely identify the device in the network
- Allocate the required keys by installing a Device Domain Certificate
- Accommodate the device in the domain

A device which is already bootstrapped acts as a proxy to bootstrap the new device which is trying to join the domain.

- Neighbour Invite phase - When a proxy device detects a new neighbor bootStrap connect message is initiated on behalf of the New device –**NEIGHBOUR CONNECT** Msg. The message is sent to the registrar to authenticate the device UDI against the whitelist of devices. The source IPv6 address is the proxy IPv6 address and the destination IPv6 address is the registrar IPv6 address. The SNBI Registrar provides appropriate device ID and IPv6 address to uniquely identify the device in the network and then invites the device to join the domain.—**NEIGHBOUR INVITE** Msg.
- Neighbour Reject - If the Device UDI is not in the white list of devices, then the device is rejected and is not accepted into the domain. The proxy device just updates its DB with the reject information but still maintains the Neighbour relationship.
- Neighbour BootStrap Phase - Once the new device gets a neighbour invite message, it tries to boot strap itself by generating a key pair. The device generates a Certificate Sign Request (CSR) PKCS10 request and gets it signed by the CA running at the SNBI Registrar.—**BS REQ** Msg. Once the certificate is enrolled and signed by the CA, the generated x.509 certificate is returned to the new device to complete the bootstrap process.—**BS RESP** Msg.

Host Configuration

Host configuration involves configuring a host to create a secure overlay network, assigning appropriate IPv6 address, setting up GRE tunnels, securing the tunnels traffic via IPsec and enabling connectivity via a routing protocol. Docker is used to package all the required dependent software modules.

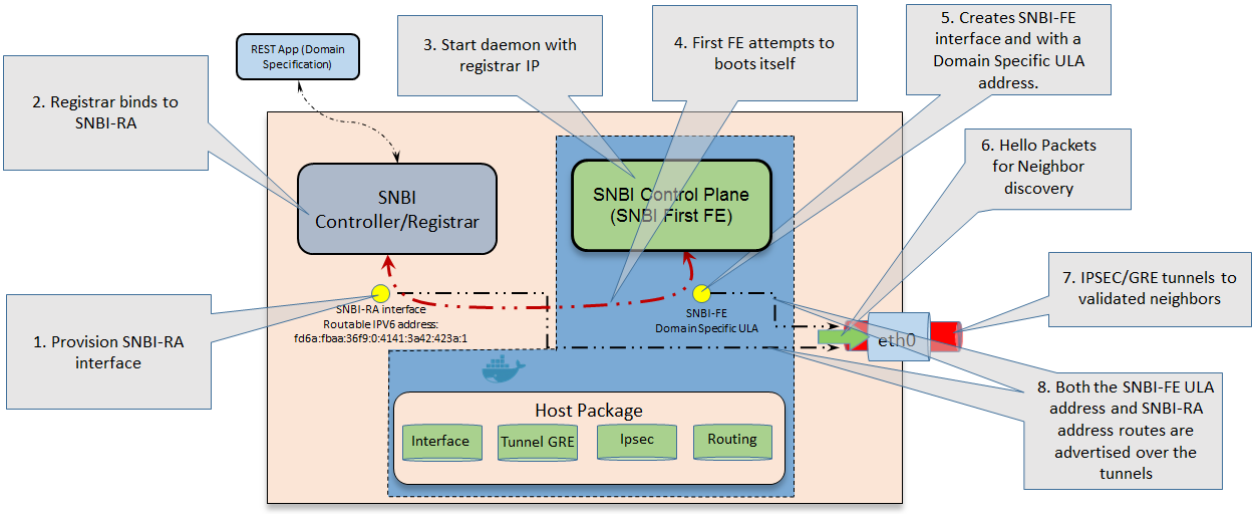

Fig. 50: SNBI Bootstrap Process

- Interace configuration: The Iproute2 package, which comes by default packaged in the Linux distributions, is used to configure the required interface (snbi-fe) and assign the appropriate IPv6 address.
- GRE Tunnel Creation: LinkLocal GRE tunnels are created to each of the discovered devices that are part of the domain. The GRE tunnels are used to create the overlay network for the domain.
- Routing over the Overlay: To enable reachability of devices within the overlay network a light weight routing protocol is used. The routing protocol of choice is the RPL (Routing Protocol for Low-Power and Lossy Networks) protocol. The routing protocol advertises the device domain IPv6 address over the overlay network. **Unstrung** is the open source implementation of RPL and is packaged within the docker image. More details on unstrung is available at http://unstrung.sandelman.ca/
- IPsec: IPsec is used to secure any traffic routed over the tunnels. StrongSwan is used to encrypt traffic using IPsec. More details on StrongSwan is available at https://www.strongswan.org/

Docker Image

The SNBI Forwarding Element is packaged in a docker container available at this link: https://hub.docker.com/r/snbi/boron/. For more information on docker, refer to this link: https://docs.docker.com/linux/.

To update an SNBI FE Daemon, build the image and copy the image to /home/snbi directory. When the docker image is run, it autoamtically generates a startup configuration file for the SNBI FE daemon. The startup configuration script is also available at /home/snbi.

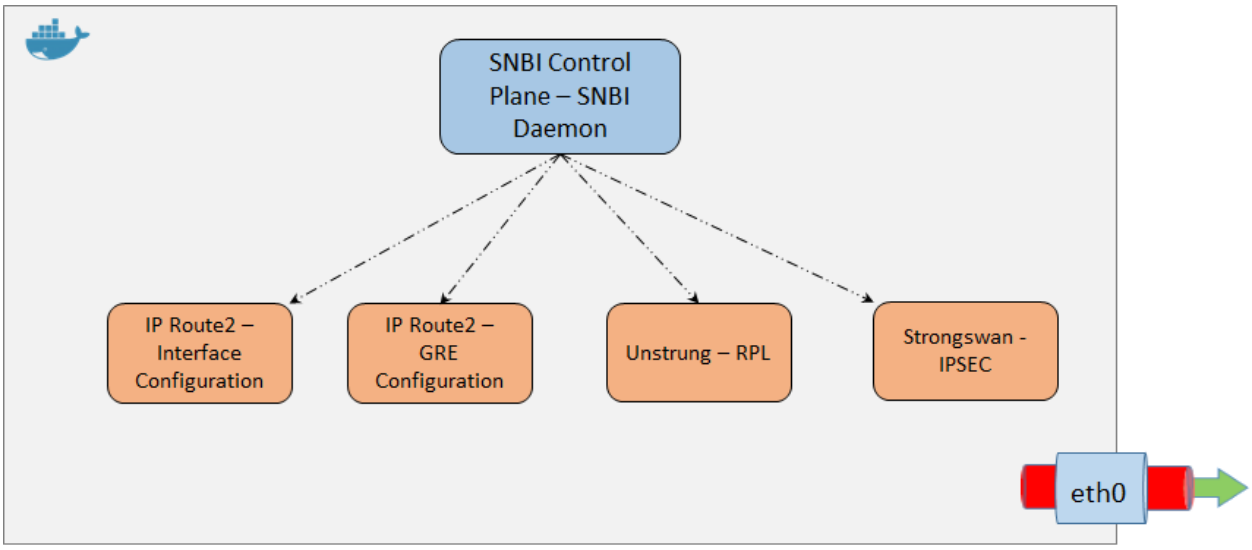

Fig. 51: SNBI Docker Image

Key APIs and Interfaces

The only API that SNBI exposes is to configure the whitelist of devices for a domain.

The POST method below configures a domain - "secure-domain" and configures a whitelist set of devices to be accommodated to the domain.

```
{
  "snbi-domain": {
    "domain-name": "secure-domain",
    "device-list": [
      {
        "list-name": "demo list",
        "list-type": "white",
        "active": true,
        "devices": [
          {
            "device-id": "UDI-FirstFE"
          },
          {
            "device-id": "UDI-dev1"
          },
          {
            "device-id": "UDI-dev2"
          }
```

(continues on next page)

(continued from previous page)

```
            ]
          }
        ]
    }
}
```

The associated device ID must be configured on the SNBI FE (see above).

API Reference Documentation

See the generated RESTCONF API documentation at: http://localhost:8181/apidoc/explorer/index.html

Look for the SNBI module to expand and see the various RESTCONF APIs.

SNMP4SDN Developer Guide

Overview

We propose a southbound plugin that can control the off-the-shelf commodity Ethernet switches for the purpose of building SDN using Ethernet switches. For Ethernet switches, forwarding table, VLAN table, and ACL are where one can install flow configuration on, and this is done via SNMP and CLI in the proposed plugin. In addition, some settings required for Ethernet switches in SDN, e.g., disabling STP and flooding, are proposed.

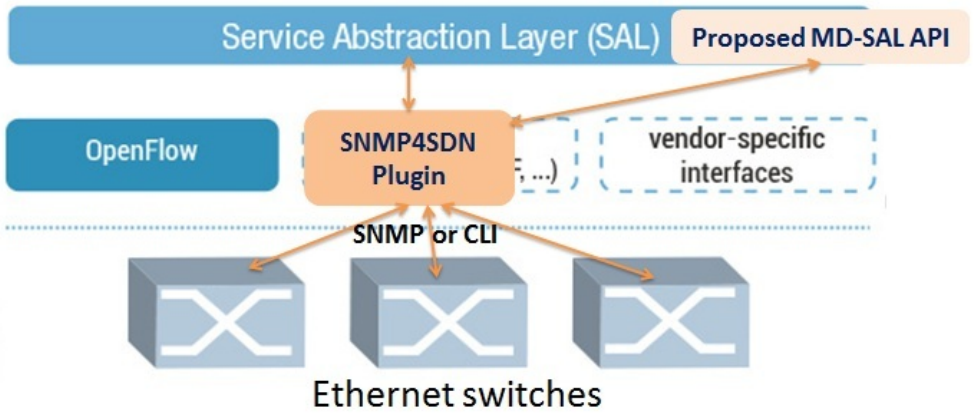

Fig. 52: SNMP4SDN as an OpenDaylight southbound plugin

Architecture

The modules in the plugin are depicted as the following figure.

- AclService: add/remove ACL profile and rule on the switches.
- FdbService: add/modify/remove FDB table entry on the switches.
- VlanService: add/modify/remove VLAN table entry on the switches.
- TopologyService: query and acquire the subnet topology.

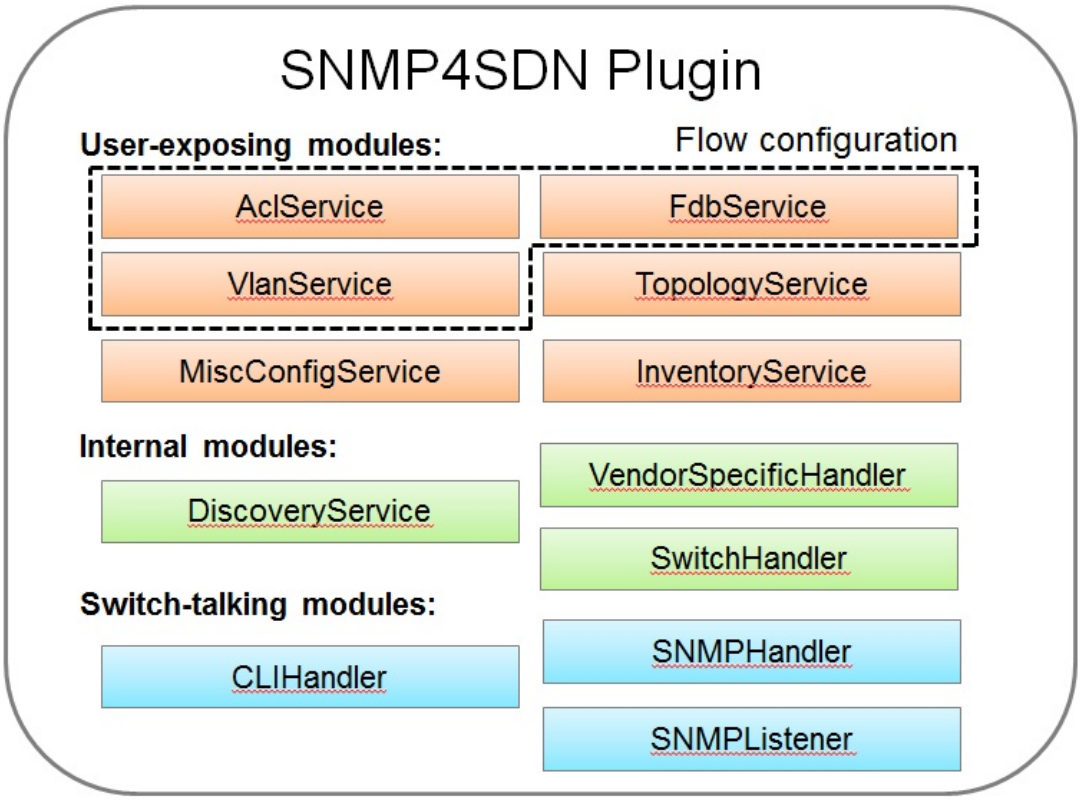

Fig. 53: Modules in the SNMP4SDN Plugin

- InventoryService: acquire the switches and their ports.
- DiscoveryService: probe and resolve the underlying switches as well as the port pairs connecting the switches. The probing is realized by SNMP queries. The updates from discovery will also be reflected to the TopologyService.
- MiscConfigService: do kinds of settings on switches
  - Supported STP and ARP settings such as enable/disable STP, get port's STP state, get ARP table, set ARP entry, and others
- VendorSpecificHandler: to assist the flow configuration services to call the switch-talking modules with correct parameters value and order.
- Switch-talking modules
  - For the services above, when they need to read or configure the underlying switches via SNMP or CLI, these queries are dealt with the modules SNMPHandler and CLIHandler which directly talk with the switches. The SNMPListener is to listen to snmp trap such as link up/down event or switch on/off event.

Design

In terms of the architecture of the SNMP4SDN Plugin's features, the features include flow configuration, topology discovery, and multi-vendor support. Their architectures please refer to Wiki (Developer Guide - Design).

Installation and Configuration Guide

- Please refer to the *Getting Started Guide* in https://www.opendaylight.org/downloads, find the SNMP4SDN section.
- For the latest full guide, please refer to Wiki (Installation Guide, User Guide - Configuration).

Tutorial

- For the latest full guide, please refer to Wiki (User Guide - Tutorial).

Programmatic Interface(s)

SNMP4SDN Plugin exposes APIs via MD-SAL with YANG model. The methods (RPC call) and data structures for them are listed below.

TopologyService

- RPC call
  - get-edge-list
  - get-node-list
  - get-node-connector-list
  - set-discovery-interval (given interval time in seconds)
  - rediscover

- Data structure
  - node: composed of node-id, node-type
  - node-connector: composed of node-connector-id, node-connector-type, node
  - topo-edge: composed of head-node-connector-id, head-node-connector-type, head-node-id, head-node-type, tail-node-connector-id, tail-node-connector-type, tail-node-id, tail-node-type

VlanService

- RPC call
  - add-vlan (given node ID, VLAN ID, VLAN name)
  - add-vlan-and-set-ports (given node ID, VLAN ID, VLAN name, tagged ports, untagged ports)
  - set-vlan-ports (given node ID, VLAN ID, tagged ports, untagged ports)
  - delete-vlan (given node ID, VLAN ID)
  - get-vlan-table (given node ID)

AclService

- RPC call
  - create-acl-profile (given node ID, acl-profile-index, acl-profile)
  - del-acl-profile (given node ID, acl-profile-index)
  - set-acl-rule (given node ID, acl-index, acl-rule)
  - del-acl-rule (given node ID, acl-index)
  - clear-acl-table (given node ID)
- Data structure
  - acl-profile-index: composed of profile-id, profile name
  - acl-profile: composed of acl-layer, vlan-mask, src-ip-mask, dst-ip-mask
  - acl-layer: IP or ETHERNET
  - acl-index: composed of acl-profile-index, acl-rule-index
  - acl-rule-index: composed of rule-id, rule-name
  - acl-rule: composed of port-list, acl-layer, acl-field, acl-action
  - acl-field: composed of vlan-id, src-ip, dst-ip
  - acl-action: PERMIT or DENY

FdbService

- RPC call
  - set-fdb-entry (given fdb-entry)
  - del-fdb-entry (given node-id, vlan-id, dest-mac-adddr)

  - get-fdb-entry (given node-id, vlan-id, dest-mac-adddr)
  - get-fdb-table (given node-id)
- Data structure
  - fdb-entry: composed of node-id, vlan-id, dest-mac-addr, port, fdb-entry-type
  - fdb-entry-type: OTHER/INVALID/LEARNED/SELF/MGMT

MiscConfigService

- RPC call
  - set-stp-port-state (given node-id, port, is_nable)
  - get-stp-port-state (given node-id, port)
  - get-stp-port-root (given node-id, port)
  - enable-stp (given node-id)
  - disable-stp (given node-id)
  - delete-arp-entry (given node-id, ip-address)
  - set-arp-entry (given node-id, arp-entry)
  - get-arp-entry (given node-id, ip-address)
  - get-arp-table (given node-id)
- Data structure
  - stp-port-state: DISABLE/BLOCKING/LISTENING/LEARNING/FORWARDING/BROKEN
  - arp-entry: composed of ip-address and mac-address

SwitchDbService

- RPC call
  - reload-db (The following 4 RPC implemention is TBD)
  - add-switch-entry
  - delete-switch-entry
  - clear-db
  - update-db
- Data structure
  - switch-info: compose of node-ip, node-mac, community, cli-user-name, cli-password, model

Help

- SNMP4SDN Wiki
- SNMP4SDN Mailing List (user, developer)
- Latest troubleshooting in Wiki

SXP Developer Guide

Overview

SXP (Source-Group Tag eXchange Protocol) project is an effort to enhance OpenDaylight platform with IP-SGT (IP Address to Source Group Tag) bindings that can be learned from connected SXP-aware network nodes. The current implementation supports SXP protocol version 4 according to the Smith, Kandula - SXP IETF draft and grouping of peers and creating filters based on ACL/Prefix-list syntax for filtering outbound and inbound IP-SGT bindings. All protocol legacy versions 1-3 are supported as well. Additionally, version 4 adds bidirectional connection type as an extension of a unidirectional one.

SXP Architecture

The SXP Server manages all connected clients in separate threads and a common SXP protocol agreement is used between connected peers. Each SXP network peer is modelled with its pertaining class, e.g., SXP Server represents the SXP Speaker, SXP Listener the Client. The server program creates the ServerSocket object on a specified port and waits until a client starts up and requests connect on the IP address and port of the server. The client program opens a Socket that is connected to the server running on the specified host IP address and port.

The SXP Listener maintains connection with its speaker peer. From an opened channel pipeline, all incoming SXP messages are processed by various handlers. Message must be decoded, parsed and validated.

The SXP Speaker is a counterpart to the SXP Listener. It maintains a connection with its listener peer and sends composed messages.

The SXP Binding Handler extracts the IP-SGT binding from a message and pulls it into the SXP-Database. If an error is detected during the IP-SGT extraction, an appropriate error code and sub-code is selected and an error message is sent back to the connected peer. All transitive messages are routed directly to the output queue of SXP Binding Dispatcher.

The SXP Binding Dispatcher represents a selector that will decides how many data from the SXP-database will be sent and when. It is responsible for message content composition based on maximum message length.

The SXP Binding Filters handles filtering of outgoing and incoming IP-SGT bindings according to BGP filtering using ACL and Prefix List syntax for specifying filter or based on Peer-sequence length.

The SXP Domains feature provides isolation of SXP peers and bindings learned between them, also exchange of Bindings is possible across SXP-Domains by ACL, Prefix List or Peer-Sequence filters

Key APIs and Interfaces

As this project is fairly small, it provides only few features that install and provide all APIs and implementations for this project.

- sxp-controller
- sxp-api
- spx-core

sxp-controller

RPC request handling

sxp-api

Contains data holders and entities

spx-core

Main logic and core features

API Reference Documentation

RESTCONF Interface and Dynamic Tree Specification and Architecture

Topology Processing Framework Developer Guide

Overview

The Topology Processing Framework allows developers to aggregate and filter topologies according to defined correlations. It also provides functionality, which you can use to make your own topology model by automating the translation from one model to another. For example to translate from the opendaylight-inventory model to only using the network-topology model.

Architecture

Chapter Overview

In this chapter we describe the architecture of the Topology Processing Framework. In the first part, we provide information about available features and basic class relationships. In the second part, we describe our model specific approach, which is used to provide support for different models.

Basic Architecture

The Topology Processing Framework consists of several Karaf features:

- odl-topoprocessing-framework
- odl-topoprocessing-inventory
- odl-topoprocessing-network-topology
- odl-topoprocessing-i2rs
- odl-topoprocessing-inventory-rendering

The feature odl-topoprocessing-framework contains the topoprocessing-api, topoprocessing-spi and topoprocessing-impl bundles. This feature is the core of the Topology Processing Framework and is required by all others features.

- topoprocessing-api - contains correlation definitions and definitions required for rendering
- topoprocessing-spi - entry point for topoprocessing service (start and close)
- topoprocessing-impl - contains base implementations of handlers, listeners, aggregators and filtrators

TopoProcessingProvider is the entry point for Topology Processing Framework. It requires a DataBroker instance. The DataBroker is needed for listener registration. There is also the TopologyRequestListener which listens on aggregated topology requests (placed into the configuration datastore) and UnderlayTopologyListeners which listen on underlay topology data changes (made in operational datastore). The TopologyRequestHandler saves toporequest data and provides a method for translating a path to the specified leaf. When a change in the topology occurs, the registered UnderlayTopologyListener processes this information for further aggregation and/or filtration. Finally, after an overlay topology is created, it is passed to the TopologyWriter, which writes this topology into operational datastore.

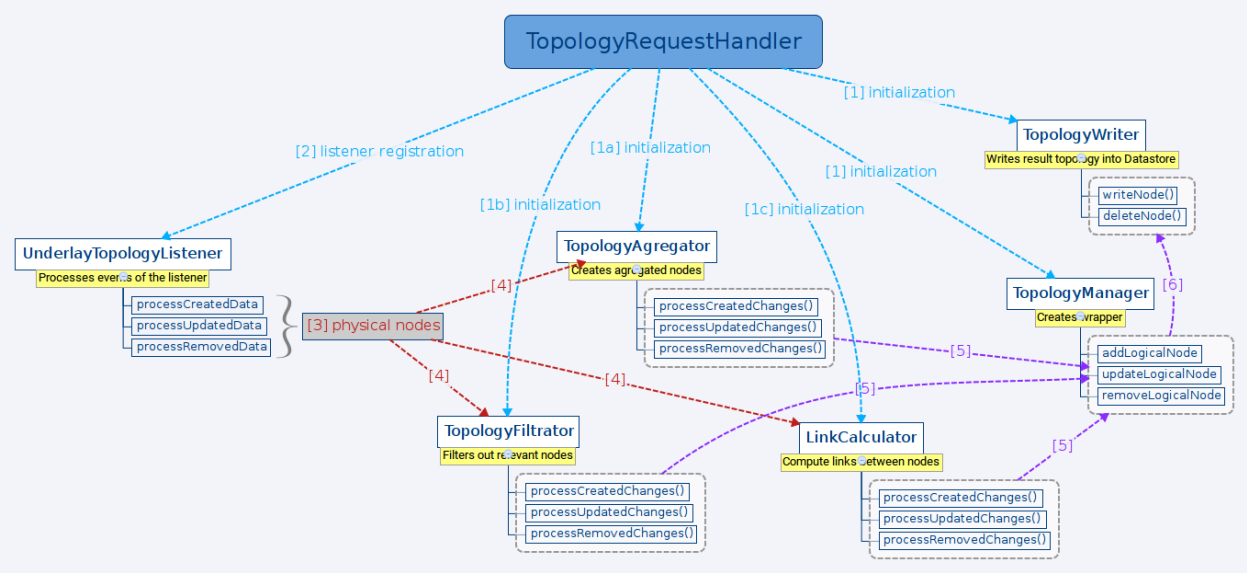

Fig. 54: Class relationship

[1] TopologyRequestHandler instantiates TopologyWriter and TopologyManager. Then, according to the request, initializes either TopologyAggregator, TopologyFiltrator or LinkCalculator.

[2] It creates as many instances of UnderlayTopologyListener as there are underlay topologies.

[3] PhysicalNodes are created for relevant incoming nodes (those having node ID).

[4a] It performs aggregation and creates logical nodes.

[4b] It performs filtration and creates logical nodes.

[4c] It performs link computation and creates links between logical nodes.

[5] Logical nodes are put into wrapper.

[6] The wrapper is translated into the appropriate format and written into datastore.

Model Specific Approach

The Topology Processing Framework consists of several modules and Karaf features, which provide support for different input models. Currently we support the network-topology, opendaylight-inventory and i2rs models. For each of these input models, the Topology Processing Framework has one module and one Karaf feature.

How it works

User point of view:

When you start the odl-topoprocessing-framework feature, the Topology Processing Framework starts without knowledge how to work with any input models. In order to allow the Topology Processing Framework to process some kind of input model, you must install one (or more) model specific features. Installing these features will also start odl-topoprocessing-framework feature if it is not already running. These features inject appropriate logic into the odl-topoprocessing-framework feature. From that point, the Topology Processing Framework is able to process different kinds of input models, specifically those that you install features for.

Developer point of view:

The topoprocessing-impl module contains (among other things) classes and interfaces, which are common for every model specific topoprocessing module. These classes and interfaces are implemented and extended by classes in particular model specific modules. Model specific modules also depend on the TopoProcessingProvider class in the topoprocessing-spi module. This dependency is injected during installation of model specific features in Karaf. When a model specific feature is started, it calls the registerAdapters(adapters) method of the injected TopoProcessingProvider object. After this step, the Topology Processing Framework is able to use registered model adapters to work with input models.

To achieve the described functionality we created a ModelAdapter interface. It represents installed feature and provides methods for creating crucial structures specific to each model.

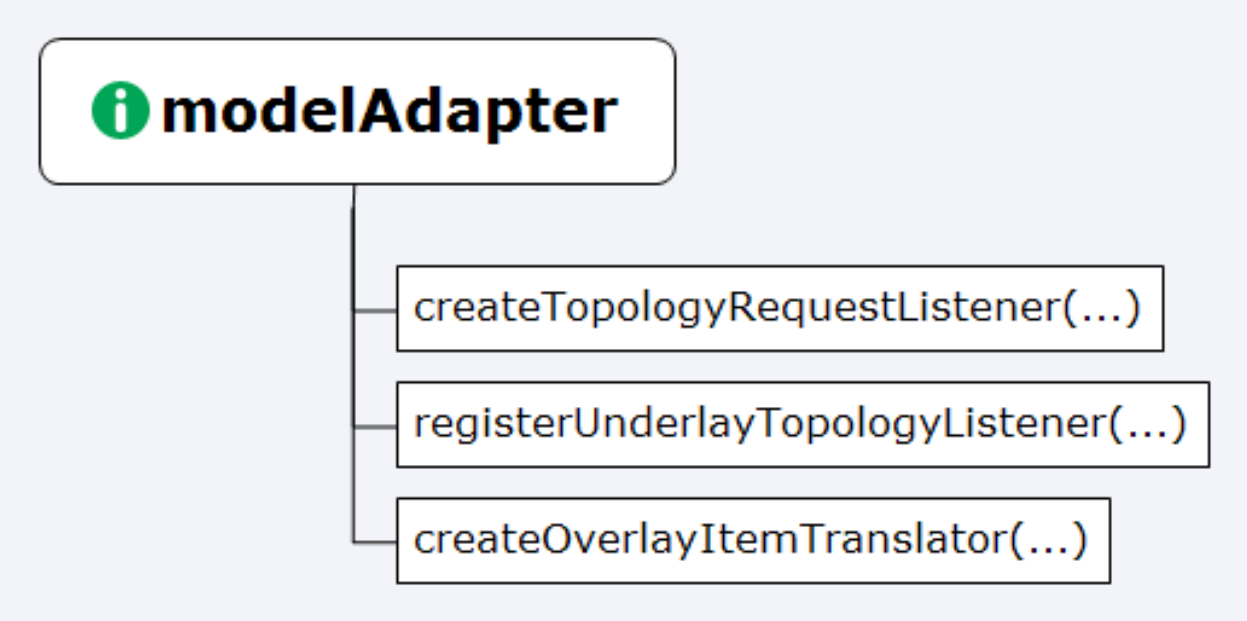

Fig. 55: ModelAdapter interface

Model Specific Features

- odl-topoprocessing-network-topology - this feature contains logic to work with network-topology model
- odl-topoprocessing-inventory - this feature contains logic to work with opendaylight-inventory model
- odl-topoprocessing-i2rs - this feature contains logic to work with i2rs model

Inventory Model Support

The opendaylight-inventory model contains only nodes, termination points, information regarding these structures. This model co-operates with network-topology model, where other topology related information is stored. This means

that we have to handle two input models at once. To support the inventory model, InventoryListener and NotificationInterConnector classes were introduced. Please see the flow diagrams below.

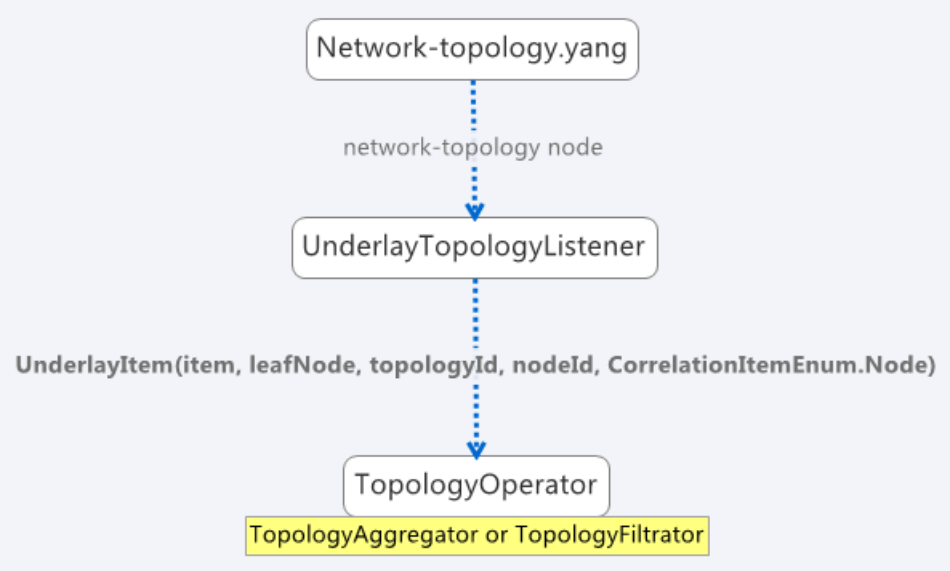

Fig. 56: Network topology model

Here we can see the InventoryListener and NotificationInterConnector classes. InventoryListener listens on data changes in the inventory model and passes these changes wrapped as an UnderlayItem for further processing to NotificationInterConnector. It doesn't contain node information - it contains a leafNode (node based on which aggregation occurs) instead. The node information is stored in the topology model, where UnderlayTopologyListener is registered as usual. This listener delivers the missing information.

Then the NotificationInterConnector combines the two notifications into a complete UnderlayItem (no null values) and delivers this UnderlayItem for further processing (to next TopologyOperator).

Aggregation and Filtration

Chapter Overview

The Topology Processing Framework allows the creation of aggregated topologies and filtered views over existing topologies. Currently, aggregation and filtration is supported for topologies that follow network-topology, opendaylight-inventory or i2rs model. When a request to create an aggregated or filtered topology is received, the framework creates one listener per underlay topology. Whenever any specified underlay topology is changed, the appropriate listener is triggered with the change and the change is processed. Two types of correlations (functionalities) are currently supported:

- Aggregation
  - Unification
  - Equality
- Filtration

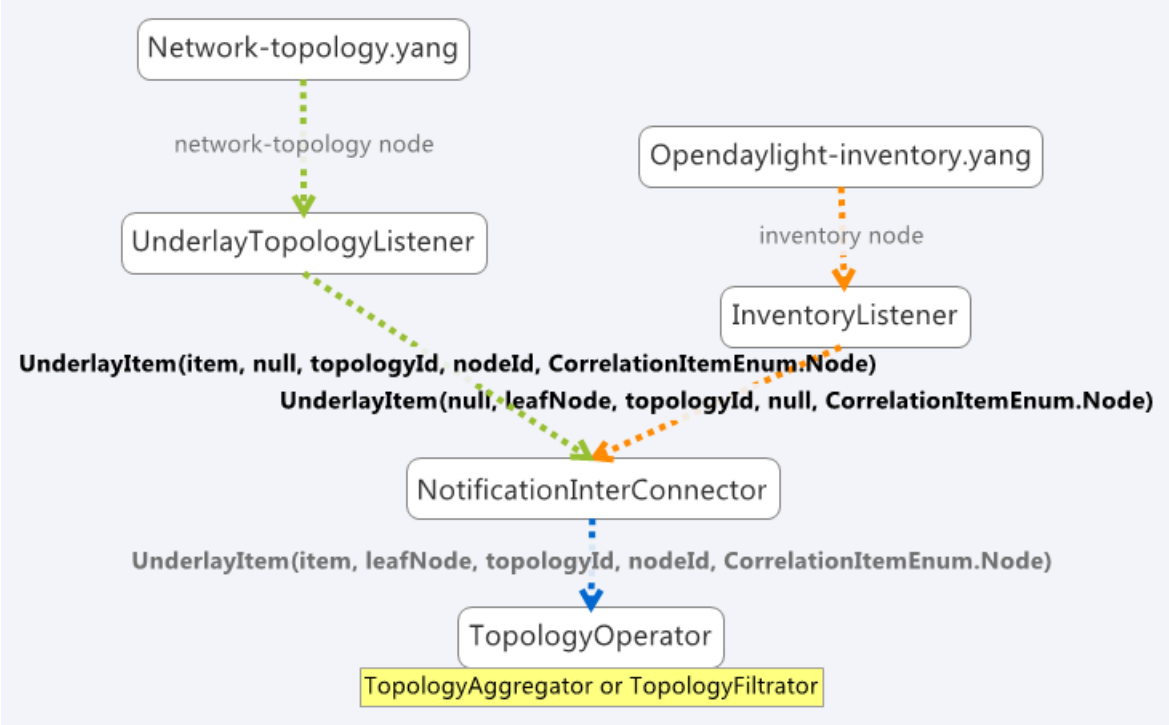

Fig. 57: Inventory model

Terminology

We use the term underlay item (physical node) for items (nodes, links, termination-points) from underlay and overlay item (logical node) for items from overlay topologies regardless of whether those are actually physical network elements.

Aggregation

Aggregation is an operation which creates an aggregated item from two or more items in the underlay topology if the aggregation condition is fulfilled. Requests for aggregated topologies must specify a list of underlay topologies over which the overlay (aggregated) topology will be created and a target field in the underlay item that the framework will check for equality.

Create Overlay Node

First, each new underlay item is inserted into the proper topology store. Once the item is stored, the framework compares it (using the target field value) with all stored underlay items from underlay topologies. If there is a target-field match, a new overlay item is created containing pointers to all *equal* underlay items. The newly created overlay item is also given new references to its supporting underlay items.

Equality case:

If an item doesn't fulfill the equality condition with any other items, processing finishes after adding the item into topology store. It will stay there for future use, ready to create an aggregated item with a new underlay item, with which it would satisfy the equality condition.

Unification case:

An overlay item is created for all underlay items, even those which don't fulfill the equality condition with any other items. This means that an overlay item is created for every underlay item, but for items which satisfy the equality condition, an aggregated item is created.

Update Node

Processing of updated underlay items depends on whether the target field has been modified. If yes, then:

- if the underlay item belonged to some overlay item, it is removed from that item. Next, if the aggregation condition on the target field is satisfied, the item is inserted into another overlay item. If the condition isn't met then:
  - in equality case - the item will not be present in overlay topology.
  - in unification case - the item will create an overlay item with a single underlay item and this will be written into overlay topology.
- if the item didn't belong to some overlay item, it is checked again for aggregation with other underlay items.

Remove Node

The underlay item is removed from the corresponding topology store, from it's overlay item (if it belongs to one) and this way it is also removed from overlay topology.

Equality case:

If there is only one underlay item left in the overlay item, the overlay item is removed.

Unification case:

The overlay item is removed once it refers to no underlay item.

Filtration

Filtration is an operation which results in creation of overlay topology containing only items fulfilling conditions set in the topoprocessing request.

Create Underlay Item

If a newly created underlay item passes all filtrators and their conditions, then it is stored in topology store and a creation notification is delivered into topology manager. No operation otherwise.

Update Underlay Item

First, the updated item is checked for presence in topology store:

- if it is present in topology store:
  - if it meets the filtering conditions, then processUpdatedData notification is triggered
  - else processRemovedData notification is triggered
- if item isn't present in topology store
  - if item meets filtering conditions, then processCreatedData notification is triggered

- else it is ignored

Remove Underlay Item

If an underlay node is supporting some overlay node, the overlay node is simply removed.

Default Filtrator Types

There are seven types of default filtrators defined in the framework:

- IPv4-address filtrator - checks if specified field meets IPv4 address + mask criteria
- IPv6-address filtrator - checks if specified field meets IPv6 address + mask criteria
- Specific number filtrator - checks for specific number
- Specific string filtrator - checks for specific string
- Range number filtrator - checks if specified field is higher than provided minimum (inclusive) and lower than provided maximum (inclusive)
- Range string filtrator - checks if specified field is alphabetically greater than provided minimum (inclusive) and alphabetically lower than provided maximum (inclusive)
- Script filtrator - allows a user or application to implement their own filtrator

Register Custom Filtrator

There might be some use case that cannot be achieved with the default filtrators. In these cases, the framework offers the possibility for a user or application to register a custom filtrator.

Pre-Filtration / Filtration & Aggregation

This feature was introduced in order to lower memory and performance demands. It is a combination of the filtration and aggregation operations. First, uninteresting items are filtered out and then aggregation is performed only on items that passed filtration. This way the framework saves on compute time. The PreAggregationFiltrator and TopologyAggregator share the same TopoStoreProvider (and thus topology store) which results in lower memory demands (as underlay items are stored only in one topology store - they aren't stored twice).

Link Computation

Chapter Overview

While processing the topology request, we create overlay nodes with lists of supporting underlay nodes. Because these overlay nodes have completely new identifiers, we lose link information. To regain this link information, we provide Link Computation functionality. Its main purpose is to create new overlay links based on the links from the underlay topologies and underlay items from overlay items. The required information for Link Computation is provided via the Link Computation model in (topology-link-computation.yang).

Link Computation Functionality

Let us consider two topologies with following components:

Topology 1:

- Node: `node:1:1`
- Node: `node:1:2`
- Node: `node:1:3`
- Link: `link:1:1` (from `node:1:1` to `node:1:2`)
- Link: `link:1:2` (from `node:1:3` to `node:1:2`)

Topology 2:

- Node: `node:2:1`
- Node: `node:2:2`
- Node: `node:2:3`
- Link: `link:2:1` (from `node:2:1` to `node:2:3`)

Now let's say that we applied some operations over these topologies that results into aggregating together

- `node:1:1` and `node:2:3` (`node:1`)
- `node:1:2` and `node:2:2` (`node:2`)
- `node:1:3` and `node:2:1` (`node:3`)

At this point we can no longer use available links in new topology because of the node ID change, so we must create new overlay links with source and destination node set to new nodes IDs. It means that `link:1:1` from topology 1 will create new link `link:1`. Since original source (`node:1:1`) is already aggregated under `node:1`, it will become source node for `link:1`. Using same method the destination will be `node:2`. And the final output will be three links:

- `link:1`, from `node:1` to `node:2`
- `link:2`, from `node:3` to `node:2`
- `link:3`, from `node:3` to `node:1`

In-Depth Look

The main logic behind Link Computation is executed in the LinkCalculator operator. The required information is passed to LinkCalculator through the LinkComputation section of the topology request. This section is defined in the topology-link-computation.yang file. The main logic also covers cases when some underlay nodes may not pass through other topology operators.

Link Computation Model

There are three essential pieces of information for link computations. All of them are provided within the LinkComputation section. These pieces are:

- output model
- overlay topology with new nodes
- underlay topologies with original links

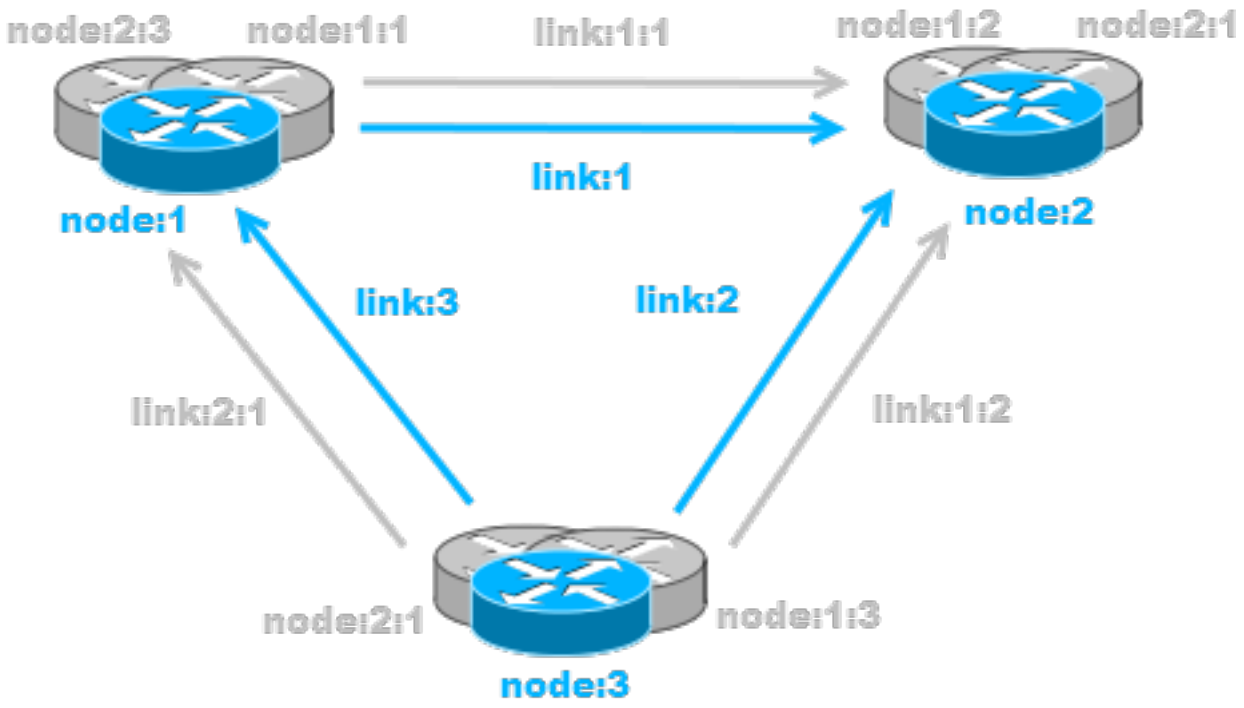

Fig. 58: Overlay topology with computed links

This whole section is augmented into `network-topology:topology`. By placing this section out of correlations section, it allows us to send link computation request separately from topology operations request.

Main Logic

Taking into consideration that some of the underlay nodes may not transform into overlay nodes (e.g. they are filtered out), we created two possible states for links:

- matched - a link is considered as matched when both original source and destination node were transformed to overlay nodes
- waiting - a link is considered as waiting if original source, destination or both nodes are missing from the overlay topology

All links in waiting the state are stored in waitingLinks list, already matched links are stored in matchedLinks list and overlay nodes are stored in the storedOverlayNodes list. All processing is based only on information in these lists. Processing created, updated and removed underlay items is slightly different and described in next sections separately.

Processing Created Items

Created items can be either nodes or links, depending on the type of listener from which they came. In the case of a link, it is immediately added to waitingLinks and calculation for possible overlay link creations (calculatePossibleLink) is started. The flow diagram for this process is shown in the following picture:

Searching for the source and destination nodes in the calculatePossibleLink method runs over each node in storedOverlayNodes and the IDs of each supporting node is compared against IDs from the underlay link's source and destination nodes. If there are any nodes missing, the link remains in the waiting state. If both the source and destination nodes are found, the corresponding overlay nodes is recorded as the new source and destination. The link is then removed from waitingLinks and a new CalculatedLink is added to the matched links. At the end, the new link (if it exists) is written into the datastore.

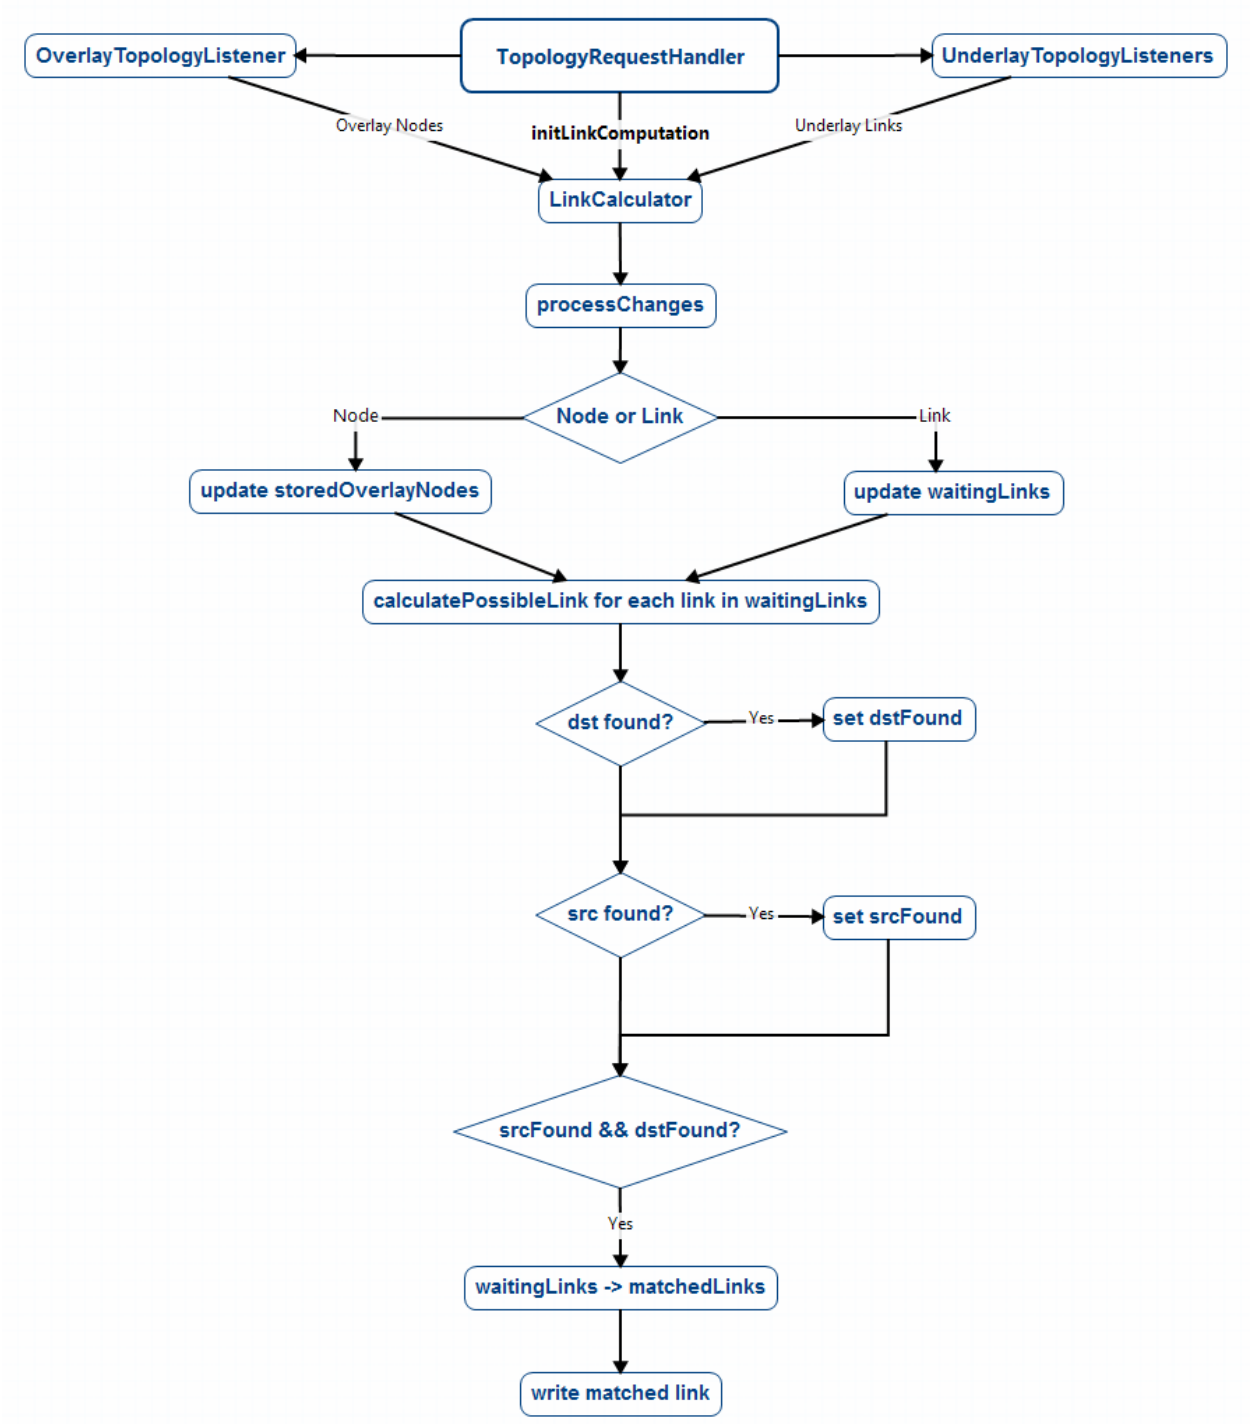

Fig. 59: Flow diagram of processing created items

If the created item is an overlayNode, this is added to storedOverlayNodes and we call calculatePossibleLink for every link in waitingLinks.

Processing Updated Items

The difference from processing created items is that we have three possible types of updated items: overlay nodes, waiting underlay links, and matched underlay links.

- In the case of a change in a matched link, this must be recalculated and based on the result it will either be matched with new source and destination or will be returned to waiting links. If the link is moved back to a waiting state, it must also be removed from the datastore.
- In the case of change in a waiting link, it is passed to the calculation process and based on the result will either remain in waiting state or be promoted to the matched state.
- In the case of a change in an overlay node, storedOverlayNodes must be updated properly and all links must be recalculated in case of changes.

Processing Removed items

Same as for processing updated item. There can be three types of removed items:

- In case of waiting link removal, the link is just removed from waitingLinks
- In case of matched link removal, the link is removed from matchingLinks and datastore
- In case of overlay node removal, the node must be removed form storedOverlayNodes and all matching links must be recalculated

Wrapper, RPC Republishing, Writing Mechanism

Chapter Overview

During the process of aggregation and filtration, overlay items (so called logical nodes) were created from underlay items (physical nodes). In the topology manager, overlay items are put into a wrapper. A wrapper is identified with unique ID and contains list of logical nodes. Wrappers are used to deal with transitivity of underlay items - which permits grouping of overlay items (into wrappers).

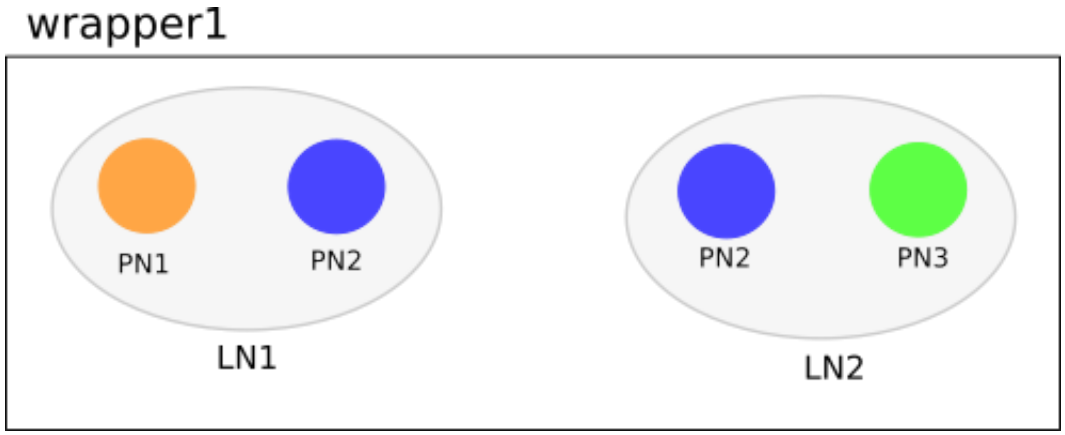

Fig. 60: Wrapper

PN1, PN2, PN3 = physical nodes

LN1, LN2 = logical nodes

RPC Republishing

All RPCs registered to handle underlay items are re-registered under their corresponding wrapper ID. RPCs of underlay items (belonging to an overlay item) are gathered, and registered under ID of their wrapper.

RPC Call

When RPC is called on overlay item, this call is delegated to it's underlay items, this means that the RPC is called on all underlay items of this overlay item.

Writing Mechanism

When a wrapper (containing overlay item(s) with it's underlay item(s)) is ready to be written into data store, it has to be converted into DOM format. After this translation is done, the result is written into datastore. Physical nodes are stored as supporting-nodes. In order to use resources responsibly, writing operation is divided into two steps. First, a set of threads registers prepared operations (deletes and puts) and one thread makes actual write operation in batch.

Topology Rendering Guide - Inventory Rendering

Chapter Overview

In the most recent OpenDaylight release, the opendaylight-inventory model is marked as deprecated. To facilitate migration from it to the network-topology model, there were requests to render (translate) data from inventory model (whether augmented or not) to another model for further processing. The Topology Processing Framework was extended to provide this functionality by implementing several rendering-specific classes. This chapter is a step-by-step guide on how to implement your own topology rendering using our inventory rendering as an example.

Use case

For the purpose of this guide we are going to render the following augmented fields from the OpenFlow model:

- from inventory node:
  - manufacturer
  - hardware
  - software
  - serial-number
  - description
  - ip-address
- from inventory node-connector:
  - name
  - hardware-address
  - current-speed
  - maximum-speed

We also want to preserve the node ID and termination-point ID from opendaylight-topology-inventory model, which is network-topology part of the inventory model.

Implementation

There are two ways to implement support for your specific topology rendering:

- add a module to your project that depends on the Topology Processing Framework
- add a module to the Topology Processing Framework itself

Regardless, a successful implementation must complete all of the following steps.

Step1 - Target Model Creation

Because the network-topology node does not have fields to store all desired data, it is necessary to create new model to render this extra data in to. For this guide we created the inventory-rendering model. The picture below shows how data will be rendered and stored.

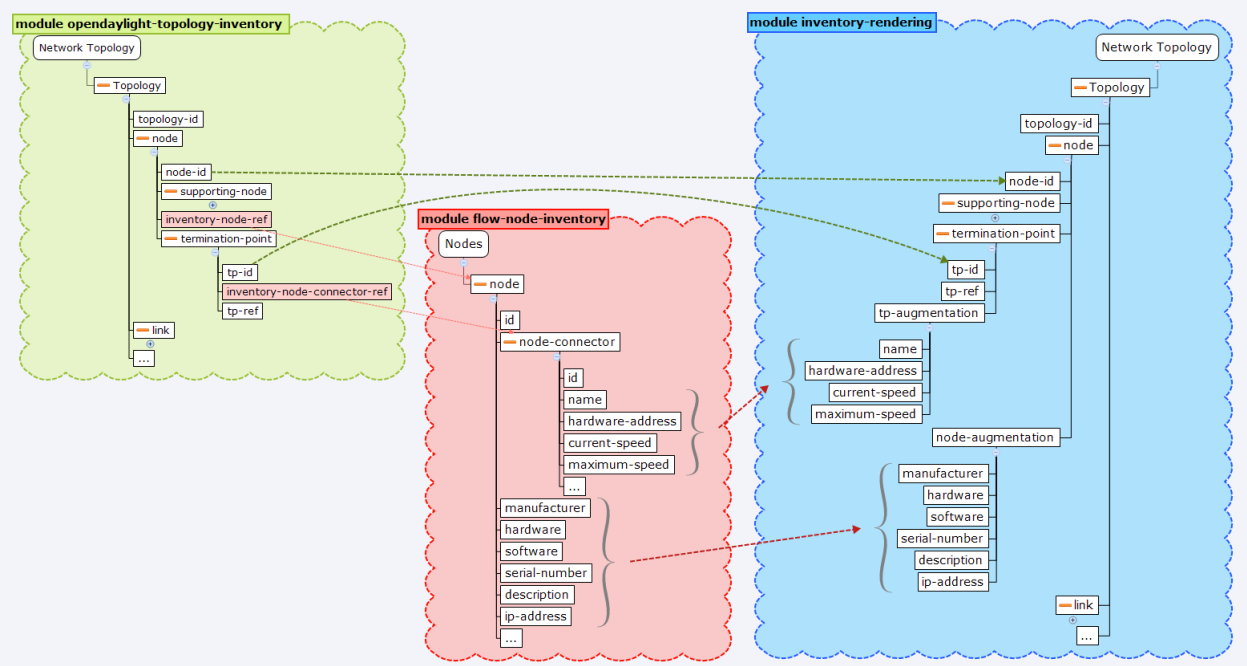

Fig. 61: Rendering to the inventory-rendering model

Important: When implementing your version of the topology-rendering model in the Topology Processing Framework, the source file of the model (.yang) must be saved in /topoprocessing-api/src/main/yang folder so corresponding structures can be generated during build and can be accessed from every module through dependencies.

When the target model is created you have to add an identifier through which you can set your new model as output model. To do that you have to add another identity item to topology-correlation.yang file. For our inventory-rendering model identity looks like this:

After that you will be able to set inventory-rendering-model as output model in XML.

Step2 - Module and Feature Creation

Important: This and following steps are based on the *model specific approach* in the Topology Processing Framework. We highly recommend that you familiarize yourself with this approach in advance.

To create a base module and add it as a feature to Karaf in the Topology Processing Framework we made the changes in following commit. Changes in other projects will likely be similar.

File	Changes
pom.xml	add new module to topoprocessing
features.xml	add feature to topoprocessing
features/pom.xml	add dependencies needed by features
topoprocessing-artifacts/pom.xml	add artifact
topoprocessing-config/pom.xml	add configuration file
81-topoprocessing-inventory-renderin g-config.xml	configuration file for new module
topoprocessing-inventory-rendering/p om.xml	main pom for new module
TopoProcessingProviderIR.java	contains startup method which register new model adapter
TopoProcessingProviderIRMod-ule.java	generated class which contains createInstance method. You should call your startup method from here.
TopoProcessingProviderIRModule-Factor y.java	generated class. You will probably not need to edit this file
log4j.xml	configuration file for logger topoprocessing-inventory-rendering-p rovider-impl.yang

Step3 - Module Adapters Creation

There are seven mandatory interfaces or abstract classes that needs to be implemented in each module. They are:

- TopoProcessingProvider - provides module registration
- ModelAdapter - provides model specific instances
- TopologyRequestListener - listens on changes in the configuration datastore
- TopologyRequestHandler - processes configuration datastore changes
- UnderlayTopologyListener - listens for changes in the specific model
- LinkTransaltor and NodeTranslator - used by OverlayItemTranslator to create NormalizedNodes from OverlayItems

The name convention we used was to add an abbreviation for the specific model to the beginning of implementing class name (e.g. the IRModelAdapter refers to class which implements ModelAdapter in module Inventory Rendering). In the case of the provider class, we put the abbreviation at the end.

Important:

- In the next sections, we use the terms TopologyRequestListener, TopologyRequestHandler, etc. without a prepended or appended abbreviation because the steps apply regardless of which specific model you are targeting.

- If you want to implement rendering from inventory to network-topology, you can just copy-paste our module and additional changes will be required only in the output part.

Provider part

This part is the starting point of the whole module. It is responsible for creating and registering TopologyRequestListeners. It is necessary to create three classes which will import:

- **TopoProcessingProviderModule** - is a generated class from topoprocessing-inventory-rendering-provider-impl.yang (created in previous step, file will appear after first build). Its method `createInstance()` is called at the feature start and must be modified to create an instance of TopoProcessingProvider and call its `startup(TopoProcessingProvider topoProvider)` function.
- **TopoProcessingProvider** - in `startup(TopoProcessingProvider topoProvider)` function provides ModelAdapter registration to TopoProcessingProviderImpl.
- **ModelAdapter** - provides creation of corresponding module specific classes.

Input part

This includes the creation of the classes responsible for input data processing. In this case, we had to create five classes implementing:

- **TopologyRequestListener** and **TopologyRequestHandler** - when notified about a change in the configuration datastore, verify if the change contains a topology request (has correlations in it) and creates UnderlayTopologyListeners if needed. The implementation of these classes will differ according to the model in which are correlations saved (network-topology or i2rs). In the case of using network-topology, as the input model, you can use our classes IRTopologyRequestListener and IRTopologyRequestHandler.
- **UnderlayTopologyListener** - registers underlay listeners according to input model. In our case (listening in the inventory model), we created listeners for the network-topology model and inventory model, and set the NotificationInterConnector as the first operator and set the IRRenderingOperator as the second operator (after NotificationInterConnector). Same as for TopologyRequestListener/Handler, if you are rendering from the inventory model, you can use our class IRUnderlayTopologyListener.
- **InventoryListener** - a new implementation of this class is required only for inventory input model. This is because the InventoryListener from topoprocessing-impl requires pathIdentifier which is absent in the case of rendering.
- **TopologyOperator** - replaces classic topoprocessing operator. While the classic operator provides specific operations on topology, the rendering operator just wraps each received UnderlayItem to OverlayItem and sends them to write.

Important: For purposes of topology rendering from inventory to network-topology, there are misused fields in UnderlayItem as follows:

- item - contains node from network-topology part of inventory
- leafItem - contains node from inventory

In case of implementing UnderlayTopologyListener or InventoryListener you have to carefully adjust UnderlayItem creation to these terms.

Output part

The output part of topology rendering is responsible for translating received overlay items to normalized nodes. In the case of inventory rendering, this is where node information from inventory are combined with node information from network-topology. This combined information is stored in our inventory-rendering model normalized node and passed to the writer.

The output part consists of two translators implementing the NodeTranslator and LinkTranslator interfaces.

NodeTranslator implementation - The NodeTranslator interface has one `translate(OverlayItemWrapper wrapper)` method. For our purposes, there is one important thing in wrapper - the list of OverlayItems which have one or more common UnderlayItems. Regardless of this list, in the case of rendering it will always contains only one OverlayItem. This item has list of UnderlayItems, but again in case of rendering there will be only one UnderlayItem item in this list. In NodeTranslator, the OverlayItem and corresponding UnderlayItem represent nodes from the translating model.

The UnderlayItem has several attributes. How you will use these attributes in your rendering is up to you, as you create this item in your topology operator. For example, as mentioned above, in our inventory rendering example is an inventory node normalized node stored in the UnderlayItem leafNode attribute, and we also store node-id from network-topology model in UnderlayItem itemId attribute. You can now use these attributes to build a normalized node for your new model. How to read and create normalized nodes is out of scope of this document.

LinkTranslator implementation - The LinkTranslator interface also has one `translate(OverlayItemWrapper wrapper)` method. In our inventory rendering this method returns `null`, because the inventory model doesn't have links. But if you also need links, this is the place where you should translate it into a normalized node for your model. In LinkTranslator, the OverlayItem and corresponding UnderlayItem represent links from the translating model. As in NodeTranslator, there will be only one OverlayItem and one UnderlayItem in the corresponding lists.

Testing

If you want to test topoprocessing with some manually created underlay topologies (like in this guide), than you have to tell Topoprocessing to listen for underlay topologies on Configuration datastore instead of Operational.

You can do this in this config file
`<topoprocessing_directory>/topoprocessing-config/src/main/resources/80-topoprocessing-config.xml`.
Here you have to change
`<datastore-type>OPERATIONAL</datastore-type>`
to
`<datastore-type>CONFIGURATION</datastore-type>`.

Also you have to add dependency required to test "inventory" topologies.

In `<topoprocessing_directory>/features/pom.xml`
add `<openflowplugin.version>latest_snapshot</openflowplugin.version>` to properties section
and add this dependency to dependencies section

```
<dependency>
        <groupId>org.opendaylight.openflowplugin</groupId>
        <artifactId>features-openflowplugin</artifactId>
        <version>${openflowplugin.version}</version>
        <classifier>features</classifier><type>xml</type>
</dependency>
```

`latest_snapshot` in `<openflowplugin.version>` replace with latest snapshot, which can be found here.

And in `<topoprocessing_directory>/features/src/main/resources/features.xml` add `<repository>mvn:org.opendaylight.openflowplugin/features-openflowplugin/${openflowplugin.version}/xml/features</repository>` to repositories section.

Now after you rebuild project and start Karaf, you can install necessary features.

You can install all with one command:

```
feature:install odl-restconf-noauth odl-topoprocessing-inventory-rendering
odl-openflowplugin-southbound odl-openflowplugin-nsf-model
```

Now you can send messages to REST from any REST client (e.g. Postman in Chrome). Messages have to have following headers:

Header	Value
Content-Type:	application/xml
Accept:	application/xml
username:	admin
password:	admin

Firstly send topology request to http://localhost:8181/restconf/config/network-topology:network-topology/topology/render:1 with method PUT. Example of simple rendering request:

```
<topology xmlns="urn:TBD:params:xml:ns:yang:network-topology">
  <topology-id>render:1</topology-id>
    <correlations xmlns="urn:opendaylight:topology:correlation" >
      <output-model>inventory-rendering-model</output-model>
      <correlation>
         <correlation-id>1</correlation-id>
          <type>rendering-only</type>
          <correlation-item>node</correlation-item>
          <rendering>
            <underlay-topology>und-topo:1</underlay-topology>
        </rendering>
      </correlation>
    </correlations>
</topology>
```

This request says that we want create topology with name render:1 and this topology should be stored in the inventory-rendering-model and it should be created from topology flow:1 by node rendering.

Next we send the network-topology part of topology flow:1. So to the URL http://localhost:8181/restconf/config/network-topology:network-topology/topology/und-topo:1 we PUT:

```
<topology xmlns="urn:TBD:params:xml:ns:yang:network-topology"
          xmlns:it="urn:opendaylight:model:topology:inventory"
          xmlns:i="urn:opendaylight:inventory">
    <topology-id>und-topo:1</topology-id>
    <node>
        <node-id>openflow:1</node-id>
        <it:inventory-node-ref>
    /i:nodes/i:node[i:id="openflow:1"]
```

(continues on next page)

(continued from previous page)

```
        </it:inventory-node-ref>
        <termination-point>
            <tp-id>tp:1</tp-id>
            <it:inventory-node-connector-ref>
                /i:nodes/i:node[i:id="openflow:1"]/i:node-connector[i:id="openflow:1:1
↪"]
            </it:inventory-node-connector-ref>
        </termination-point>
    </node>
</topology>
```

And the last input will be inventory part of topology. To the URL http://localhost:8181/restconf/config/opendaylight-inventory:nodes we PUT:

```
<nodes
    xmlns="urn:opendaylight:inventory">
    <node>
        <id>openflow:1</id>
        <node-connector>
            <id>openflow:1:1</id>
            <port-number
                xmlns="urn:opendaylight:flow:inventory">1
            </port-number>
            <current-speed
                xmlns="urn:opendaylight:flow:inventory">10000000
            </current-speed>
            <name
                xmlns="urn:opendaylight:flow:inventory">s1-eth1
            </name>
            <supported
                xmlns="urn:opendaylight:flow:inventory">
            </supported>
            <current-feature
                xmlns="urn:opendaylight:flow:inventory">copper ten-gb-fd
            </current-feature>
            <configuration
                xmlns="urn:opendaylight:flow:inventory">
            </configuration>
            <peer-features
                xmlns="urn:opendaylight:flow:inventory">
            </peer-features>
            <maximum-speed
                xmlns="urn:opendaylight:flow:inventory">0
            </maximum-speed>
            <advertised-features
                xmlns="urn:opendaylight:flow:inventory">
            </advertised-features>
            <hardware-address
                xmlns="urn:opendaylight:flow:inventory">0E:DC:8C:63:EC:D1
            </hardware-address>
            <state
                xmlns="urn:opendaylight:flow:inventory">
                <link-down>false</link-down>
                <blocked>false</blocked>
                <live>false</live>
            </state>
```

(continues on next page)

(continued from previous page)

```
            <flow-capable-node-connector-statistics
                xmlns="urn:opendaylight:port:statistics">
                <receive-errors>0</receive-errors>
                <receive-frame-error>0</receive-frame-error>
                <receive-over-run-error>0</receive-over-run-error>
                <receive-crc-error>0</receive-crc-error>
                <bytes>
                    <transmitted>595</transmitted>
                    <received>378</received>
                </bytes>
                <receive-drops>0</receive-drops>
                <duration>
                    <second>28</second>
                    <nanosecond>410000000</nanosecond>
                </duration>
                <transmit-errors>0</transmit-errors>
                <collision-count>0</collision-count>
                <packets>
                    <transmitted>7</transmitted>
                    <received>5</received>
                </packets>
                <transmit-drops>0</transmit-drops>
            </flow-capable-node-connector-statistics>
        </node-connector>
        <node-connector>
            <id>openflow:1:LOCAL</id>
            <port-number
                xmlns="urn:opendaylight:flow:inventory">4294967294
            </port-number>
            <current-speed
                xmlns="urn:opendaylight:flow:inventory">0
            </current-speed>
            <name
                xmlns="urn:opendaylight:flow:inventory">s1
            </name>
            <supported
                xmlns="urn:opendaylight:flow:inventory">
            </supported>
            <current-feature
                xmlns="urn:opendaylight:flow:inventory">
            </current-feature>
            <configuration
                xmlns="urn:opendaylight:flow:inventory">
            </configuration>
            <peer-features
                xmlns="urn:opendaylight:flow:inventory">
            </peer-features>
            <maximum-speed
                xmlns="urn:opendaylight:flow:inventory">0
            </maximum-speed>
            <advertised-features
                xmlns="urn:opendaylight:flow:inventory">
            </advertised-features>
            <hardware-address
                xmlns="urn:opendaylight:flow:inventory">BA:63:87:0C:76:41
            </hardware-address>
            <state
```

(continues on next page)

(continued from previous page)

```
                xmlns="urn:opendaylight:flow:inventory">
                <link-down>false</link-down>
                <blocked>false</blocked>
                <live>false</live>
            </state>
            <flow-capable-node-connector-statistics
                xmlns="urn:opendaylight:port:statistics">
                <receive-errors>0</receive-errors>
                <receive-frame-error>0</receive-frame-error>
                <receive-over-run-error>0</receive-over-run-error>
                <receive-crc-error>0</receive-crc-error>
                <bytes>
                    <transmitted>576</transmitted>
                    <received>468</received>
                </bytes>
                <receive-drops>0</receive-drops>
                <duration>
                    <second>28</second>
                    <nanosecond>426000000</nanosecond>
                </duration>
                <transmit-errors>0</transmit-errors>
                <collision-count>0</collision-count>
                <packets>
                    <transmitted>6</transmitted>
                    <received>6</received>
                </packets>
                <transmit-drops>0</transmit-drops>
            </flow-capable-node-connector-statistics>
        </node-connector>
        <serial-number
            xmlns="urn:opendaylight:flow:inventory">None
        </serial-number>
        <manufacturer
            xmlns="urn:opendaylight:flow:inventory">Nicira, Inc.
        </manufacturer>
        <hardware
            xmlns="urn:opendaylight:flow:inventory">Open vSwitch
        </hardware>
        <software
            xmlns="urn:opendaylight:flow:inventory">2.1.3
        </software>
        <description
            xmlns="urn:opendaylight:flow:inventory">None
        </description>
        <ip-address
            xmlns="urn:opendaylight:flow:inventory">10.20.30.40
    </ip-address>
        <meter-features
            xmlns="urn:opendaylight:meter:statistics">
            <max_bands>0</max_bands>
            <max_color>0</max_color>
            <max_meter>0</max_meter>
        </meter-features>
        <group-features
            xmlns="urn:opendaylight:group:statistics">
            <group-capabilities-supported
                xmlns:x="urn:opendaylight:group:types">x:chaining
```

(continues on next page)

(continued from previous page)

```
            </group-capabilities-supported>
            <group-capabilities-supported
                xmlns:x="urn:opendaylight:group:types">x:select-weight
            </group-capabilities-supported>
            <group-capabilities-supported
                xmlns:x="urn:opendaylight:group:types">x:select-liveness
            </group-capabilities-supported>
            <max-groups>4294967040</max-groups>
            <actions>67082241</actions>
            <actions>0</actions>
        </group-features>
    </node>
</nodes>
```

After this, the expected result from a GET request to http://127.0.0.1:8181/restconf/operational/network-topology:network-topology is:

```
<network-topology
    xmlns="urn:TBD:params:xml:ns:yang:network-topology">
    <topology>
        <topology-id>render:1</topology-id>
        <node>
            <node-id>openflow:1</node-id>
            <node-augmentation
                xmlns="urn:opendaylight:topology:inventory:rendering">
                <ip-address>10.20.30.40</ip-address>
                <serial-number>None</serial-number>
                <manufacturer>Nicira, Inc.</manufacturer>
                <description>None</description>
                <hardware>Open vSwitch</hardware>
                <software>2.1.3</software>
            </node-augmentation>
            <termination-point>
                <tp-id>openflow:1:1</tp-id>
                <tp-augmentation
                    xmlns="urn:opendaylight:topology:inventory:rendering">
                    <hardware-address>0E:DC:8C:63:EC:D1</hardware-address>
                    <current-speed>10000000</current-speed>
                    <maximum-speed>0</maximum-speed>
                    <name>s1-eth1</name>
                </tp-augmentation>
            </termination-point>
            <termination-point>
                <tp-id>openflow:1:LOCAL</tp-id>
                <tp-augmentation
                    xmlns="urn:opendaylight:topology:inventory:rendering">
                    <hardware-address>BA:63:87:0C:76:41</hardware-address>
                    <current-speed>0</current-speed>
                    <maximum-speed>0</maximum-speed>
                    <name>s1</name>
                </tp-augmentation>
            </termination-point>
        </node>
    </topology>
</network-topology>
```

Use Cases

You can find use case examples on this wiki page.

Key APIs and Interfaces

The basic provider class is TopoProcessingProvider which provides startup and shutdown methods. Otherwise, the framework communicates via requests and outputs stored in the MD-SAL datastores.

API Reference Documentation

You can find API examples on this wiki page.

TTP Model Developer Guide

Overview

Table Type Patterns are a specification developed by the Open Networking Foundation to enable the description and negotiation of subsets of the OpenFlow protocol. This is particularly useful for hardware switches that support OpenFlow as it enables the to describe what features they do (and thus also what features they do not) support. More details can be found in the full specification listed on the OpenFlow specifications page.

TTP Model Architecture

The TTP Model provides a YANG-modeled type for a TTP and allows them to be associated with a master list of known TTPs, as well as active and supported TTPs with nodes in the MD-SAL inventory model.

Key APIs and Interfaces

The key API provided by the TTP Model feature is the ability to store a set of TTPs in the MD-SAL as well as associate zero or one active TTPs and zero or more supported TTPs along with a given node in the MD-SAL inventory model.

API Reference Documentation

RESTCONF

See the generated RESTCONF API documentation at: http://localhost:8181/apidoc/explorer/index.html

Look for the onf-ttp module to expand and see the various RESTCONF APIs.

Java Bindings

As stated above there are 3 locations where a Table Type Pattern can be placed into the MD-SAL Data Store. They correspond to 3 different REST API URIs:

1. `restconf/config/onf-ttp:opendaylight-ttps/onf-ttp:table-type-patterns/`

2. `restconf/config/opendaylight-inventory:nodes/node/{id}/ttp-inventory-node:active_ttp/`
3. `restconf/config/opendaylight-inventory:nodes/node/{id}/ttp-inventory-node:supported_ttps/`

Note: Typically, these URIs are running on the machine the controller is on at port 8181. If you are on the same machine they can thus be accessed at `http://localhost:8181/<uri>`

Using the TTP Model RESTCONF APIs

Setting REST HTTP Headers

Authentication

The REST API calls require authentication by default. The default method is to use basic auth with a user name and password of 'admin'.

Content-Type and Accept

RESTCONF supports both xml and json. This example focuses on JSON, but xml can be used just as easily. When doing a PUT or POST be sure to specify the appropriate `Conetnt-Type` header: either `application/json` or `application/xml`.

When doing a GET be sure to specify the appropriate `Accept` header: again, either `application/json` or `application/xml`.

Content

The contents of a PUT or POST should be a OpenDaylight Table Type Pattern. An example of one is provided below. The example can also be found at parser/sample-TTP-from-tests.ttp in the TTP git repository.

Sample Table Type Pattern (json).

```
{
    "table-type-patterns": {
        "table-type-pattern": [
            {
                "security": {
                    "doc": [
                        "This TTP is not published for use by ONF. It is an example␣
↪and for",
                        "illustrative purposes only.",
                        "If this TTP were published for use it would include",
                        "guidance as to any security considerations in this doc␣
↪member."
                    ]
                },
                "NDM_metadata": {
                    "authority": "org.opennetworking.fawg",
                    "OF_protocol_version": "1.3.3",
```

(continues on next page)

(continued from previous page)

```
                    "version": "1.0.0",
                    "type": "TTPv1",
                    "doc": [
                        "Example of a TTP supporting L2 (unicast, multicast,␣
↪flooding), L3 (unicast only),",
                        "and an ACL table."
                    ],
                    "name": "L2-L3-ACLs"
                },
                "identifiers": [
                    {
                        "doc": [
                            "The VLAN ID of a locally attached L2 subnet on a Router."
                        ],
                        "var": "<subnet_VID>"
                    },
                    {
                        "doc": [
                            "An OpenFlow group identifier (integer) identifying a␣
↪group table entry",
                            "of the type indicated by the variable name"
                        ],
                        "var": "<<group_entry_types/name>>"
                    }
                ],
                "features": [
                    {
                        "doc": [
                            "Flow entry notification Extension - notification of␣
↪changes in flow entries"
                        ],
                        "feature": "ext187"
                    },
                    {
                        "doc": [
                            "Group notifications Extension - notification of changes␣
↪in group or meter entries"
                        ],
                        "feature": "ext235"
                    }
                ],
                "meter_table": {
                    "meter_types": [
                        {
                            "name": "ControllerMeterType",
                            "bands": [
                                {
                                    "type": "DROP",
                                    "rate": "1000..10000",
                                    "burst": "50..200"
                                }
                            ]
                        },
                        {
                            "name": "TrafficMeter",
                            "bands": [
                                {
```

(continues on next page)

(continued from previous page)

```
                                "type": "DSCP_REMARK",
                                "rate": "10000..500000",
                                "burst": "50..500"
                            },
                            {
                                "type": "DROP",
                                "rate": "10000..500000",
                                "burst": "50..500"
                            }
                        ]
                    }
                ],
                "built_in_meters": [
                    {
                        "name": "ControllerMeter",
                        "meter_id": 1,
                        "type": "ControllerMeterType",
                        "bands": [
                            {
                                "rate": 2000,
                                "burst": 75
                            }
                        ]
                    },
                    {
                        "name": "AllArpMeter",
                        "meter_id": 2,
                        "type": "ControllerMeterType",
                        "bands": [
                            {
                                "rate": 1000,
                                "burst": 50
                            }
                        ]
                    }
                ]
            },
            "table_map": [
                {
                    "name": "ControlFrame",
                    "number": 0
                },
                {
                    "name": "IngressVLAN",
                    "number": 10
                },
                {
                    "name": "MacLearning",
                    "number": 20
                },
                {
                    "name": "ACL",
                    "number": 30
                },
                {
                    "name": "L2",
                    "number": 40
```

(continues on next page)

(continued from previous page)


```
                    },
                    {
                        "name": "ProtoFilter",
                        "number": 50
                    },
                    {
                        "name": "IPv4",
                        "number": 60
                    },
                    {
                        "name": "IPv6",
                        "number": 80
                    }
                ],
                "parameters": [
                    {
                        "doc": [
                            "documentation"
                        ],
                        "name": "Showing-curt-how-this-works",
                        "type": "type1"
                    }
                ],
                "flow_tables": [
                    {
                        "doc": [
                            "Filters L2 control reserved destination addresses and",
                            "may forward control packets to the controller.",
                            "Directs all other packets to the Ingress VLAN table."
                        ],
                        "name": "ControlFrame",
                        "flow_mod_types": [
                            {
                                "doc": [
                                    "This match/action pair allows for flow_mods that␣
↪match on either",
                                    "ETH_TYPE or ETH_DST (or both) and send the␣
↪packet to the",
                                    "controller, subject to metering."
                                ],
                                "name": "Frame-To-Controller",
                                "match_set": [
                                    {
                                        "field": "ETH_TYPE",
                                        "match_type": "all_or_exact"
                                    },
                                    {
                                        "field": "ETH_DST",
                                        "match_type": "exact"
                                    }
                                ],
                                "instruction_set": [
                                    {
                                        "doc": [
                                            "This meter may be used to limit the rate␣
↪of PACKET_IN frames",
                                            "sent to the controller"
```


(continues on next page)

(continued from previous page)

```
                                        ],
                                        "instruction": "METER",
                                        "meter_name": "ControllerMeter"
                                    },
                                    {
                                        "instruction": "APPLY_ACTIONS",
                                        "actions": [
                                            {
                                                "action": "OUTPUT",
                                                "port": "CONTROLLER"
                                            }
                                        ]
                                    }
                                ]
                            }
                        ],
                        "built_in_flow_mods": [
                            {
                                "doc": [
                                    "Mandatory filtering of control frames with C-
↪VLAN Bridge reserved DA."
                                ],
                                "name": "Control-Frame-Filter",
                                "priority": "1",
                                "match_set": [
                                    {
                                        "field": "ETH_DST",
                                        "mask": "0xfffffffffff0",
                                        "value": "0x0180C2000000"
                                    }
                                ]
                            },
                            {
                                "doc": [
                                    "Mandatory miss flow_mod, sends packets to␣
↪IngressVLAN table."
                                ],
                                "name": "Non-Control-Frame",
                                "priority": "0",
                                "instruction_set": [
                                    {
                                        "instruction": "GOTO_TABLE",
                                        "table": "IngressVLAN"
                                    }
                                ]
                            }
                        ]
                    }
                ],
                "group_entry_types": [
                    {
                        "doc": [
                            "Output to a port, removing VLAN tag if needed.",
                            "Entry per port, plus entry per untagged VID per port."
                        ],
                        "name": "EgressPort",
                        "group_type": "INDIRECT",
```

(continues on next page)

(continued from previous page)

```
                        "bucket_types": [
                            {
                                "name": "OutputTagged",
                                "action_set": [
                                    {
                                        "action": "OUTPUT",
                                        "port": "<port_no>"
                                    }
                                ]
                            },
                            {
                                "name": "OutputUntagged",
                                "action_set": [
                                    {
                                        "action": "POP_VLAN"
                                    },
                                    {
                                        "action": "OUTPUT",
                                        "port": "<port_no>"
                                    }
                                ]
                            },
                            {
                                "opt_tag": "VID-X",
                                "name": "OutputVIDTranslate",
                                "action_set": [
                                    {
                                        "action": "SET_FIELD",
                                        "field": "VLAN_VID",
                                        "value": "<local_vid>"
                                    },
                                    {
                                        "action": "OUTPUT",
                                        "port": "<port_no>"
                                    }
                                ]
                            }
                        ]
                    }
                ],
                "flow_paths": [
                    {
                        "doc": [
                            "This object contains just a few examples of flow paths,␣
↪it is not",
                            "a comprehensive list of the flow paths required for this␣
↪TTP.  It is",
                            "intended that the flow paths array could include either␣
↪a list of",
                            "required flow paths or a list of specific flow paths␣
↪that are not",
                            "required (whichever is more concise or more useful."
                        ],
                        "name": "L2-2",
                        "path": [
                            "Non-Control-Frame",
                            "IV-pass",
```

(continues on next page)

(continued from previous page)

```
                "Known-MAC",
                "ACLskip",
                "L2-Unicast",
                "EgressPort"
            ]
        },
        {
            "name": "L2-3",
            "path": [
                "Non-Control-Frame",
                "IV-pass",
                "Known-MAC",
                "ACLskip",
                "L2-Multicast",
                "L2Mcast",
                "[EgressPort]"
            ]
        },
        {
            "name": "L2-4",
            "path": [
                "Non-Control-Frame",
                "IV-pass",
                "Known-MAC",
                "ACL-skip",
                "VID-flood",
                "VIDflood",
                "[EgressPort]"
            ]
        },
        {
            "name": "L2-5",
            "path": [
                "Non-Control-Frame",
                "IV-pass",
                "Known-MAC",
                "ACLskip",
                "L2-Drop"
            ]
        },
        {
            "name": "v4-1",
            "path": [
                "Non-Control-Frame",
                "IV-pass",
                "Known-MAC",
                "ACLskip",
                "L2-Router-MAC",
                "IPv4",
                "v4-Unicast",
                "NextHop",
                "EgressPort"
            ]
        },
        {
            "name": "v4-2",
            "path": [
```

(continues on next page)

(continued from previous page)

```
                                "Non-Control-Frame",
                                "IV-pass",
                                "Known-MAC",
                                "ACLskip",
                                "L2-Router-MAC",
                                "IPv4",
                                "v4-Unicast-ECMP",
                                "L3ECMP",
                                "NextHop",
                                "EgressPort"
                            ]
                        }
                    ]
                }
            ]
        }
}
```

Making a REST Call

In this example we'll do a PUT to install the sample TTP from above into OpenDaylight and then retrieve it both as json and as xml. We'll use the Postman - REST Client for Chrome in the examples, but any method of accessing REST should work.

First, we'll fill in the basic information:

1. Set the URL to `http://localhost:8181/restconf/config/onf-ttp:opendaylight-ttps/onf-ttp:table-type-patterns/`
2. Set the action to `PUT`
3. Click Headers and
4. Set a header for `Content-Type` to `application/json`
5. Make sure the content is set to raw and
6. Copy the sample TTP from above into the content
7. Click the Basic Auth tab and
8. Set the username and password to admin
9. Click Refresh headers

After clicking Refresh headers, we can see that a new header (`Authorization`) has been created and this will allow us to authenticate to make the REST call.

At this point, clicking send should result in a Status response of `200 OK` indicating we've successfully PUT the TTP into OpenDaylight.

We can now retrieve the TTP by:

1. Changing the action to `GET`
2. Setting an `Accept` header to `application/json` and
3. Pressing send

The same process can retrieve the content as xml by setting the `Accept` header to `application/xml`.

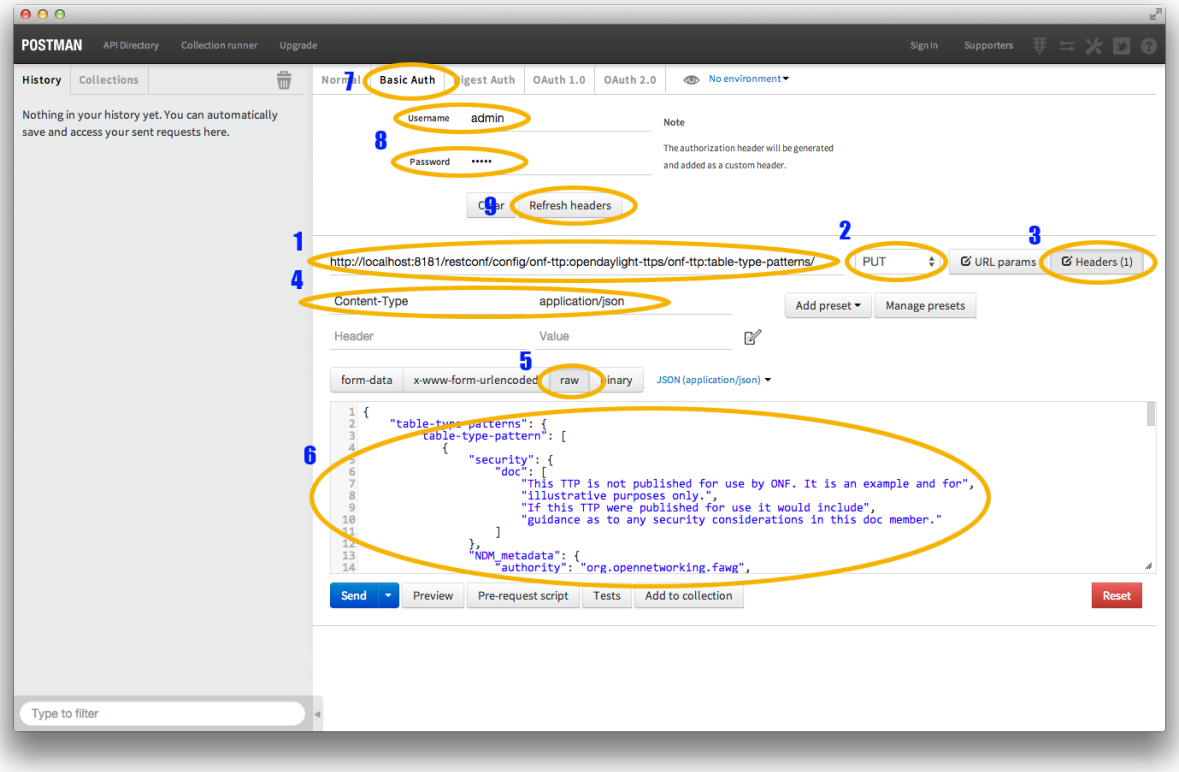

Fig. 62: Filling in URL, content, Content-Type and basic auth

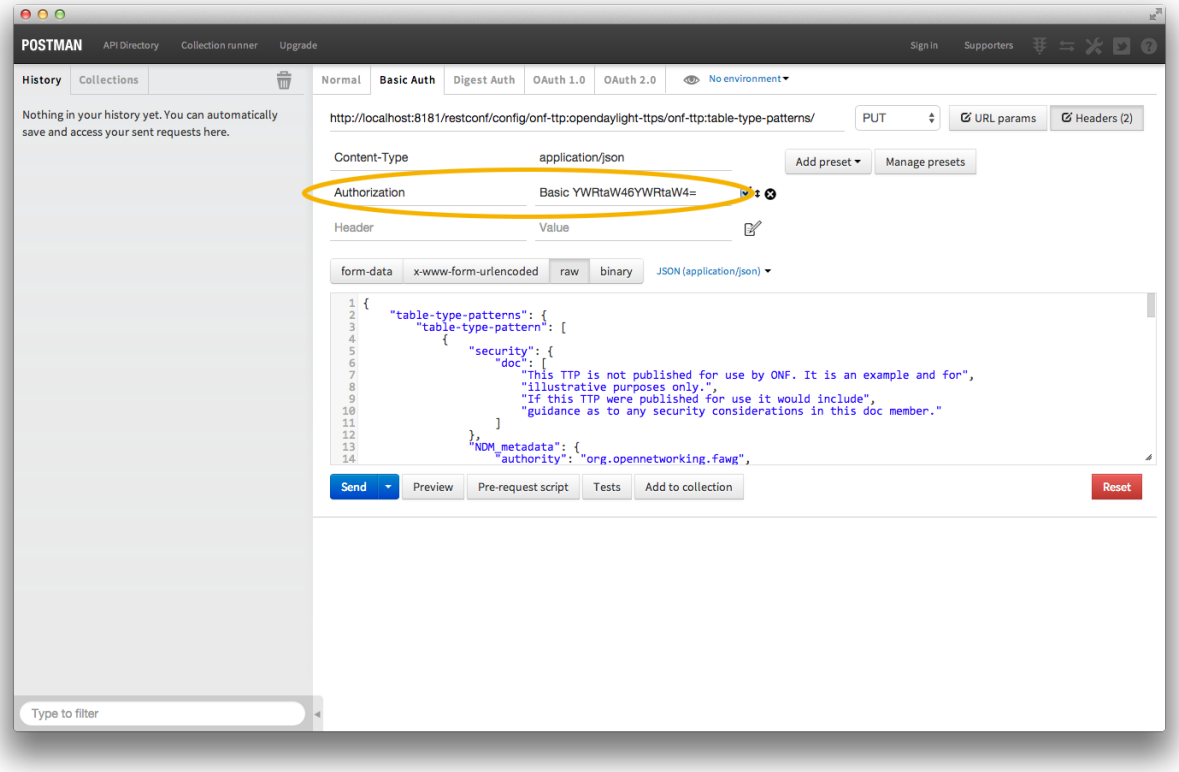

Fig. 63: Refreshing basic auth headers

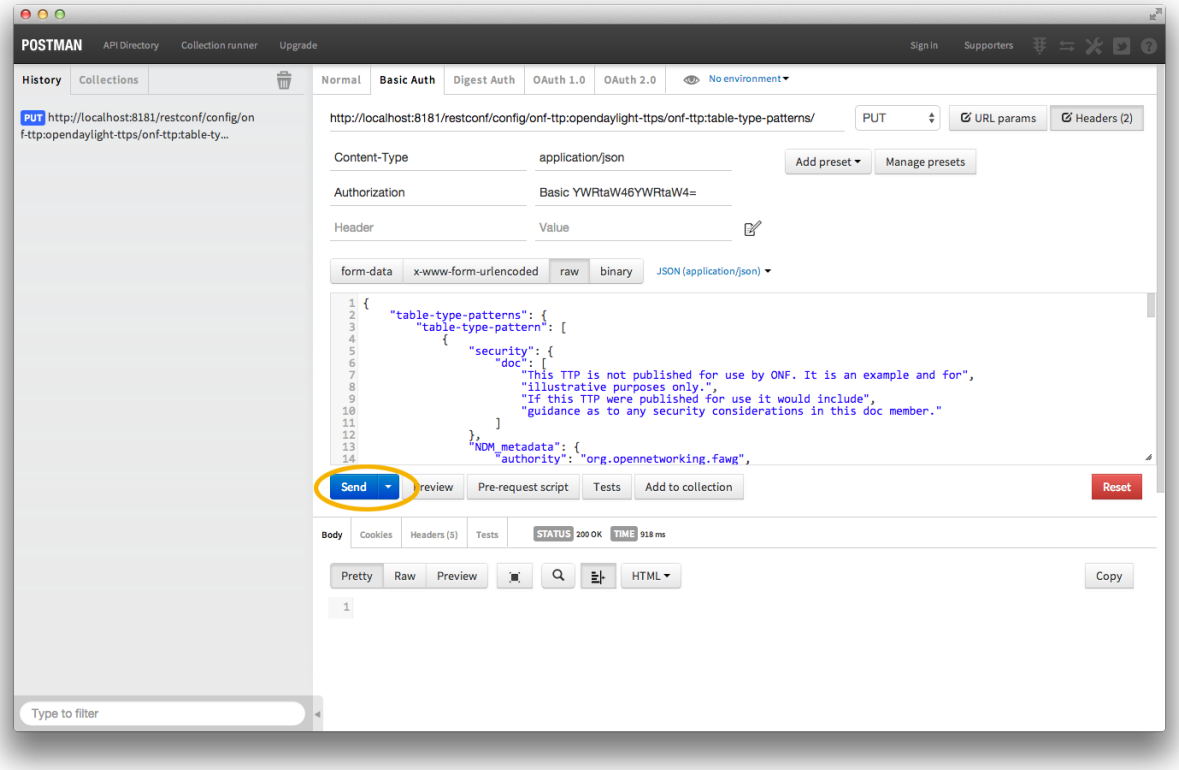

Fig. 64: PUTting a TTP

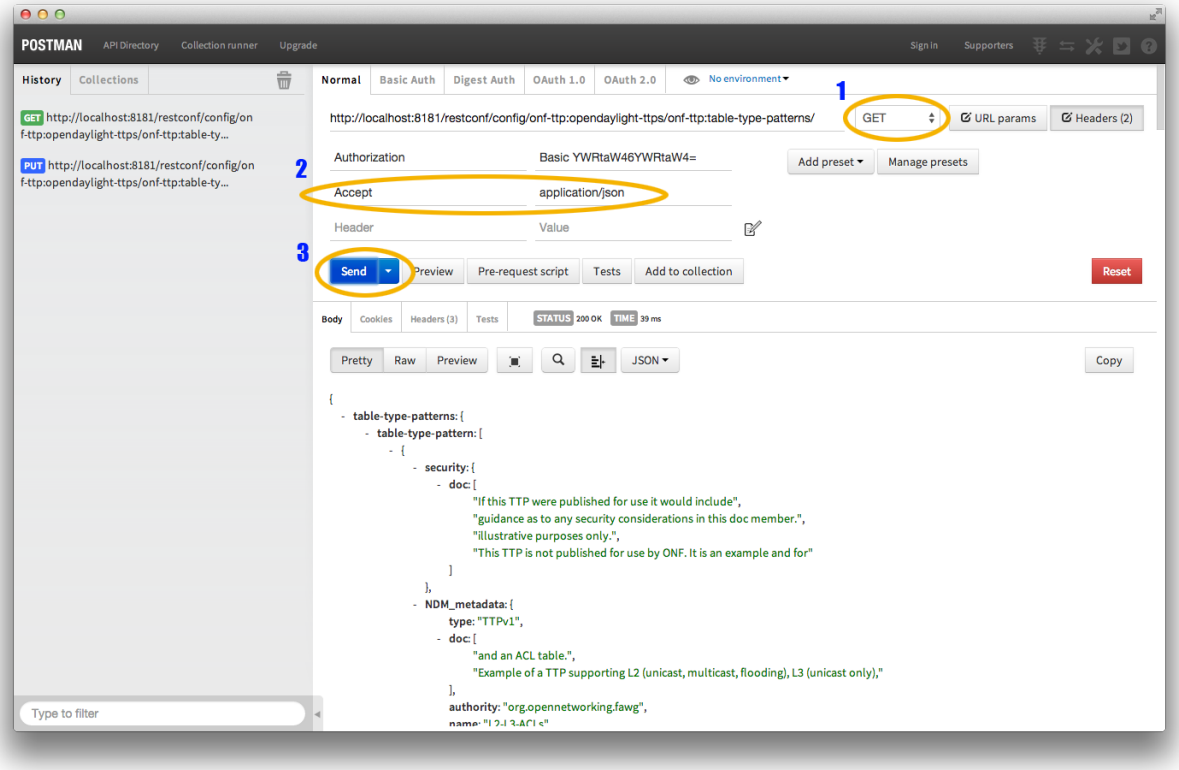

Fig. 65: Retrieving the TTP as json via a GET

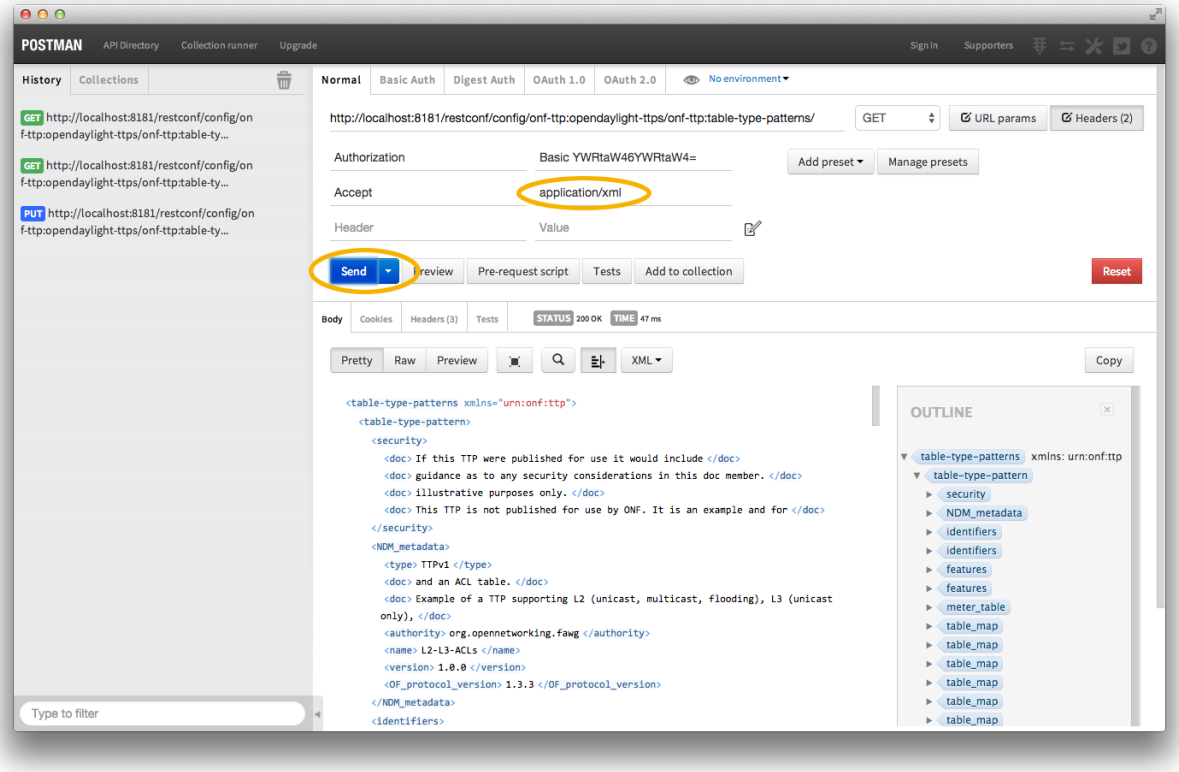

Fig. 66: Retrieving the TTP as xml via a GET

TTP CLI Tools Developer Guide

Overview

Table Type Patterns are a specification developed by the Open Networking Foundation to enable the description and negotiation of subsets of the OpenFlow protocol. This is particularly useful for hardware switches that support OpenFlow as it enables the to describe what features they do (and thus also what features they do not) support. More details can be found in the full specification listed on the OpenFlow specifications page.

The TTP CLI Tools provide a way for people interested in TTPs to read in, validate, output, and manipulate TTPs as a self-contained, executable jar file.

TTP CLI Tools Architecture

The TTP CLI Tools use the TTP Model and the YANG Tools/RESTCONF codecs to translate between the Data Transfer Objects (DTOs) and JSON/XML.

Command Line Options

This will cover the various options for the CLI Tools. For now, there are no options and it merely outputs fixed data using the codecs.

User Network Interface Manager Plug-in (Unimgr) Developer Guide

User Network Interface Manager Plug-in (Unimgr) is an experimental/proof of concept (PoC) project formed to initiate the development of data models and APIs facilitating the use by software applications and/or service orchestrators of OpenDaylight to configure and provision connectivity services, in particular Carrier Ethernet services as defined by Metro Ethernet Forum (MEF), in physical or virtual network elements.

Functionality

Unimgr provides support for both service orchestration, via the Legato API, and network resource provisioning, via the Presto API. These APIs, and the interfaces they provide, are defined by YANG models developed within MEF in collaboration with ONF and IETF. An application/user can interact with Unimgr at ether layer. For the Boron release, the YANG models are as follows:

Legato YANG models

https://git.opendaylight.org/gerrit/gitweb?p=unimgr.git;a=tree;f=legato-api/src/main/yang;hb=refs/heads/stable/boron

Presto YANG models

https://git.opendaylight.org/gerrit/gitweb?p=unimgr.git;a=tree;f=presto-api/src/main/yang;hb=refs/heads/stable/boron

Legato API Tree

module: mef-services

```
+--rw mef-services
   +--rw mef-service* [svc-id]
      +--rw evc
      |  +--rw unis
      |  |  +--rw uni* [uni-id]
      |  |     +--rw evc-uni-ce-vlans
      |  |     |  +--rw evc-uni-ce-vlan* [vid]
      |  |     |     +--rw vid    -> /mef-interfaces:mef-interfaces/unis/uni[mef-
↪interfaces:uni-id = current()/../../../uni-id]/ce-vlans/ce-vlan/vid
      |  |     +--rw ingress-bwp-flows-per-cos!
      |  |     |  +--rw coupling-enabled?   boolean
      |  |     |  +--rw bwp-flow-per-cos* [cos-name]
      |  |     |     +--rw cos-name      -> /mef-global:mef-global/profiles/cos-names/
↪cos-name/name
      |  |     |     +--rw bw-profile    -> /mef-interfaces:mef-interfaces/unis/
↪uni[mef-interfaces:uni-id = current()/../../../uni-id]/ingress-envelopes/envelope/
↪env-id
      |  |     +--rw egress-bwp-flows-per-eec!
      |  |     |  +--rw coupling-enabled?   boolean
      |  |     |  +--rw bwp-flow-per-eec* [eec-name]
      |  |     |     +--rw eec-name      -> /mef-global:mef-global/profiles/eec-names/
↪eec-name/name
      |  |     |     +--rw bw-profile    -> /mef-interfaces:mef-interfaces/unis/
↪uni[mef-interfaces:uni-id = current()/../../../uni-id]/egress-envelopes/envelope/
↪env-id
      |  |     +--rw status
      |  |     |  +--ro oper-state-enabled?   boolean
      |  |     |  +--ro available-status?     mef-types:svc-endpoint-availability-type
      |  |     +--rw uni-id                        -> /mef-interfaces:mef-interfaces/
↪unis/uni/uni-id
      |  |     +--rw role                          mef-types:evc-uni-role-type
      |  |     +--rw admin-state-enabled?          boolean
      |  |     +--rw color-id?                     mef-types:cos-color-identifier-
↪type
      |  |     +--rw data-svc-frm-cos?             -> /mef-global:mef-global/
↪profiles/cos/cos-profile/id
      |  |     +--rw l2cp-svc-frm-cos?             -> /mef-global:mef-global/
↪profiles/l2cp-cos/l2cp-profile/id
      |  |     +--rw soam-svc-frm-cos?             -> /mef-global:mef-global/
↪profiles/cos/cos-profile/id
      |  |     +--rw data-svc-frm-eec?             -> /mef-global:mef-global/
↪profiles/eec/eec-profile/id
      |  |     +--rw l2cp-svc-frm-eec?             -> /mef-global:mef-global/
↪profiles/l2cp-eec/l2cp-profile/id
      |  |     +--rw soam-svc-frm-eec?             -> /mef-global:mef-global/
↪profiles/eec/eec-profile/id
      |  |     +--rw ingress-bw-profile-per-evc?    -> /mef-interfaces:mef-interfaces/
↪unis/uni[mef-interfaces:uni-id = current()/../uni-id]/ingress-envelopes/envelope/
↪env-id
      |  |     +--rw egress-bw-profile-per-evc?     -> /mef-interfaces:mef-interfaces/
↪unis/uni[mef-interfaces:uni-id = current()/../uni-id]/egress-envelopes/envelope/env-
↪id
      |  |     +--rw src-mac-addr-limit-enabled?    boolean
```

(continues on next page)

(continued from previous page)

```
      |  |     +--rw src-mac-addr-limit?            uint32
      |  |     +--rw src-mac-addr-limit-interval?   yang:timeticks
      |  |     +--rw test-meg-enabled?              boolean
      |  |     +--rw test-meg?                      mef-types:identifier45
      |  |     +--rw subscriber-meg-mip-enabled?    boolean
      |  |     +--rw subscriber-meg-mip?            mef-types:identifier45
      |  +--rw status
      |  |  +--ro oper-state-enabled?   boolean
      |  |  +--ro available-status?     mef-types:virt-cx-availability-type
      |  +--rw sls-inclusions-by-cos
      |  |  +--rw sls-inclusion-by-cos* [cos-name]
      |  |     +--rw cos-name    -> /mef-global:mef-global/profiles/cos-names/cos-
↪name/name
      |  +--rw sls-uni-inclusions!
      |  |  +--rw sls-uni-inclusion-set* [pm-type pm-id uni-id1 uni-id2]
      |  |     +--rw pm-type    -> /mef-global:mef-global/slss/sls[mef-global:sls-id␣
↪= current()/../../../evc-performance-sls]/perf-objs/perf-obj/pm-type
      |  |     +--rw pm-id      -> /mef-global:mef-global/slss/sls[mef-global:sls-id␣
↪= current()/../../../evc-performance-sls]/perf-objs/perf-obj[mef-global:pm-type =␣
↪current()/../pm-type]/pm-id
      |  |     +--rw uni-id1    -> ../../../unis/uni/uni-id
      |  |     +--rw uni-id2    -> ../../../unis/uni/uni-id
      |  +--rw sls-uni-exclusions!
      |  |  +--rw sls-uni-exclusion-set* [pm-type pm-id uni-id1 uni-id2]
      |  |     +--rw pm-type    -> /mef-global:mef-global/slss/sls[mef-global:sls-id␣
↪= current()/../../../evc-performance-sls]/perf-objs/perf-obj/pm-type
      |  |     +--rw pm-id      -> /mef-global:mef-global/slss/sls[mef-global:sls-id␣
↪= current()/../../../evc-performance-sls]/perf-objs/perf-obj[mef-global:pm-type =␣
↪current()/../pm-type]/pm-id
      |  |     +--rw uni-id1    -> ../../../unis/uni/uni-id
      |  |     +--rw uni-id2    -> ../../../unis/uni/uni-id
      |  +--rw evc-id                          mef-types:evc-id-type
      |  +--ro evc-status?                     mef-types:evc-status-type
      |  +--rw evc-type                        mef-types:evc-type
      |  +--rw admin-state-enabled?            boolean
      |  +--rw elastic-enabled?                boolean
      |  +--rw elastic-service?                mef-types:identifier45
      |  +--rw max-uni-count?                  uint32
      |  +--rw preserve-ce-vlan-id?            boolean
      |  +--rw cos-preserve-ce-vlan-id?        boolean
      |  +--rw evc-performance-sls?            -> /mef-global:mef-global/slss/sls/sls-id
      |  +--rw unicast-svc-frm-delivery?       mef-types:data-svc-frame-delivery-type
      |  +--rw multicast-svc-frm-delivery?     mef-types:data-svc-frame-delivery-type
      |  +--rw broadcast-svc-frm-delivery?     mef-types:data-svc-frame-delivery-type
      |  +--rw evc-meg-id?                     mef-types:identifier45
      |  +--rw max-svc-frame-size?             mef-types:max-svc-frame-size-type
      +--rw svc-id        mef-types:retail-svc-id-type
      +--rw sp-id?        -> /mef-global:mef-global/svc-providers/svc-provider/sp-id
      +--rw svc-type?     mef-types:mef-service-type
      +--rw user-label?   mef-types:identifier45
      +--rw svc-entity?   mef-types:service-entity-type
```

module: mef-global

```
+--rw mef-global
   +--rw svc-providers!
   |  +--rw svc-provider* [sp-id]
```

(continues on next page)

(continued from previous page)

```
   |     +--rw sp-id    mef-types:svc-provider-type
   +--rw cens!
   |  +--rw cen* [cen-id]
   |     +--rw cen-id    mef-types:cen-type
   |     +--rw sp-id?    -> /mef-global/svc-providers/svc-provider/sp-id
   +--rw slss!
   |  +--rw sls* [sls-id]
   |     +--rw perf-objs
   |     |  +--rw pm-time-interval                    uint64
   |     |  +--rw pm-time-interval-increment          uint64
   |     |  +--rw unavail-flr-threshold-pp            mef-types:simple-percent
   |     |  +--rw consecutive-small-time-intervals    uint64
   |     |  +--rw perf-obj* [pm-type pm-id]
   |     |     +--rw pm-type                                mef-types:performance-
↪metric-type
   |     |     +--rw pm-id                                  mef-types:identifier45
   |     |     +--rw cos-name                               -> /mef-global/profiles/
↪cos-names/cos-name/name
   |     |     +--rw fd-pp                                  mef-types:simple-percent
   |     |     +--rw fd-range-pp                            mef-types:simple-percent
   |     |     +--rw fd-perf-obj                            uint64
   |     |     +--rw fd-range-perf-obj                      uint64
   |     |     +--rw fd-mean-perf-obj                       uint64
   |     |     +--rw ifdv-pp                                mef-types:simple-percent
   |     |     +--rw ifdv-pair-interval                     mef-types:simple-percent
   |     |     +--rw ifdv-perf-obj                          uint64
   |     |     +--rw flr-perf-obj                           uint64
   |     |     +--rw avail-pp                               mef-types:simple-percent
   |     |     +--rw hli-perf-obj                           uint64
   |     |     +--rw chli-consecutive-small-time-intervals  uint64
   |     |     +--rw chli-perf-obj                          uint64
   |     |     +--rw min-uni-pairs-avail                    uint64
   |     |     +--rw gp-avail-pp                            mef-types:simple-percent
   |     +--rw sls-id       mef-types:cen-type
   |     +--rw sp-id?       -> /mef-global/svc-providers/svc-provider/sp-id
   +--rw subscribers!
   |  +--rw subscriber* [sub-id]
   |     +--rw sub-id    mef-types:subscriber-type
   |     +--rw sp-id?    -> /mef-global/svc-providers/svc-provider/sp-id
   |     +--rw cen-id?   -> /mef-global/cens/cen/cen-id
   +--rw profiles!
      +--rw cos-names
      |  +--rw cos-name* [name]
      |     +--rw name    mef-types:identifier45
      +--rw eec-names
      |  +--rw eec-name* [name]
      |     +--rw name    mef-types:identifier45
      +--rw ingress-bwp-flows
      |  +--rw bwp-flow* [bw-profile]
      |     +--rw bw-profile          mef-types:identifier45
      |     +--rw user-label?         mef-types:identifier45
      |     +--rw cir?                mef-types:bwp-cir-type
      |     +--rw cir-max?            mef-types:bwp-cir-type
      |     +--rw cbs?                mef-types:bwp-cbs-type
      |     +--rw eir?                mef-types:bwp-eir-type
      |     +--rw eir-max?            mef-types:bwp-eir-type
      |     +--rw ebs?                mef-types:bwp-ebs-type
```

(continues on next page)

(continued from previous page)

```
      |     +--rw coupling-enabled?   boolean
      |     +--rw color-mode?         mef-types:bwp-color-mode-type
      |     +--rw coupling-flag?      mef-types:bwp-coupling-flag-type
      +--rw egress-bwp-flows
      |  +--rw bwp-flow* [bw-profile]
      |     +--rw bw-profile          mef-types:identifier45
      |     +--rw user-label?         mef-types:identifier45
      |     +--rw cir?                mef-types:bwp-cir-type
      |     +--rw cir-max?            mef-types:bwp-cir-type
      |     +--rw cbs?                mef-types:bwp-cbs-type
      |     +--rw eir?                mef-types:bwp-eir-type
      |     +--rw eir-max?            mef-types:bwp-eir-type
      |     +--rw ebs?                mef-types:bwp-ebs-type
      |     +--rw coupling-enabled?   boolean
      |     +--rw color-mode?         mef-types:bwp-color-mode-type
      |     +--rw coupling-flag?      mef-types:bwp-coupling-flag-type
      +--rw l2cp-cos
      |  +--rw l2cp-profile* [id]
      |     +--rw l2cps
      |     |  +--rw l2cp* [dest-mac-addr peering-proto-name]
      |     |     +--rw dest-mac-addr         yang:mac-address
      |     |     +--rw peering-proto-name    mef-types:identifier45
      |     |     +--rw protocol?             mef-types:l2cp-peering-protocol-type
      |     |     +--rw protocol-id?          yang:hex-string
      |     |     +--rw cos-name?             -> /mef-global/profiles/cos-names/cos-
↪name/name
      |     |     +--rw handling?             mef-types:l2cp-handling-type
      |     |     +--rw subtype*              yang:hex-string
      |     +--rw id            mef-types:identifier45
      |     +--rw user-label?   mef-types:identifier45
      +--rw l2cp-eec
      |  +--rw l2cp-profile* [id]
      |     +--rw l2cps
      |     |  +--rw l2cp* [dest-mac-addr peering-proto-name]
      |     |     +--rw dest-mac-addr         yang:mac-address
      |     |     +--rw peering-proto-name    mef-types:identifier45
      |     |     +--rw protocol?             mef-types:l2cp-peering-protocol-type
      |     |     +--rw protocol-id?          yang:hex-string
      |     |     +--rw eec-name?             -> /mef-global/profiles/eec-names/eec-
↪name/name
      |     |     +--rw handling?             mef-types:l2cp-handling-type
      |     |     +--rw subtype*              yang:hex-string
      |     +--rw id            mef-types:identifier45
      |     +--rw user-label?   mef-types:identifier45
      +--rw l2cp-peering
      |  +--rw l2cp-profile* [id]
      |     +--rw l2cps
      |     |  +--rw l2cp* [dest-mac-addr peering-proto-name]
      |     |     +--rw dest-mac-addr         yang:mac-address
      |     |     +--rw peering-proto-name    mef-types:identifier45
      |     |     +--rw protocol?             mef-types:l2cp-peering-protocol-type
      |     |     +--rw protocol-id?          yang:hex-string
      |     |     +--rw subtype*              yang:hex-string
      |     +--rw id            mef-types:identifier45
      |     +--rw user-label?   mef-types:identifier45
      +--rw elmi
      |  +--rw elmi-profile* [id]
```

(continues on next page)

(continued from previous page)

```
      |     +--rw id                              mef-types:identifier45
      |     +--rw user-label?                     mef-types:identifier45
      |     +--rw polling-counter?                mef-types:elmi-polling-counter-type
      |     +--rw status-error-threshold?         mef-types:elmi-status-error-threshold-
↪type
      |     +--rw polling-timer?                  mef-types:elmi-polling-timer-type
      |     +--rw polling-verification-timer?     mef-types:elmi-polling-verification-
↪timer-type
      +--rw eec
      |  +--rw eec-profile* [id]
      |     +--rw id          mef-types:identifier45
      |     +--rw (eec-id)?
      |        +--:(pcp)
      |        |  +--rw eec-pcp!
      |        |     +--rw default-pcp-eec-name?   -> /mef-global/profiles/eec-names/
↪eec-name/name
      |        |     +--rw default-pcp-color?      mef-types:cos-color-type
      |        |     +--rw pcp* [pcp-value]
      |        |        +--rw pcp-value        mef-types:ieee8021p-priority-type
      |        |        +--rw discard-value?   boolean
      |        |        +--rw eec-name?        -> /mef-global/profiles/eec-names/eec-
↪name/name
      |        |        +--rw color?           mef-types:cos-color-type
      |        +--:(dscp)
      |           +--rw eec-dscp!
      |              +--rw default-ipv4-eec-name?   -> /mef-global/profiles/eec-names/
↪eec-name/name
      |              +--rw default-ipv4-color?      mef-types:cos-color-type
      |              +--rw default-ipv6-eec-name?   -> /mef-global/profiles/eec-names/
↪eec-name/name
      |              +--rw default-ipv6-color?      mef-types:cos-color-type
      |              +--rw ipv4-dscp* [dscp-value]
      |              |  +--rw dscp-value       inet:dscp
      |              |  +--rw discard-value?   boolean
      |              |  +--rw eec-name?        -> /mef-global/profiles/eec-names/eec-
↪name/name
      |              |  +--rw color?           mef-types:cos-color-type
      |              +--rw ipv6-dscp* [dscp-value]
      |                 +--rw dscp-value       inet:dscp
      |                 +--rw discard-value?   boolean
      |                 +--rw eec-name?        -> /mef-global/profiles/eec-names/eec-
↪name/name
      |                 +--rw color?           mef-types:cos-color-type
      +--rw cos
         +--rw cos-profile* [id]
            +--rw id          mef-types:identifier45
            +--rw (cos-id)?
               +--:(evc)
               |  +--rw cos-evc!
               |     +--rw default-evc-cos-name?   -> /mef-global/profiles/cos-names/
↪cos-name/name
               |     +--rw default-evc-color?      mef-types:cos-color-type
               +--:(pcp)
               |  +--rw cos-pcp!
               |     +--rw default-pcp-cos-name?   -> /mef-global/profiles/cos-names/
↪cos-name/name
               |     +--rw default-pcp-color?      mef-types:cos-color-type
```

(continues on next page)

(continued from previous page)

```
               |     +--rw pcp* [pcp-value]
               |        +--rw pcp-value        mef-types:ieee8021p-priority-type
               |        +--rw discard-value?   boolean
               |        +--rw cos-name?        -> /mef-global/profiles/cos-names/cos-
↪name/name
               |        +--rw color?           mef-types:cos-color-type
               +--:(dscp)
                  +--rw cos-dscp!
                     +--rw default-ipv4-cos-name?   -> /mef-global/profiles/cos-names/
↪cos-name/name
                     +--rw default-ipv4-color?      mef-types:cos-color-type
                     +--rw default-ipv6-cos-name?   -> /mef-global/profiles/cos-names/
↪cos-name/name
                     +--rw default-ipv6-color?      mef-types:cos-color-type
                     +--rw ipv4-dscp* [dscp-value]
                     |  +--rw dscp-value       inet:dscp
                     |  +--rw discard-value?   boolean
                     |  +--rw cos-name?        -> /mef-global/profiles/cos-names/cos-
↪name/name
                     |  +--rw color?           mef-types:cos-color-type
                     +--rw ipv6-dscp* [dscp-value]
                        +--rw dscp-value       inet:dscp
                        +--rw discard-value?   boolean
                        +--rw cos-name?        -> /mef-global/profiles/cos-names/cos-
↪name/name
                        +--rw color?           mef-types:cos-color-type
```

Presto API Tree

module: onf-core-network-module

```
+--rw forwarding-constructs
   +--rw forwarding-construct* [uuid]
      +--rw uuid                    string
      +--rw layerProtocolName?      onf-cnt:LayerProtocolName
      +--rw lowerLevelFc*           -> /forwarding-constructs/forwarding-construct/uuid
      +--rw fcRoute* [uuid]
      |  +--rw uuid    string
      |  +--rw fc*     -> /forwarding-constructs/forwarding-construct/uuid
      +--rw fcPort* [topology node tp]
      |  +--rw topology           nt:topology-ref
      |  +--rw node               nt:node-ref
      |  +--rw tp                 nt:tp-ref
      |  +--rw role?              onf-cnt:PortRole
      |  +--rw fcPortDirection?   onf-cnt:PortDirection
      +--rw fcSpec
      |  +--rw uuid?                      string
      |  +--rw fcPortSpec* [uuid]
      |  |  +--rw uuid                string
      |  |  +--rw ingressFcPortSet* [topology node tp]
      |  |  |  +--rw topology    nt:topology-ref
      |  |  |  +--rw node        nt:node-ref
      |  |  |  +--rw tp          nt:tp-ref
      |  |  +--rw egressFcPortSet* [topology node tp]
      |  |  |  +--rw topology    nt:topology-ref
```

(continues on next page)

(continued from previous page)

```
      |  |  |  +--rw node        nt:node-ref
      |  |  |  +--rw tp          nt:tp-ref
      |  |  +--rw role?               string
      |  +--rw nrp:nrp-ce-fcspec-attrs
      |     +--rw nrp:connectionType?           nrp-types:NRP_ConnectionType
      |     +--rw nrp:unicastFrameDelivery?     nrp-types:NRP_ServiceFrameDelivery
      |     +--rw nrp:multicastFrameDelivery?   nrp-types:NRP_ServiceFrameDelivery
      |     +--rw nrp:broadcastFrameDelivery?   nrp-types:NRP_ServiceFrameDelivery
      |     +--rw nrp:vcMaxServiceFrame?        nrp-types:NRP_PositiveInteger
      |     +--rw nrp:vcId?                     nrp-types:NRP_PositiveInteger
      +--rw forwardingDirection?   onf-cnt:ForwardingDirection
```

augment /nt:network-topology/nt:topology/nt:node/nt:termination-point:

```
+--rw ltp-attrs
   +--rw lpList* [uuid]
   |  +--rw uuid                        string
   |  +--rw layerProtocolName?          onf-cnt:LayerProtocolName
   |  +--rw lpSpec
   |  |  +--rw adapterSpec
   |  |  |  +--rw nrp:nrp-conn-adapt-spec-attrs
   |  |  |  |  +--rw nrp:sourceMacAddressLimit
   |  |  |  |  |  +--rw nrp:enabled?        boolean
   |  |  |  |  |  +--rw nrp:limit?          NRP_NaturalNumber
   |  |  |  |  |  +--rw nrp:timeInterval?   NRP_NaturalNumber
   |  |  |  |  +--rw nrp:CeExternalInterface
   |  |  |  |  |  +--rw nrp:physicalLayer?             nrp-types:NRP_PhysicalLayer
   |  |  |  |  |  +--rw nrp:syncMode* [linkId]
   |  |  |  |  |  |  +--rw nrp:linkId             string
   |  |  |  |  |  |  +--rw nrp:syncModeEnabled?   boolean
   |  |  |  |  |  +--rw nrp:numberOfLinks?             nrp-types:NRP_NaturalNumber
   |  |  |  |  |  +--rw nrp:resiliency?                nrp-types:NRP_
↪InterfaceResiliency
   |  |  |  |  |  +--rw nrp:portConvsIdToAggLinkMap
   |  |  |  |  |  |  +--rw nrp:conversationId?   NRP_NaturalNumber
   |  |  |  |  |  |  +--rw nrp:linkId?           NRP_NaturalNumber
   |  |  |  |  |  +--rw nrp:maxFrameSize?              nrp-types:NRP_NaturalNumber
   |  |  |  |  |  +--rw nrp:linkOamEnabled?            boolean
   |  |  |  |  |  +--rw nrp:tokenShareEnabled?         boolean
   |  |  |  |  |  +--rw nrp:serviceProviderUniId?      string
   |  |  |  |  +--rw nrp:coloridentifier
   |  |  |  |  |  +--rw (identifier)?
   |  |  |  |  |     +--:(sap-color-id)
   |  |  |  |  |     |  +--rw nrp:serviceAccessPointColorId
   |  |  |  |  |     |     +--rw nrp:color?   nrp-types:NRP_FrameColor
   |  |  |  |  |     +--:(pcp-color-id)
   |  |  |  |  |     |  +--rw nrp:pcpColorId
   |  |  |  |  |     |     +--rw nrp:vlanTag?    nrp-types:NRP_VlanTag
   |  |  |  |  |     |     +--rw nrp:pcpValue*   nrp-types:NRP_NaturalNumber
   |  |  |  |  |     |     +--rw nrp:color?      nrp-types:NRP_FrameColor
   |  |  |  |  |     +--:(dei-color-id)
   |  |  |  |  |     |  +--rw nrp:deiColorId
   |  |  |  |  |     |     +--rw nrp:vlanTag?    nrp-types:NRP_VlanTag
   |  |  |  |  |     |     +--rw nrp:deiValue*   nrp-types:NRP_NaturalNumber
   |  |  |  |  |     |     +--rw nrp:color?      nrp-types:NRP_FrameColor
   |  |  |  |  |     +--:(desp-color-id)
   |  |  |  |  |        +--rw nrp:despColorId
```

(continues on next page)

(continued from previous page)

```
|  |  |  |  |           +--rw nrp:ipVersion?   nrp-types:NRP_IpVersion
|  |  |  |  |           +--rw nrp:dscpValue*   nrp-types:NRP_NaturalNumber
|  |  |  |  |           +--rw nrp:color?       nrp-types:NRP_FrameColor
|  |  |  |  +--rw nrp:ingressBwpFlow
|  |  |  |  |  +--rw nrp:bwpFlowIndex?         nrp-types:NRP_PositiveInteger
|  |  |  |  |  +--rw nrp:cir?                  nrp-types:NRP_NaturalNumber
|  |  |  |  |  +--rw nrp:cirMax?               nrp-types:NRP_NaturalNumber
|  |  |  |  |  +--rw nrp:cbs?                  nrp-types:NRP_NaturalNumber
|  |  |  |  |  +--rw nrp:eir?                  nrp-types:NRP_NaturalNumber
|  |  |  |  |  +--rw nrp:eirMax?               nrp-types:NRP_NaturalNumber
|  |  |  |  |  +--rw nrp:ebs?                  nrp-types:NRP_NaturalNumber
|  |  |  |  |  +--rw nrp:couplingFlag?         nrp-types:NRP_NaturalNumber
|  |  |  |  |  +--rw nrp:colorMode?            nrp-types:NRP_ColorMode
|  |  |  |  |  +--rw nrp:rank?                 nrp-types:NRP_PositiveInteger
|  |  |  |  |  +--rw nrp:tokenRequestOffset?   nrp-types:NRP_NaturalNumber
|  |  |  |  +--rw nrp:egressBwpFlow
|  |  |  |  |  +--rw nrp:bwpFlowIndex?         nrp-types:NRP_PositiveInteger
|  |  |  |  |  +--rw nrp:cir?                  nrp-types:NRP_NaturalNumber
|  |  |  |  |  +--rw nrp:cirMax?               nrp-types:NRP_NaturalNumber
|  |  |  |  |  +--rw nrp:cbs?                  nrp-types:NRP_NaturalNumber
|  |  |  |  |  +--rw nrp:eir?                  nrp-types:NRP_NaturalNumber
|  |  |  |  |  +--rw nrp:eirMax?               nrp-types:NRP_NaturalNumber
|  |  |  |  |  +--rw nrp:ebs?                  nrp-types:NRP_NaturalNumber
|  |  |  |  |  +--rw nrp:couplingFlag?         nrp-types:NRP_NaturalNumber
|  |  |  |  |  +--rw nrp:colorMode?            nrp-types:NRP_ColorMode
|  |  |  |  |  +--rw nrp:rank?                 nrp-types:NRP_PositiveInteger
|  |  |  |  |  +--rw nrp:tokenRequestOffset?   nrp-types:NRP_NaturalNumber
|  |  |  |  +--rw nrp:l2cpAddressSet?          nrp-types:NRP_L2cpAddressSet
|  |  |  |  +--rw nrp:l2cpPeering* [linkId]
|  |  |  |     +--rw nrp:destinationMacAddress?   string
|  |  |  |     +--rw nrp:protocolType?            NRP_ProtocolFrameType
|  |  |  |     +--rw nrp:linkId                   string
|  |  |  |     +--rw nrp:protocolId?              string
|  |  |  +--rw nrp:nrp-ivc-endpoint-conn-adapt-spec-attrs
|  |  |  |  +--rw nrp:ivcEndPointId?             string
|  |  |  |  +--rw nrp:testMegEnabled?            boolean
|  |  |  |  +--rw nrp:ivcEndPointRole?           nrp-types:NRP_EndPointRole
|  |  |  |  +--rw nrp:ivcEndPointMap* [vlanId]
|  |  |  |  |  +--rw nrp:vlanId        nrp-types:NRP_PositiveInteger
|  |  |  |  |  +--rw (endpoint-map-form)?
|  |  |  |  |     +--:(map-form-e)
|  |  |  |  |     |  +--rw nrp:enni-svid* [vid]
|  |  |  |  |     |     +--rw nrp:vid    nrp-types:NRP_PositiveInteger
|  |  |  |  |     +--:(map-form-t)
|  |  |  |  |     |  +--rw nrp:root-svid?    nrp-types:NRP_PositiveInteger
|  |  |  |  |     |  +--rw nrp:leaf-svid?    nrp-types:NRP_PositiveInteger
|  |  |  |  |     +--:(map-form-v)
|  |  |  |  |     |  +--rw nrp:vuni-vid?     nrp-types:NRP_PositiveInteger
|  |  |  |  |     |  +--rw nrp:enni-cevid* [vid]
|  |  |  |  |     |     +--rw nrp:vid    nrp-types:NRP_PositiveInteger
|  |  |  |  |     +--:(map-form-u)
|  |  |  |  |        +--rw nrp:cvid* [vid]
|  |  |  |  |           +--rw nrp:vid    nrp-types:NRP_PositiveInteger
|  |  |  |  +--rw nrp:subscriberMegMipEnabled?   boolean
|  |  |  +--rw nrp:nrp-evc-endpoint-conn-adapt-spec-attrs
|  |  |     +--rw nrp:sourceMacAddressLimit
|  |  |     |  +--rw nrp:enabled?        boolean
```

(continues on next page)

(continued from previous page)

```
   |  |  |     |  +--rw nrp:limit?          NRP_NaturalNumber
   |  |  |     |  +--rw nrp:timeInterval?   NRP_NaturalNumber
   |  |  |     +--rw nrp:CeExternalInterface
   |  |  |     |  +--rw nrp:physicalLayer?             nrp-types:NRP_PhysicalLayer
   |  |  |     |  +--rw nrp:syncMode* [linkId]
   |  |  |     |  |  +--rw nrp:linkId              string
   |  |  |     |  |  +--rw nrp:syncModeEnabled?   boolean
   |  |  |     |  +--rw nrp:numberOfLinks?             nrp-types:NRP_NaturalNumber
   |  |  |     |  +--rw nrp:resiliency?                nrp-types:NRP_
↪InterfaceResiliency
   |  |  |     |  +--rw nrp:portConvsIdToAggLinkMap
   |  |  |     |  |  +--rw nrp:conversationId?   NRP_NaturalNumber
   |  |  |     |  |  +--rw nrp:linkId?           NRP_NaturalNumber
   |  |  |     |  +--rw nrp:maxFrameSize?              nrp-types:NRP_NaturalNumber
   |  |  |     |  +--rw nrp:linkOamEnabled?            boolean
   |  |  |     |  +--rw nrp:tokenShareEnabled?         boolean
   |  |  |     |  +--rw nrp:serviceProviderUniId?      string
   |  |  |     +--rw nrp:coloridentifier
   |  |  |     |  +--rw (identifier)?
   |  |  |     |     +--:(sap-color-id)
   |  |  |     |     |  +--rw nrp:serviceAccessPointColorId
   |  |  |     |     |     +--rw nrp:color?   nrp-types:NRP_FrameColor
   |  |  |     |     +--:(pcp-color-id)
   |  |  |     |     |  +--rw nrp:pcpColorId
   |  |  |     |     |     +--rw nrp:vlanTag?    nrp-types:NRP_VlanTag
   |  |  |     |     |     +--rw nrp:pcpValue*   nrp-types:NRP_NaturalNumber
   |  |  |     |     |     +--rw nrp:color?      nrp-types:NRP_FrameColor
   |  |  |     |     +--:(dei-color-id)
   |  |  |     |     |  +--rw nrp:deiColorId
   |  |  |     |     |     +--rw nrp:vlanTag?    nrp-types:NRP_VlanTag
   |  |  |     |     |     +--rw nrp:deiValue*   nrp-types:NRP_NaturalNumber
   |  |  |     |     |     +--rw nrp:color?      nrp-types:NRP_FrameColor
   |  |  |     |     +--:(desp-color-id)
   |  |  |     |        +--rw nrp:despColorId
   |  |  |     |           +--rw nrp:ipVersion?   nrp-types:NRP_IpVersion
   |  |  |     |           +--rw nrp:dscpValue*   nrp-types:NRP_NaturalNumber
   |  |  |     |           +--rw nrp:color?       nrp-types:NRP_FrameColor
   |  |  |     +--rw nrp:ingressBwpFlow
   |  |  |     |  +--rw nrp:bwpFlowIndex?         nrp-types:NRP_PositiveInteger
   |  |  |     |  +--rw nrp:cir?                  nrp-types:NRP_NaturalNumber
   |  |  |     |  +--rw nrp:cirMax?               nrp-types:NRP_NaturalNumber
   |  |  |     |  +--rw nrp:cbs?                  nrp-types:NRP_NaturalNumber
   |  |  |     |  +--rw nrp:eir?                  nrp-types:NRP_NaturalNumber
   |  |  |     |  +--rw nrp:eirMax?               nrp-types:NRP_NaturalNumber
   |  |  |     |  +--rw nrp:ebs?                  nrp-types:NRP_NaturalNumber
   |  |  |     |  +--rw nrp:couplingFlag?         nrp-types:NRP_NaturalNumber
   |  |  |     |  +--rw nrp:colorMode?            nrp-types:NRP_ColorMode
   |  |  |     |  +--rw nrp:rank?                 nrp-types:NRP_PositiveInteger
   |  |  |     |  +--rw nrp:tokenRequestOffset?   nrp-types:NRP_NaturalNumber
   |  |  |     +--rw nrp:egressBwpFlow
   |  |  |     |  +--rw nrp:bwpFlowIndex?         nrp-types:NRP_PositiveInteger
   |  |  |     |  +--rw nrp:cir?                  nrp-types:NRP_NaturalNumber
   |  |  |     |  +--rw nrp:cirMax?               nrp-types:NRP_NaturalNumber
   |  |  |     |  +--rw nrp:cbs?                  nrp-types:NRP_NaturalNumber
   |  |  |     |  +--rw nrp:eir?                  nrp-types:NRP_NaturalNumber
   |  |  |     |  +--rw nrp:eirMax?               nrp-types:NRP_NaturalNumber
   |  |  |     |  +--rw nrp:ebs?                  nrp-types:NRP_NaturalNumber
```

(continues on next page)

(continued from previous page)

```
|  |  |     |  +--rw nrp:couplingFlag?         nrp-types:NRP_NaturalNumber
|  |  |     |  +--rw nrp:colorMode?            nrp-types:NRP_ColorMode
|  |  |     |  +--rw nrp:rank?                 nrp-types:NRP_PositiveInteger
|  |  |     |  +--rw nrp:tokenRequestOffset?   nrp-types:NRP_NaturalNumber
|  |  |     +--rw nrp:l2cpAddressSet?            nrp-types:NRP_L2cpAddressSet
|  |  |     +--rw nrp:l2cpPeering* [linkId]
|  |  |     |  +--rw nrp:destinationMacAddress?   string
|  |  |     |  +--rw nrp:protocolType?            NRP_ProtocolFrameType
|  |  |     |  +--rw nrp:linkId                   string
|  |  |     |  +--rw nrp:protocolId?              string
|  |  |     +--rw nrp:evcEndPointId?             nrp-types:NRP_PositiveInteger
|  |  |     +--rw nrp:testMegEnabled?            boolean
|  |  |     +--rw nrp:evcEndPointRole?           nrp-types:NRP_EvcEndPointRole
|  |  |     +--rw nrp:evcEndPointMap* [vid]
|  |  |     |  +--rw nrp:vid    nrp-types:NRP_PositiveInteger
|  |  |     +--rw nrp:subscriberMegMipEbabled?   boolean
|  |  +--rw terminationSpec
|  |  |  +--rw nrp:nrp-termination-spec-attrs
|  |  |  |  +--rw nrp:physicalLayer?             nrp-types:NRP_PhysicalLayer
|  |  |  |  +--rw nrp:syncMode* [linkId]
|  |  |  |  |  +--rw nrp:linkId             string
|  |  |  |  |  +--rw nrp:syncModeEnabled?   boolean
|  |  |  |  +--rw nrp:numberOfLinks?             nrp-types:NRP_NaturalNumber
|  |  |  |  +--rw nrp:resiliency?                nrp-types:NRP_InterfaceResiliency
|  |  |  |  +--rw nrp:portConvsIdToAggLinkMap
|  |  |  |  |  +--rw nrp:conversationId?   NRP_NaturalNumber
|  |  |  |  |  +--rw nrp:linkId?           NRP_NaturalNumber
|  |  |  |  +--rw nrp:maxFrameSize?              nrp-types:NRP_NaturalNumber
|  |  |  |  +--rw nrp:linkOamEnabled?            boolean
|  |  |  |  +--rw nrp:tokenShareEnabled?         boolean
|  |  |  |  +--rw nrp:serviceProviderUniId?      string
|  |  |  +--rw nrp:nrp-uni-termination-attrs
|  |  |     +--rw nrp:defaultCeVlanId?             nrp-types:NRP_PositiveInteger
|  |  |     +--rw nrp:uniMegEnabled?               boolean
|  |  |     +--rw nrp:elmiEnabled?                 boolean
|  |  |     +--rw nrp:serviceprovideruniprofile?   string
|  |  |     +--rw nrp:operatoruniprofile?          string
|  |  |     +--rw nrp:ingressBwpUni
|  |  |     |  +--rw nrp:bwpFlowIndex?         nrp-types:NRP_PositiveInteger
|  |  |     |  +--rw nrp:cir?                  nrp-types:NRP_NaturalNumber
|  |  |     |  +--rw nrp:cirMax?               nrp-types:NRP_NaturalNumber
|  |  |     |  +--rw nrp:cbs?                  nrp-types:NRP_NaturalNumber
|  |  |     |  +--rw nrp:eir?                  nrp-types:NRP_NaturalNumber
|  |  |     |  +--rw nrp:eirMax?               nrp-types:NRP_NaturalNumber
|  |  |     |  +--rw nrp:ebs?                  nrp-types:NRP_NaturalNumber
|  |  |     |  +--rw nrp:couplingFlag?         nrp-types:NRP_NaturalNumber
|  |  |     |  +--rw nrp:colorMode?            nrp-types:NRP_ColorMode
|  |  |     |  +--rw nrp:rank?                 nrp-types:NRP_PositiveInteger
|  |  |     |  +--rw nrp:tokenRequestOffset?   nrp-types:NRP_NaturalNumber
|  |  |     +--rw nrp:egressBwpUni
|  |  |        +--rw nrp:bwpFlowIndex?         nrp-types:NRP_PositiveInteger
|  |  |        +--rw nrp:cir?                  nrp-types:NRP_NaturalNumber
|  |  |        +--rw nrp:cirMax?               nrp-types:NRP_NaturalNumber
|  |  |        +--rw nrp:cbs?                  nrp-types:NRP_NaturalNumber
|  |  |        +--rw nrp:eir?                  nrp-types:NRP_NaturalNumber
|  |  |        +--rw nrp:eirMax?               nrp-types:NRP_NaturalNumber
|  |  |        +--rw nrp:ebs?                  nrp-types:NRP_NaturalNumber
```

(continues on next page)

(continued from previous page)

```
|  |  |        +--rw nrp:couplingFlag?         nrp-types:NRP_NaturalNumber
|  |  |        +--rw nrp:colorMode?            nrp-types:NRP_ColorMode
|  |  |        +--rw nrp:rank?                 nrp-types:NRP_PositiveInteger
|  |  |        +--rw nrp:tokenRequestOffset?   nrp-types:NRP_NaturalNumber
|  |  +--rw adapterPropertySpecList* [uuid]
|  |  |  +--rw uuid    string
|  |  +--rw providerViewSpec
|  |  +--rw serverSpecList* [uuid]
|  |     +--rw uuid    string
|  +--rw configuredClientCapacity?   string
|  +--rw lpDirection?                onf-cnt:TerminationDirection
|  +--rw terminationState?           string
+--rw ltpSpec
+--rw ltpDirection?   onf-cnt:TerminationDirection
```

Unified Secure Channel

Overview

The Unified Secure Channel (USC) feature provides REST API, manager, and plugin for unified secure channels. The REST API provides a northbound api. The manager monitors, maintains, and provides channel related services. The plugin handles the lifecycle of channels.

USC Channel Architecture

- USC Agent
  - The USC Agent provides proxy and agent functionality on top of all standard protocols supported by the device. It initiates call-home with the controller, maintains live connections with with the controller, acts as a demuxer/muxer for packets with the USC header, and authenticates the controller.
- USC Plugin
  - The USC Plugin is responsible for communication between the controller and the USC agent . It responds to call-home with the controller, maintains live connections with the devices, acts as a muxer/demuxer for packets with the USC header, and provides support for TLS/DTLS.
- USC Manager
  - The USC Manager handles configurations, high availability, security, monitoring, and clustering support for USC.
- USC UI
  - The USC UI is responsible for displaying a graphical user interface representing the state of USC in the OpenDaylight DLUX UI.

USC Channel APIs and Interfaces

This section describes the APIs for interacting with the unified secure channels.

USC Channel Topology API

The USC project maintains a topology that is YANG-based in MD-SAL. These models are available via RESTCONF.

- Name: view-channel
- URL: http://\protect\T1\textdollar\protect\T1\textbraceleftipaddress\protect\T1\textbraceright:8181/restconf/operations/usc-channel:view-channel
- Description: Views the current state of the USC environment.

API Reference Documentation

Go to http://\protect\T1\textdollar\protect\T1\textbraceleftipaddress\protect\T1\textbraceright:8181/apidoc/explorer/index.html, sign in, and expand the usc-channel panel. From there, users can execute various API calls to test their USC deployment.

Usecplugin-AAA Developer Guide

Overview

Usecplugin-AAA provides security related information for the AAA northbound interface.

Usecplugin-AAA Architecture

AAA plugin creates log messages about successful and failed login attempts to OpenDaylight. Usecplugin-AAA continuously reads this log file and checks for either successful and failed attempt information. Whenever Usecpluin-AAA identifies a new attempt entry in the log file it is stored in YANG Data Store and its own log file.

Usecplugin-AAA is implemented with the help of a few java classes.

UsecpluginAAAProvider Provider class for Usecplugin-AAA feature implementation.

UsecpluginAAANotifImpl Logs notification information which can be seen by log:display at the Karaf terminal

UsecpluginAAARPCImpl Implements Usecplugin RPCs

UsecpluginAAAParsingLog Parses OpenDaylight log information for identifying login attempts.

UsecpluginAAAPublishNotif Publishes failed login attempt notification.

UsecpluginAAAStore Creates login information at the YANG Data Store.

Key APIs and Interfaces

- RPC APIs

  Login Attempt from IP Returns Time and Type of Attempts (Success or Failure)

  Login Attempt at Time Returns Attempter IP Address and Type of Attempts (Success or Failure)
- Notification APIs

  On Invalid Login Attempt Notification generated on Invalid Login Attempt
- YANG Data Store APIs

Get Login Attempts Returns Source IP address of Attempter with Time of Attempts and Type of Attempts (Success or Failure)

Usecplugin-OpenFlow Developer Guide

Overview

Usecplugin-OpenFlow provides security related information for the OpenFlow southbound interface.

Usecplugin-OpenFlow Architecture

Usecplugin-OpenFlow listens on OpenFlow southbound interface for Packet_In messages. The application parses the message for header information. Usecplugin-OpenFlow has PacketHandler class that implements the PacketProcessing interface to override the OnPacketReceived notification by which the application is notified of Packet_In messages.

Usecplugin-OpenFlow is implemented with the help of a few java classes.

UsecpluginProvider Provider class for Usecplugin-OpenFlow feature implementation.

PacketHandler Receives Packet_In messages coming to the controller and process them appropriately

PacketParsing Decodes Packet_In messages for packet header information (L2, L3 & L4 information)

InventoryUtility Decodes Packet_In messages for OpenFlow Switch and Port information

UsecpluginNotifImpl Logs notification information which can be seen by log:display at the Karaf terminal

UsecpluginRPCImpl Implements Usecplugin RPCs

UsecpluginStore Stores attack information into YANG Data Store and log file.

Key APIs and Interfaces

- RPC APIs

  Attacks from DPID Number of OpenFlow Packet_In Attacks from Switch with DeviceID

  Attacks from Host Number of OpenFlow Packet_In Attacks from SrcIP Address

  Attacks to Server Number of OpenFlow Packet_In Attacks to DstIP Address

  Attacks at Time of Day Number of OpenFlow Packet_In Attacks at a Particular Time with a variable Window Time

- Notification APIs

  On Low Water Mark Breached Notification generated on breaching Low Water Mark

Virtual Tenant Network (VTN)

OpenDaylight Virtual Tenant Network (VTN) Overview

OpenDaylight Virtual Tenant Network (VTN) is an application that provides multi-tenant virtual network on an SDN controller.

Conventionally, huge investment in the network systems and operating expenses are needed because the network is configured as a silo for each department and system. Therefore various network appliances must be installed for each

tenant and those boxes cannot be shared with others. It is a heavy work to design, implement and operate the entire complex network.

The uniqueness of VTN is a logical abstraction plane. This enables the complete separation of logical plane from physical plane. Users can design and deploy any desired network without knowing the physical network topology or bandwidth restrictions.

VTN allows the users to define the network with a look and feel of conventional L2/L3 network. Once the network is designed on VTN, it will automatically be mapped into underlying physical network, and then configured on the individual switch leverage SDN control protocol. The definition of logical plane makes it possible not only to hide the complexity of the underlying network but also to better manage network resources. It achieves reducing reconfiguration time of network services and minimizing network configuration errors. OpenDaylight Virtual Tenant Network (VTN) is an application that provides multi-tenant virtual network on an SDN controller. It provides API for creating a common virtual network irrespective of the physical network.

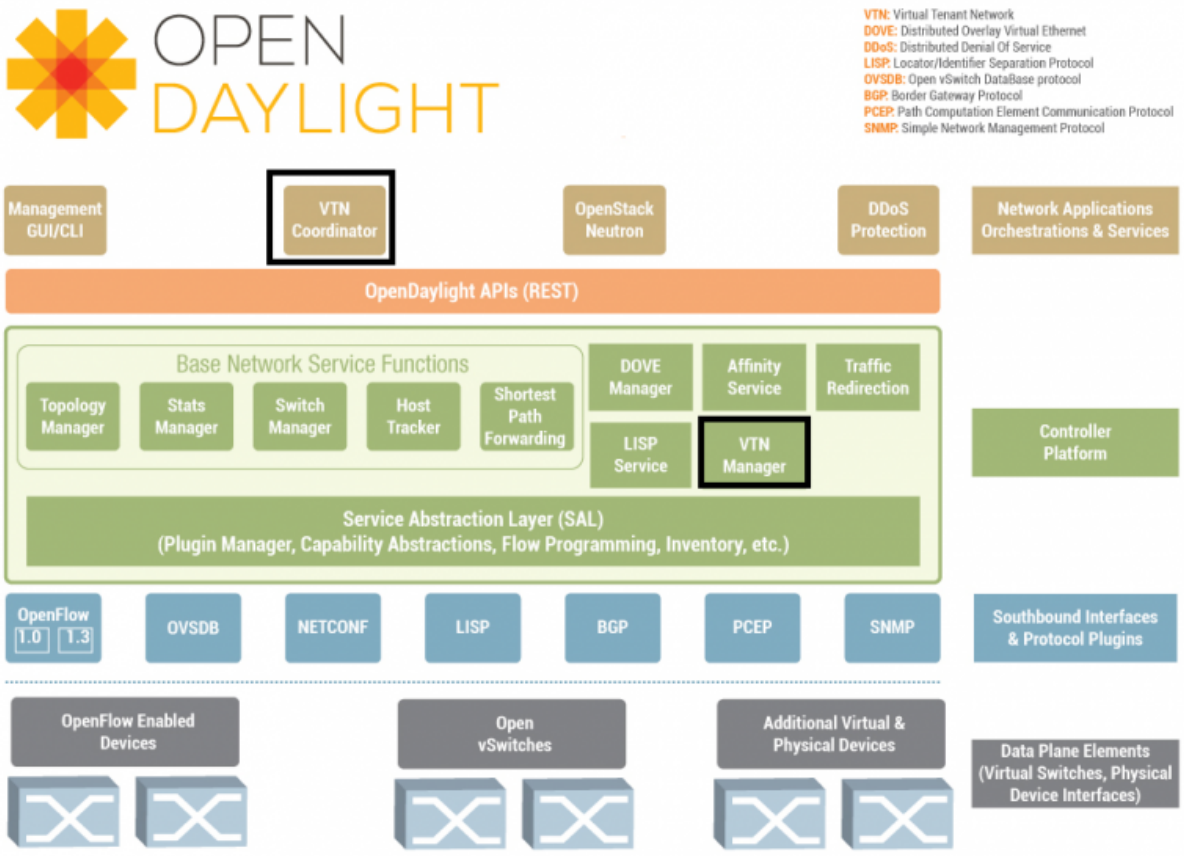

Fig. 67: VTN Architecture

It is implemented as two major components

- *VTN Manager*
- *VTN Coordinator*

VTN Manager

An OpenDaylight Plugin that interacts with other modules to implement the components of the VTN model. It also provides a REST interface to configure VTN components in OpenDaylight. VTN Manager is implemented as one plugin to the OpenDaylight. This provides a REST interface to create/update/delete VTN components. The user command in VTN Coordinator is translated as REST API to VTN Manager by the OpenDaylight Driver component. In addition to the above mentioned role, it also provides an implementation to the OpenStack L2 Network Functions API.

Function Outline

The table identifies the functions and the interface used by VTN Components:

Component	Interface	Purpose
VTN Manager	RESTful API	Configure VTN Virtualization model components in OpenDaylight
VTN Manager	Neutron API implementation	Handle Networks API from OpenStack (Neutron Interface)
VTN Coordinator	RESTful API	(1) Uses the RESTful interface of VTN Manager and configures VTN Virtualization model components in OpenDaylight. (2) Handles multiple OpenDaylight orchestration. (3) Provides API to read the physical network details. See samples for usage.

Feature Overview

There are three features

- **odl-vtn-manager** provides VTN Manager's JAVA API.
- **odl-vtn-manager-rest** provides VTN Manager's REST API.
- **odl-vtn-manager-neutron** provides the integration with Neutron interface.

REST Conf documentation for VTN Manager, please refer to: https://nexus.opendaylight.org/content/sites/site/org.opendaylight.vtn/boron/manager.model/apidocs/index.html

For VTN Java API documentation, please refer to: https://nexus.opendaylight.org/content/sites/site/org.opendaylight.vtn/boron/apidocs/index.html

Once the Karaf distribution is up, install dlux and apidocs.

```
feature:install odl-dlux-all odl-mdsal-apidocs
```

Logging In

To Log in to DLUX, after installing the application:

- Open a browser and enter the login URL as http://<OpenDaylight-IP>:8181/index.html

Note: Replace "<OpenDaylight-IP>" with the IP address of OpenDaylight based on your environment.

- Login to the application with user ID and password credentials as admin.

Note: admin is the only default user available for DLUX in this release.

- In the right hand side frame, click "Yang UI".

YANG documentation for VTN Manager, please refer to: https://nexus.opendaylight.org/content/sites/site/org.opendaylight.vtn/boron/manager.model/apidocs/index.html

VTN Coordinator

The VTN Coordinator is an external application that provides a REST interface for an user to use OpenDaylight VTN Virtualization. It interacts with the VTN Manager plugin to implement the user configuration. It is also capable of multiple OpenDaylight orchestration. It realizes VTN provisioning in OpenDaylight instances. In the OpenDaylight architecture VTN Coordinator is part of the network application, orchestration and services layer. VTN Coordinator will use the REST interface exposed by the VTN Manger to realize the virtual network using OpenDaylight. It uses OpenDaylight APIs (REST) to construct the virtual network in OpenDaylight instances. It provides REST APIs for northbound VTN applications and supports virtual networks spanning across multiple OpenDaylight by coordinating across OpenDaylight.

VTN Coordinator Components:

- Transaction Coordinator
- Unified Provider Physical Layer (UPPL)
- Unified Provider Logical LAyer (UPLL)
- OpenDaylight Controller Diver (ODC Driver)

OpenDaylight Virtual Tenant Network (VTN) API Overview

The VTN API module is a sub component of the VTN Coordinator and provides the northbound REST API interface for VTN applications. It consists of two subcomponents:

- Web Server
- VTN service Java API Library

Web Server

The Web Server module handles the REST APIs received from the VTN applications. It translates the REST APIs to the appropriate Java APIs.

The main functions of this module are:

- Starts via the startup script `catalina.sh`.
- VTN Application sends HTTP request to Web server in XML or JSON format.
- Creates a session and acquire a read/write lock.
- Invokes the VTN Service Java API Library corresponding to the specified URI.

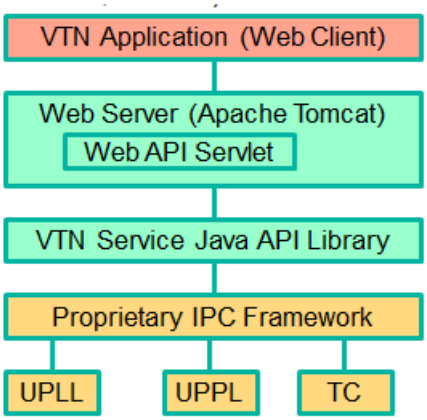

Fig. 68: VTN-Coordinator_api-architechture

- Returns the response to the VTN Application.

WebServer Class Details

The list below shows the classes available for Web Server module and their descriptions:

Init Manager It is a singleton class for executing the acquisition of configuration information from properties file, log initialization, initialization of VTN Service Java API Library. Executed by init() of VtnServiceWebAPIServlet.

Configuration Manager Maintains the configuration information acquired from properties file.

VtnServiceCommonUtil Utility class

VtnServiceWebUtil Utility class

VtnServiceWebAPIServlet Receives HTTP request from VTN Application and calls the method of corresponding VtnServiceWebAPIHandler. herits class HttpServlet, and overrides doGet(), doPut(), doDelete(), doPost().

VtnServiceWebAPIHandler Creates JsonObject(com.google.gson) from HTTP request, and calls method of corresponding VtnServiceWebAPIController.

VtnServiceWebAPIController Creates RestResource() class and calls UPLL API/UPPL API through Java API. the time of calling UPLL API/UPPL API, performs the creation/deletion of session, acquisition/release of configuration mode, acquisition/release of read lock by TC API through Java API.

Data Converter Converts HTTP request to JsonObject and JsonXML to JSON.

VTN Service Java API Library

It provides the Java API library to communicate with the lower layer modules in the VTN Coordinator. The main functions of this library are:

- Creates an IPC client session to the lower layer.
- Converts the request to IPC framework format.
- Invokes the lower layer API (i.e. UPPL API, UPLL API, TC API).
- Returns the response from the lower layer to the web server
- VTN Service Java API Library Class Details

Feature Overview

VTN Coordinator doesn't have Karaf features.

For VTN Coordinator REST API, please refer to: https://wiki.opendaylight.org/view/OpenDaylight_Virtual_Tenant_Network_%28VTN%29:VTN_Coordinator:RestApi

Usage Examples

- L2 Network using Single Controller

YANG Tools Developer Guide

Overview

YANG Tools is set of libraries and tooling providing support for use YANG for Java (or other JVM-based language) projects and applications.

YANG Tools provides following features in OpenDaylight:

- parsing of YANG sources and semantic inference of relationship across YANG models as defined in RFC6020
- representation of YANG-modeled data in Java
  - **Normalized Node** representation - DOM-like tree model, which uses conceptual meta-model more tailored to YANG and OpenDaylight use-cases than a standard XML DOM model allows for.
- serialization / deserialization of YANG-modeled data driven by YANG models
  - XML - as defined in RFC6020
  - JSON - as defined in draft-lhotka-netmod-yang-json-01
  - support for third-party generators processing YANG models.

Architecture

YANG Tools project consists of following logical subsystems:

- **Commons** - Set of general purpose code, which is not specific to YANG, but is also useful outside YANG Tools implementation.
- **YANG Model and Parser** - YANG semantic model and lexical and semantic parser of YANG models, which creates in-memory cross-referenced represenation of YANG models, which is used by other components to determine their behaviour based on the model.
- **YANG Data** - Definition of Normalized Node APIs and Data Tree APIs, reference implementation of these APIs and implementation of XML and JSON codecs for Normalized Nodes.
- **YANG Maven Plugin** - Maven plugin which integrates YANG parser into Maven build lifecycle and provides code-generation framework for components, which wants to generate code or other artefacts based on YANG model.

Concepts

Project defines base concepts and helper classes which are project-agnostic and could be used outside of YANG Tools project scope.

Components

- yang-common
- yang-data-api
- yang-data-codec-gson
- yang-data-codec-xml
- yang-data-impl
- yang-data-jaxen
- yang-data-transform
- yang-data-util
- yang-maven-plugin
- yang-maven-plugin-it
- yang-maven-plugin-spi
- yang-model-api
- yang-model-export
- yang-model-util
- yang-parser-api
- yang-parser-impl

YANG Model API

Class diagram of yang model API

YANG Parser

Yang Statement Parser works on the idea of statement concepts as defined in RFC6020, section 6.3. We come up here with basic ModelStatement and StatementDefinition, following RFC6020 idea of having sequence of statements, where every statement contains keyword and zero or one argument. ModelStatement is extended by DeclaredStatement (as it comes from source, e.g. YANG source) and EffectiveStatement, which contains other substatements and tends to represent result of semantic processing of other statements (uses, augment for YANG). IdentifierNamespace represents common superclass for YANG model namespaces.

Input of the Yang Statement Parser is a collection of StatementStreamSource objects. StatementStreamSource interface is used for inference of effective model and is required to emit its statements using supplied StatementWriter. Each source (e.g. YANG source) has to be processed in three steps in order to emit different statements for each step. This package provides support for various namespaces used across statement parser in order to map relations during declaration phase process.

Currently, there are two implementations of StatementStreamSource in Yangtools:

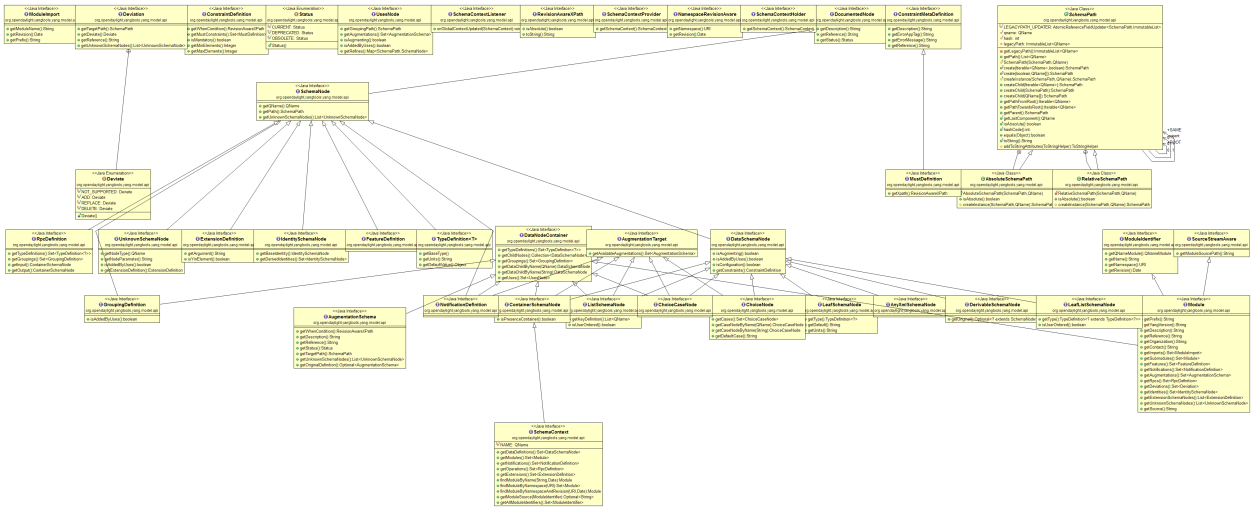

Fig. 69: YANG Model API

- YangStatementSourceImpl - intended for yang sources
- YinStatementSourceImpl - intended for yin sources

YANG Data API

Class diagram of yang data API

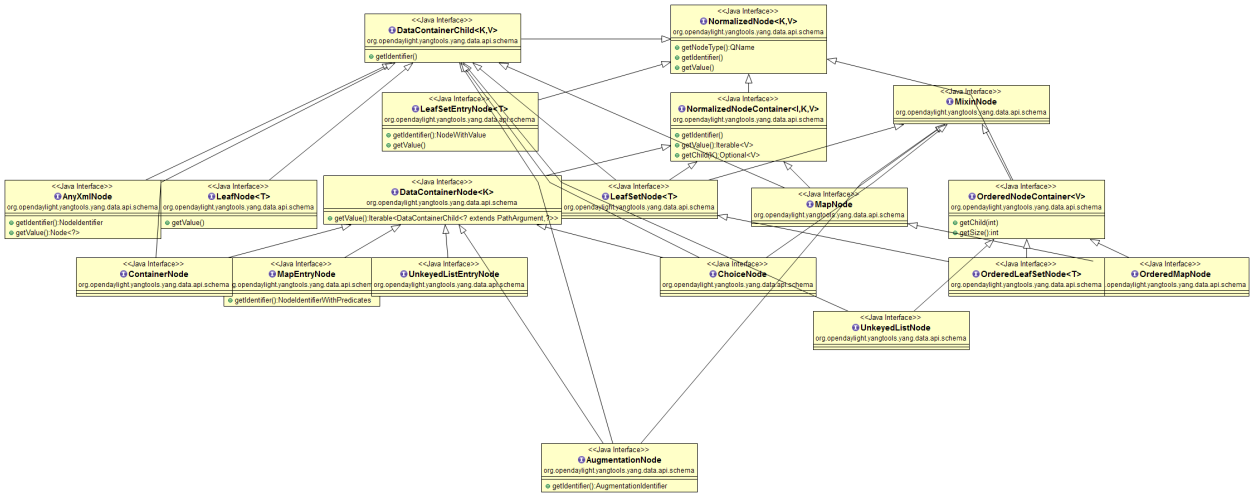

Fig. 70: YANG Data API

YANG Data Codecs

Codecs which enable serialization of NormalizedNodes into YANG-modeled data in XML or JSON format and deserialization of YANG-modeled data in XML or JSON format into NormalizedNodes.

YANG Maven Plugin

Maven plugin which integrates YANG parser into Maven build lifecycle and provides code-generation framework for components, which wants to generate code or other artefacts based on YANG model.

How to / Tutorials

Working with YANG Model

First thing you need to do if you want to work with YANG models is to instantiate a SchemaContext object. This object type describes one or more parsed YANG modules.

In order to create it you need to utilize YANG statement parser which takes one or more StatementStreamSource objects as input and then produces the SchemaContext object.

StatementStreamSource object contains the source file information. It has two implementations, one for YANG sources - YangStatementSourceImpl, and one for YIN sources - YinStatementSourceImpl.

Here is an example of creating StatementStreamSource objects for YANG files, providing them to the YANG statement parser and building the SchemaContext:

```
StatementStreamSource yangModuleSource == new YangStatementSourceImpl("/example.yang",
↪ false);
StatementStreamSource yangModuleSource2 == new YangStatementSourceImpl("/example2.yang
↪", false);

CrossSourceStatementReactor.BuildAction reactor == YangInferencePipeline.RFC6020_
↪REACTOR.newBuild();
reactor.addSources(yangModuleSource, yangModuleSource2);

SchemaContext schemaContext == reactor.buildEffective();
```

First, StatementStreamSource objects with two constructor arguments should be instantiated: path to the yang source file (which is a regular String object) and a boolean which determines if the path is absolute or relative.

Next comes the initiation of new yang parsing cycle - which is represented by CrossSourceStatementReactor.BuildAction object. You can get it by calling method newBuild() on CrossSourceStatementReactor object (RFC6020_REACTOR) in YangInferencePipeline class.

Then you should feed yang sources to it by calling method addSources() that takes one or more StatementStreamSource objects as arguments.

Finally you call the method buildEffective() on the reactor object which returns EffectiveSchemaContext (that is a concrete implementation of SchemaContext). Now you are ready to work with contents of the added yang sources.

Let us explain how to work with models contained in the newly created SchemaContext. If you want to get all the modules in the schemaContext, you have to call method getModules() which returns a Set of modules. If you want to get all the data definitions in schemaContext, you need to call method getDataDefinitions, etc.

```
Set<Module> modules == schemaContext.getModules();
Set<DataSchemaNodes> dataSchemaNodes == schemaContext.getDataDefinitions();
```

Usually you want to access specific modules. Getting a concrete module from SchemaContext is a matter of calling one of these methods:

- findModuleByName(),
- findModuleByNamespace(),

- findModuleByNamespaceAndRevision().

In the first case, you need to provide module name as it is defined in the yang source file and module revision date if it specified in the yang source file (if it is not defined, you can just pass a null value). In order to provide the revision date in proper format, you can use a utility class named SimpleDateFormatUtil.

```
Module exampleModule == schemaContext.findModuleByName("example-module", null);
// or
Date revisionDate == SimpleDateFormatUtil.getRevisionFormat().parse("2015-09-02");
Module exampleModule == schemaContext.findModuleByName("example-module",␣
↪revisionDate);
```

In the second case, you have to provide module namespace in form of an URI object.

```
Module exampleModule == schema.findModuleByNamespace(new URI("opendaylight.org/
↪example-module"));
```

In the third case, you provide both module namespace and revision date as arguments.

Once you have a Module object, you can access its contents as they are defined in YANG Model API. One way to do this is to use method like getIdentities() or getRpcs() which will give you a Set of objects. Otherwise you can access a DataSchemaNode directly via the method getDataChildByName() which takes a QName object as its only argument. Here are a few examples.

```
Set<AugmentationSchema> augmentationSchemas == exampleModule.getAugmentations();
Set<ModuleImport> moduleImports == exampleModule.getImports();

ChoiceSchemaNode choiceSchemaNode == (ChoiceSchemaNode) exampleModule.
↪getDataChildByName(QName.create(exampleModule.getQNameModule(), "example-choice"));

ContainerSchemaNode containerSchemaNode == (ContainerSchemaNode) exampleModule.
↪getDataChildByName(QName.create(exampleModule.getQNameModule(), "example-container
↪"));
```

The YANG statement parser can work in three modes:

- default mode
- mode with active resolution of if-feature statements
- mode with active semantic version processing

The default mode is active when you initialize the parsing cycle as usual by calling the method newBuild() without passing any arguments to it. The second and third mode can be activated by invoking the newBuild() with a special argument. You can either activate just one of them or both by passing proper arguments. Let us explain how these modes work.

Mode with active resolution of if-features makes yang statements containing an if-feature statement conditional based on the supported features. These features are provided in the form of a QName-based java.util.function.Predicate object. In the example below, only two features are supported: example-feature-1 and example-feature-2. The Predicate which contains this information is passed to the method newBuild() and the mode is activated.

```
Predicate<QName> isFeatureSupported == qName -> {
    Set<QName> supportedFeatures == new HashSet<>();
    supportedFeatures.add(QName.create("example-namespace", "2016-08-31", "example-
↪feature-1"));
    supportedFeatures.add(QName.create("example-namespace", "2016-08-31", "example-
↪feature-2"));
    return supportedFeatures.contains(qName);
```

(continues on next page)

(continued from previous page)

```
}

CrossSourceStatementReactor.BuildAction reactor == YangInferencePipeline.RFC6020_
↪REACTOR.newBuild(isFeatureSupported);
```

In case when no features should be supported, you should provide a Predicate<QName> object whose test() method just returns false.

```
Predicate<QName> isFeatureSupported == qName -> false;

CrossSourceStatementReactor.BuildAction reactor == YangInferencePipeline.RFC6020_
↪REACTOR.newBuild(isFeatureSupported);
```

When this mode is not activated, all features in the processed YANG sources are supported.

Mode with active semantic version processing changes the way how YANG import statements work - each module import is processed based on the specified semantic version statement and the revision-date statement is ignored. In order to activate this mode, you have to provide StatementParserMode.SEMVER_MODE enum constant as argument to the method newBuild().

```
CrossSourceStatementReactor.BuildAction reactor == YangInferencePipeline.RFC6020_
↪REACTOR.newBuild(StatementParserMode.SEMVER_MODE);
```

Before you use a semantic version statement in a YANG module, you need to define an extension for it so that the YANG statement parser can recognize it.

In the example above, you see a YANG module which defines semantic version as an extension. This extension can be imported to other modules in which we want to utilize the semantic versioning concept.

Below is a simple example of the semantic versioning usage. With semantic version processing mode being active, the foo module imports the bar module based on its semantic version. Notice how both modules import the module with the semantic-version extension.

Every semantic version must have the following form: x.y.z. The x corresponds to a major version, the y corresponds to a minor version and the z corresponds to a patch version. If no semantic version is specified in a module or an import statement, then the default one is used - 0.0.0.

A major version number of 0 indicates that the model is still in development and is subject to change.

Following a release of major version 1, all modules will increment major version number when backwards incompatible changes to the model are made.

The minor version is changed when features are added to the model that do not impact current clients use of the model.

The patch version is incremented when non-feature changes (such as bugfixes or clarifications of human-readable descriptions that do not impact model functionality) are made that maintain backwards compatibility.

When importing a module with activated semantic version processing mode, only the module with the newest (highest) compatible semantic version is imported. Two semantic versions are compatible when all of the following conditions are met:

- the major version in the import statement and major version in the imported module are equal. For instance, 1.5.3 is compatible with 1.5.3, 1.5.4, 1.7.2, etc., but it is not compatible with 0.5.2 or 2.4.8, etc.
- the combination of minor version and patch version in the import statement is not higher than the one in the imported module. For instance, 1.5.2 is compatible with 1.5.2, 1.5.4, 1.6.8 etc. In fact, 1.5.2 is also compatible with versions like 1.5.1, 1.4.9 or 1.3.7 as they have equal major version. However, they will not be imported because their minor and patch version are lower (older).

If the import statement does not specify a semantic version, then the default one is chosen - 0.0.0. Thus, the module is imported only if it has a semantic version compatible with the default one, for example 0.0.0, 0.1.3, 0.3.5 and so on.

Working with YANG Data

If you want to work with YANG Data you are going to need NormalizedNode objects that are specified in the YANG Data API. NormalizedNode is an interface at the top of the YANG Data hierarchy. It is extended through sub-interfaces which define the behaviour of specific NormalizedNode types like AnyXmlNode, ChoiceNode, LeafNode, ContainerNode, etc. Concrete implemenations of these interfaces are defined in yang-data-impl module. Once you have one or more NormalizedNode instances, you can perform CRUD operations on YANG data tree which is an in-memory database designed to store normalized nodes in a tree-like structure.

In some cases it is clear which NormalizedNode type belongs to which yang statement (e.g. AnyXmlNode, ChoiceNode, LeafNode). However, there are some normalized nodes which are named differently from their yang counterparts. They are listed below:

- LeafSetNode - leaf-list
- OrderedLeafSetNode - leaf-list that is ordered-by user
- LeafSetEntryNode - concrete entry in a leaf-list
- MapNode - keyed list
- OrderedMapNode - keyed list that is ordered-by user
- MapEntryNode - concrete entry in a keyed list
- UnkeyedListNode - unkeyed list
- UnkeyedListEntryNode - concrete entry in an unkeyed list

In order to create a concrete NormalizedNode object you can use the utility class Builders or ImmutableNodes. These classes can be found in yang-data-impl module and they provide methods for building each type of normalized node. Here is a simple example of building a normalized node:

```
\\ example 1
ContainerNode containerNode == Builders.containerBuilder().withNodeIdentifier(new␣
↪YangInstanceIdentifier.NodeIdentifier(QName.create(moduleQName, "example-container
↪")).build();

\\ example 2
ContainerNode containerNode2 == Builders.containerBuilder(containerSchemaNode).
↪build();
```

Both examples produce the same result. NodeIdentifier is one of the four types of YangInstanceIdentifier (these types are described in the javadoc of YangInstanceIdentifier). The purpose of YangInstanceIdentifier is to uniquely identify a particular node in the data tree. In the first example, you have to add NodeIdentifier before building the resulting node. In the second example it is also added using the provided ContainerSchemaNode object.

ImmutableNodes class offers similar builder methods and also adds an overloaded method called fromInstanceId() which allows you to create a NormalizedNode object based on YangInstanceIdentifier and SchemaContext. Below is an example which shows the use of this method.

```
YangInstanceIdentifier.NodeIdentifier contId == new YangInstanceIdentifier.
↪NodeIdentifier(QName.create(moduleQName, "example-container");

NormalizedNode<?, ?> contNode == ImmutableNodes.fromInstanceId(schemaContext,␣
↪YangInstanceIdentifier.create(contId));
```

Let us show a more complex example of creating a NormalizedNode. First, consider the following YANG module:

In the following example, two normalized nodes based on the module above are written to and read from the data tree.

```
TipProducingDataTree inMemoryDataTree ==     InMemoryDataTreeFactory.getInstance().
↪create(TreeType.OPERATIONAL);
inMemoryDataTree.setSchemaContext(schemaContext);

// first data tree modification
MapEntryNode parentOrderedListEntryNode == Builders.mapEntryBuilder().
↪withNodeIdentifier(
new YangInstanceIdentifier.NodeIdentifierWithPredicates(
parentOrderedListQName, parentKeyLeafQName, "pkval1"))
.withChild(Builders.leafBuilder().withNodeIdentifier(
new YangInstanceIdentifier.NodeIdentifier(parentOrdinaryLeafQName))
.withValue("plfval1").build()).build();

OrderedMapNode parentOrderedListNode == Builders.orderedMapBuilder().
↪withNodeIdentifier(
new YangInstanceIdentifier.NodeIdentifier(parentOrderedListQName))
.withChild(parentOrderedListEntryNode).build();

ContainerNode parentContainerNode == Builders.containerBuilder().withNodeIdentifier(
new YangInstanceIdentifier.NodeIdentifier(parentContainerQName))
.withChild(Builders.containerBuilder().withNodeIdentifier(
new NodeIdentifier(childContainerQName)).withChild(parentOrderedListNode).build()).
↪build();

YangInstanceIdentifier path1 == YangInstanceIdentifier.of(parentContainerQName);

DataTreeModification treeModification == inMemoryDataTree.takeSnapshot().
↪newModification();
treeModification.write(path1, parentContainerNode);

// second data tree modification
MapEntryNode childOrderedListEntryNode == Builders.mapEntryBuilder().
↪withNodeIdentifier(
new YangInstanceIdentifier.NodeIdentifierWithPredicates(
childOrderedListQName, childKeyLeafQName, "chkval1"))
.withChild(Builders.leafBuilder().withNodeIdentifier(
new YangInstanceIdentifier.NodeIdentifier(childOrdinaryLeafQName))
.withValue("chlfval1").build()).build();

OrderedMapNode childOrderedListNode == Builders.orderedMapBuilder().
↪withNodeIdentifier(
new YangInstanceIdentifier.NodeIdentifier(childOrderedListQName))
.withChild(childOrderedListEntryNode).build();

ImmutableMap.Builder<QName, Object> builder == ImmutableMap.builder();
ImmutableMap<QName, Object> keys == builder.put(parentKeyLeafQName, "pkval1").build();

YangInstanceIdentifier path2 == YangInstanceIdentifier.of(parentContainerQName).
↪node(childContainerQName)
.node(parentOrderedListQName).node(new␣
↪NodeIdentifierWithPredicates(parentOrderedListQName, keys)).
↪node(childOrderedListQName);

treeModification.write(path2, childOrderedListNode);
```

(continues on next page)

(continued from previous page)

```
treeModification.ready();
inMemoryDataTree.validate(treeModification);
inMemoryDataTree.commit(inMemoryDataTree.prepare(treeModification));

DataTreeSnapshot snapshotAfterCommits == inMemoryDataTree.takeSnapshot();
Optional<NormalizedNode<?, ?>> readNode == snapshotAfterCommits.readNode(path1);
Optional<NormalizedNode<?, ?>> readNode2 == snapshotAfterCommits.readNode(path2);
```

First comes the creation of in-memory data tree instance. The schema context (containing the model mentioned above) of this tree is set. After that, two normalized nodes are built. The first one consists of a parent container, a child container and a parent ordered list which contains a key leaf and an ordinary leaf. The second normalized node is a child ordered list that also contains a key leaf and an ordinary leaf.

In order to add a child node to a node, method withChild() is used. It takes a NormalizedNode as argument. When creating a list entry, YangInstanceIdentifier.NodeIdentifierWithPredicates should be used as its identifier. Its arguments are the QName of the list, QName of the list key and the value of the key. Method withValue() specifies a value for the ordinary leaf in the list.

Before writing a node to the data tree, a path (YangInstanceIdentifier) which determines its place in the data tree needs to be defined. The path of the first normalized node starts at the parent container. The path of the second normalized node points to the child ordered list contained in the parent ordered list entry specified by the key value "pkval1".

Write operation is performed with both normalized nodes mentioned earlier. It consist of several steps. The first step is to instantiate a DataTreeModification object based on a DataTreeSnapshot. DataTreeSnapshot gives you the current state of the data tree. Then comes the write operation which writes a normalized node at the provided path in the data tree. After doing both write operations, method ready() has to be called, marking the modification as ready for application to the data tree. No further operations within the modification are allowed. The modification is then validated - checked whether it can be applied to the data tree. Finally we commit it to the data tree.

Now you can access the written nodes. In order to do this, you have to create a new DataTreeSnapshot instance and call the method readNode() with path argument pointing to a particular node in the tree.

Serialization / deserialization of YANG Data

If you want to deserialize YANG-modeled data which have the form of an XML document, you can use the XML parser found in the module yang-data-codec-xml. The parser walks through the XML document containing YANG-modeled data based on the provided SchemaContext and emits node events into a NormalizedNodeStreamWriter. The parser disallows multiple instances of the same element except for leaf-list and list entries. The parser also expects that the YANG-modeled data in the XML source are wrapped in a root element. Otherwise it will not work correctly.

Here is an example of using the XML parser.

```
InputStream resourceAsStream == ExampleClass.class.getResourceAsStream("/example-
↪module.yang");

XMLInputFactory factory == XMLInputFactory.newInstance();
XMLStreamReader reader == factory.createXMLStreamReader(resourceAsStream);

NormalizedNodeResult result == new NormalizedNodeResult();
NormalizedNodeStreamWriter streamWriter == ImmutableNormalizedNodeStreamWriter.
↪from(result);

XmlParserStream xmlParser == XmlParserStream.create(streamWriter, schemaContext);
xmlParser.parse(reader);

NormalizedNode<?, ?> transformedInput == result.getResult();
```

The XML parser utilizes the javax.xml.stream.XMLStreamReader for parsing an XML document. First, you should create an instance of this reader using XMLInputFactory and then load an XML document (in the form of InputStream object) into it.

In order to emit node events while parsing the data you need to instantiate a NormalizedNodeStreamWriter. This writer is actually an interface and therefore you need to use a concrete implementation of it. In this example it is the ImmutableNormalizedNodeStreamWriter, which constructs immutable instances of NormalizedNodes.

There are two ways how to create an instance of this writer using the static overloaded method from(). One version of this method takes a NormalizedNodeResult as argument. This object type is a result holder in which the resulting NormalizedNode will be stored. The other version takes a NormalizedNodeContainerBuilder as argument. All created nodes will be written to this builder.

Next step is to create an instance of the XML parser. The parser itself is represented by a class named XmlParserStream. You can use one of two versions of the static overloaded method create() to construct this object. One version accepts a NormalizedNodeStreamWriter and a SchemaContext as arguments, the other version takes the same arguments plus a SchemaNode. Node events are emitted to the writer. The SchemaContext is used to check if the YANG data in the XML source comply with the provided YANG model(s). The last argument, a SchemaNode object, describes the node that is the parent of nodes defined in the XML data. If you do not provide this argument, the parser sets the SchemaContext as the parent node.

The parser is now ready to walk through the XML. Parsing is initiated by calling the method parse() on the XmlParserStream object with XMLStreamReader as its argument.

Finally you can access the result of parsing - a tree of NormalizedNodes containg the data as they are defined in the parsed XML document - by calling the method getResult() on the NormalizedNodeResult object.

Introducing schema source repositories

Writing YANG driven generators

Introducing specific extension support for YANG parser

Diagnostics

YANG-PUSH Developer Guide

Overview

The YANG PUBSUB project allows subscriptions to be placed on targeted subtrees of YANG datastores residing on remote devices. Changes in YANG objects within the remote subtree can be pushed to an OpenDaylight controller as specified without a requiring the controller to make a continuous set of fetch requests.

YANG-PUSH capabilities available

This module contains the base code which embodies the intent of YANG-PUSH requirements for subscription as defined in {i2rs-pub-sub-requirements} [https://datatracker.ietf.org/doc/draft-ietf-i2rs-pub-sub-requirements/]. The mechanism for delivering on these YANG-PUSH requirements over Netconf transport is defined in {netconf-yang-push} [netconf-yang-push: https://tools.ietf.org/html/draft-ietf-netconf-yang-push-00].

Note that in the current release, not all capabilities of draft-ietf-netconf-yang-push are realized. Currently only implemented is **create-subscription** RPC support from ietf-datastore-push@2015-10-15.yang; and this will be for periodic subscriptions only. There of course is intent to provide much additional functionality in future OpenDaylight releases.

Future YANG-PUSH capabilities

Over time, the intent is to flesh out more robust capabilities which will allow OpenDaylight applications to subscribe to YANG-PUSH compliant devices. Capabilities for future releases will include:

Support for subscription change/delete: **modify-subscription** rpc support for all mountpoint devices or particular mountpoint device **delete-subscription** rpc support for all mountpoint devices or particular mountpoint device

Support for static subscriptions: This will enable the receipt of subscription updates pushed from publishing devices where no signaling from the controller has been used to establish the subscriptions.

Support for additional transports: NETCONF is not the only transport of interest to OpenDaylight or the subscribed devices. Over time this code will support Restconf and HTTP/2 transport requirements defined in {netconf-restconf-yang-push} [https://tools.ietf.org/html/draft-voit-netconf-restconf-yang-push-01]

YANG-PUSH Architecture

The code architecture of Yang push consists of two main elements

YANGPUSH Provider YANGPUSH Listener

YANGPUSH Provider receives create-subscription requests from applications and then establishes/registers the corresponding listener which will receive information pushed by a publisher. In addition, YANGPUSH Provider also invokes an augmented OpenDaylight create-subscription RPC which enables applications to register for notification as per rfc5277. This augmentation adds periodic time period (duration) and subscription-id values to the existing RPC parameters. The Java package supporting this capability is "org.opendaylight.yangpush.impl". Below class supports the YANGPUSH Provider capability:

(1) YangpushDomProvider The Binding Independent version. It uses a neutral data Document Object Model format for data and API calls, which is independent of any generated Java language bindings from the YANG model.

The YANGPUSH Listener accepts update notifications from a device after they have been de-encapsulated from the NETCONF transport. The YANGPUSH Listener then passes these updates to MD-SAL. This function is implemented via the YangpushDOMNotificationListener class within the "org.opendaylight.yangpush.listner" Java package.

Key APIs and Interfaces

YangpushDomProvider

Central to this is onSessionInitiated which acquires the Document Object Model format based versions of MD-SAL services, including the MountPoint service and RPCs. Via these acquired services, invoke registerDataChangeListener over in YangpushDOMNotificationListener.

YangpushDOMNotificationListener

This API handles instances of a received Push Updates which are inbound to the listener and places these in MD-SAL. Key classes in include:

onPushUpdate Converts and validates the encoding of the pushed subscription update. If the subscription exists and is active, calls updateDataStoreForPushUpdate so that the information can be put in MD-SAL. Finally logs the pushed subscription update as well as some additional context information.

updateDataStoreForPushUpdate Used to put the published information into MD-SAL. This pushed information will also include elements such as the subscription-id, the identity of the publisher, the time of the update, the incoming encoding type, and the pushed YANG subtree information.

YangpushDOMNotificationListener Starts the listener tracking a new Subscription ID from a particular publisher.

API Reference Documentation

Javadocs are generated while creating mvn:site and they are located in target/ directory in each module.

CHAPTER 3

Content for OpenDaylight Contributors

The following content is intended for developers who either currently participate in the development of OpenDaylight or would like to start.

3.1 Gerrit Guide

3.1.1 How to push to Gerrit

It is highly recommended to use ssh to push to Gerrit to push code to Gerrit. In the event that you cannot use ssh such as corporate firewall blocking you then falling back to pushing via https should work.

Using ssh to push to Gerrit

TODO

Using https to push to Gerrit

Gerrit does not allow you to use your regular account credentials when pushing via https. Instead it requires you to first generate a http password via the Web U and use that as the password when pushing via https.

To do this:

1. navigate to https://git.opendaylight.org/gerrit/#/settings/http-password (Steps 1, 2 and 3 in the image above.)
2. click the **Generate Password** button.

Gerrit will then generate a random password which you will need to use as your password when using git to push code to Gerrit via https.

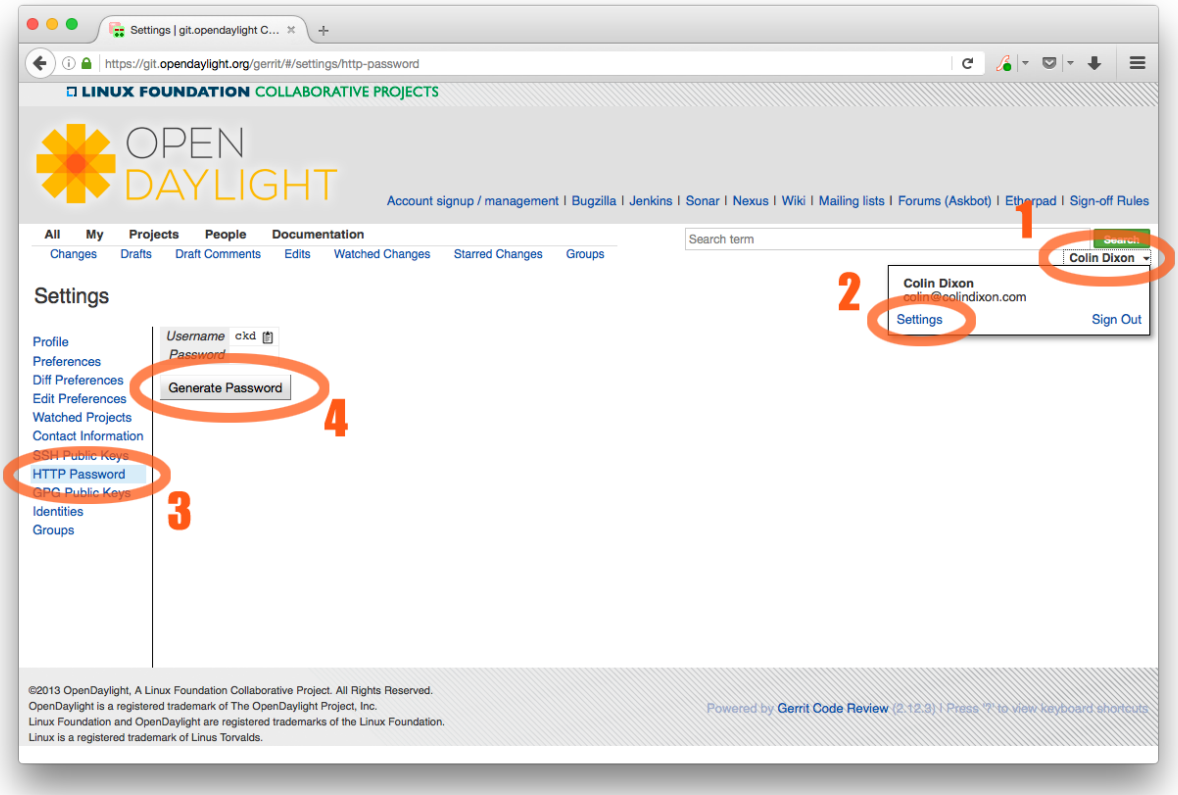

Fig. 1: Setting up an https password to push using https instead of ssh.

3.1.2 Signing Gerrit Commits

1. Generate your GPG key.

   The following instructions work on a Mac, but the general approach should be the same on other OSes.

```
brew install gpg2  # If you don't have homebrew, get that here: http://brew.sh/
gpg2 --gen-key
# pick 1 for "RSA and RSA"
# enter 4096 to creat a 4096-bit key
# enter an expiration time, I picked 2y for 2 years
# enter y to accept the expiration time
# pick O or Q to accept your name/email/comment
# enter a pass phrase twice. it seems like backspace doesn't work, so type␣
↪carefully
gpg2 --fingerprint
# you'll get something like this:
# spectre:~ ckd$ gpg2 --fingerprint
# /Users/ckd/.gnupg/pubring.gpg
# ------------------------------
# pub   4096R/F566C9B1 2015-04-06 [expires: 2017-04-05]
#       Key fingerprint = 7C37 02AC D651 1FA7 9209  48D3 5DD5 0C4B F566 C9B1
# uid       [ultimate] Colin Dixon <colin at colindixon.com>
# sub   4096R/DC1497E1 2015-04-06 [expires: 2017-04-05]
# you're looking for the part after 4096R, which is your key ID
gpg2 --send-keys $KEY_ID
# in the above example, the $KEY_ID would be F566C9B1
# you should see output like this:
# gpg: sending key F566C9B1 to hkp server keys.gnupg.net
```

   If you're trying to participate in an OpenDaylight keysigning, then send the output of `gpg2 --fingerprint $KEY_ID` to keysigning@opendaylight.org

```
gpg2 --fingerprint $KEY_ID
# in the above example, the $KEY_ID would be F566C9B1
# in my case, the output was:
# pub   4096R/F566C9B1 2015-04-06 [expires: 2017-04-05]
#       Key fingerprint = 7C37 02AC D651 1FA7 9209  48D3 5DD5 0C4B F566 C9B1
# uid       [ultimate] Colin Dixon <colin at colindixon.com>
# sub   4096R/DC1497E1 2015-04-06 [expires: 2017-04-05]
```

2. Install gpg, instead of or addition to gpg2. It appears as though gpg2 has annoying things that it does when asking for your passphrase, which I haven't debugged yet.

   Note: you can tell git to use gpg by doing: `git config --global gpg.program gpg2` but that then will seem to struggle asking for your passphrase unless you have your gpg-agent set up right.

3. Add you GPG to Gerrit

   (a) Run the following at the CLI:

```
gpg --export -a $FINGER_PRINT
# e.g., gpg --export -a F566C9B1
# in my case the output looked like:
# -----BEGIN PGP PUBLIC KEY BLOCK-----
# Version: GnuPG v2
#
```

(continues on next page)

(continued from previous page)

```
# mQINBFUisGABEAC/DkcjNUhxQkRLdfbfdlq9NlfDusWri0cXLVz4YN1cTUTF5HiW
# ...
# gJT+FwDvCGgaE+JGlmXgjv0WSd4f9cNXkgYqfb6mpji0F3TF2HXXiVPqbwJ1V3I2
# NA+l+/koCW0aMReK
# =A/ql
# -----END PGP PUBLIC KEY BLOCK-----
```

(b) Browse to https://git.opendaylight.org/gerrit/#/settings/gpg-keys

(c) Click Add Key...

(d) Copy the output from the above command, paste it into the box, and click Add

4. Set up your git to sign commits and push signatures

```
git config commit.gpgsign true
git config push.gpgsign true
git config user.signingkey $FINGER_PRINT
# e.g., git config user.signingkey F566C9B1
```

Note: you can do this instead with `git commit -S` You can use `git commit -S` and `git push --signed` on the CLI instead of configuring it in config if you want to control which commits use your signature.

---

5. Commit and push a change

(a) change a file

(b) `git commit -asm "test commit"`

Note: this should result in git asking you for your passphrase

---

(c) `git review`

Note: this should result in git asking you for your passphrase

---

Note: annoyingly, the presence of a gpgp signature or pushing of a gpg signature isn't recognized as a "change" by Gerrit, so if you forget to do either, you need to change something about the commit to get Gerrit to accept the patch again. Slightly tweaking the commit message is a good way.

Note: this assumes you have git review set up and push.gpgsign set to true. Otherwise:

`git push --signed gerrit HEAD:refs/for/master`

Note: this assumes you have your gerrit remote set up, if not it's something like: `ssh://ckd@git.opendaylight.org:29418/<repo>.git` where repo is something like docs or controller

---

---

6. Verify that your commit is signed by going to the change in Gerrit and checking for a green check (instead of a blue ?) next to your name.

Setting up gpg-agent on a Mac

1. Install `gpg-agent` and `pinentry-mac` using `brew`:

```
brew install gpg-agent pinentry-mac
```

2. Edit your `~/.gnupg/gpg.conf` contain the line:

```
use-agent
```

3. Edit your `~/.gnupg/gpg-agent.conf` to something like:

```
use-standard-socket
enable-ssh-support
default-cache-ttl 600
max-cache-ttl 7200
pinentry-program /usr/local/bin/pinentry-mac
```

4. Edit your `.bash_profile` or equivalent file to contain the following:

```
[ -f ~/.gpg-agent-info ] && source ~/.gpg-agent-info
if [ -S "${GPG_AGENT_INFO%%:*}" ]; then
  export GPG_AGENT_INFO
else
  eval $( gpg-agent --daemon --write-env-file ~/.gpg-agent-info )
fi
```

5. Kill any stray `gpg-agent` daemons running:

```
sudo killall gpg-agent
```

6. Restart your terminal (or log in and out) to reload the your `.bash_profile` or equivalent file
7. The next time a git operation makes a call to gpg, it should use your gpg-agent to run a GUI window to ask for your passphrase and give you an option to save your passphrase in the keychain.

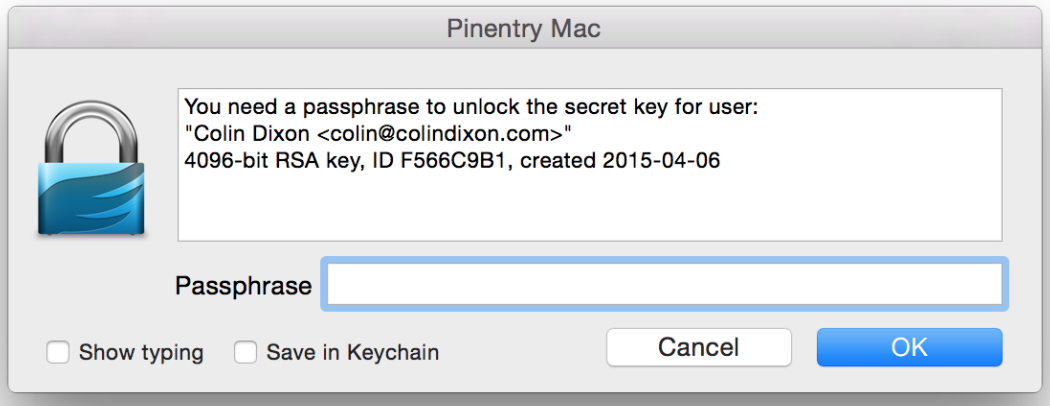

3.2 Infrastructure Guide

This guide provides details into OpenDaylight Infrastructure and services.

Contents:

3.2.1 Jenkins

The Release Engineering Project consolidates the Jenkins jobs from project-specific VMs to a single Jenkins server. Each OpenDaylight project has a tab for their jobs on the jenkins-master. The system utilizes Jenkins Job Builder for the creation and management of the Jenkins jobs.

Sections:

- *New Project Quick Start*
- *Jenkins Master*
- *Build Minions*
  - *Adding New Components to the Minions*
  - *Flavors*
  - *Pool: ODLVEX*
  - *Pool: ODLVEX - HOT (Heat Orchestration Templates)*
- *Creating Jenkins Jobs*
- *Getting Jenkins Job Builder*
- *Installing Jenkins Job Builder*
- *Virtual Environments*
- *Installing JJB using pip*
  - *Updating releng/builder repo or global-jjb*
- *Installing JJB Manually*
- *Jenkins Job Templates*
- *Maven Properties*
- *Jenkins Sandbox*
  - *Notes Regarding the Sandbox*
  - *Configuration*
  - *Manual Method*
  - *Testing Jobs*
  - *Pushing Jobs*
  - *Running Jobs*

New Project Quick Start

This section attempts to provide details on how to get going as a new project quickly with minimal steps. The rest of the guide should be read and understood by those who need to create and contribute new job types that is not already covered by the existing job templates provided by OpenDaylight's JJB repo.

As a new project you will be mainly interested in getting your jobs to appear in the jenkins-master silo and this can be achieved by simply creating a <project>.yaml in the releng/builder project's jjb directory.

```
git clone --recursive https://git.opendaylight.org/gerrit/releng/builder
cd builder
mkdir jjb/<new-project>
```

Where <new-project> should be the same name as your project's git repo in Gerrit. If your project is called "aaa" then create a new jjb/aaa directory.

Next we will create <new-project>.yaml as follows:

```
---
- project:
    name: <NEW_PROJECT>-carbon
    jobs:
      - '{project-name}-clm-{stream}'
      - '{project-name}-integration-{stream}'
      - '{project-name}-merge-{stream}'
      - '{project-name}-verify-{stream}-{maven}-{jdks}'

    project: '<NEW_PROJECT>'
    project-name: '<NEW_PROJECT>'
    stream: carbon
    branch: 'master'
    jdk: openjdk8
    jdks:
      - openjdk8
    maven:
      - mvn33:
          mvn-version: 'mvn33'
    mvn-settings: '<NEW_PROJECT>-settings'
    mvn-goals: 'clean install -Dmaven.repo.local=/tmp/r -Dorg.ops4j.pax.url.mvn.
↪localRepository=/tmp/r'
    mvn-opts: '-Xmx1024m -XX:MaxPermSize=256m'
    dependencies: 'odlparent-merge-{stream},yangtools-merge-{stream},controller-merge-
↪{stream}'
    email-upstream: '[<NEW_PROJECT>] [odlparent] [yangtools] [controller]'
    archive-artifacts: ''

- project:
    name: <NEW_PROJECT>-sonar
    jobs:
      - '{project-name}-sonar'

    project: '<NEW_PROJECT>'
    project-name: '<NEW_PROJECT>'
    branch: 'master'
    mvn-settings: '<NEW_PROJECT>-settings'
    mvn-goals: 'clean install -Dmaven.repo.local=/tmp/r -Dorg.ops4j.pax.url.mvn.
↪localRepository=/tmp/r'
    mvn-opts: '-Xmx1024m -XX:MaxPermSize=256m'
```

Replace all instances of <new-project> with the name of your project. This will create the jobs with the default job types we recommend for Java projects. If your project is participating in the simultanious-release and ultimately will be included in the final distribution, it is required to add the following job types into the job list for the release you are participating.

```
- '{project-name}-distribution-check-{stream}'
- '{project-name}-validate-autorelease-{stream}'
```

If you'd like to explore the additional tweaking options available please refer to the *Jenkins Job Templates* section.

Finally we need to push these files to Gerrit for review by the releng/builder team to push your jobs to Jenkins.

```
git add jjb/<new-project>
git commit -sm "Add <new-project> jobs to Jenkins"
git review
```

This will push the jobs to Gerrit and your jobs will appear in Jenkins once the releng/builder team has reviewed and merged your patch.

Jenkins Master

The jenkins-master is the home for all project's Jenkins jobs. All maintenance and configuration of these jobs must be done via JJB through the releng-builder-repo. Project contributors can no longer edit the Jenkins jobs directly on the server.

Build Minions

The Jenkins jobs are run on build minions (executors) which are created on an as-needed basis. If no idle build minions are available a new VM is brought up. This process can take up to 2 minutes. Once the build minion has finished a job, it will be destroyed.

Our Jenkins master supports many types of dynamic build minions. If you are creating custom jobs then you will need to have an idea of what type of minions are available. The following are the current minion types and descriptions. Minion Template Names are needed for jobs that take advantage of multiple minions as they must be specifically called out by template name instead of label.

Adding New Components to the Minions

If your project needs something added to one of the minions, you can help us get things added faster by doing one of the following:

- Submit a patch to RelEng/Builder for the appropriate *jenkins-scripts* definition which configure software during minion boot up.
- Submit a patch to RelEng/Builder for the *packer/provision* scripts that configures software during minion instance imaging.
- Submit a patch to RelEng/Builder for the Packer's templates in the *packer/templates* directory that configures a new instance definition along with changes in *packer/provision*.

Going the first route will be faster in the short term as we can inspect the changes and make test modifications in the sandbox to verify that it works.

Note: The first route may add additional setup time considering this is run every time the minion is booted.

The second and third routes, however, is better for the community as a whole as it will allow others to utilize our Packer setups to replicate our systems more closely. It is, however, more time consuming as an image snapshot needs to be created based on the updated Packer definitions before it can be attached to the Jenkins configuration on sandbox for validation testing.

In either case, the changes must be validated in the sandbox with tests to make sure that we don't break current jobs and that the new software features are operating as intended. Once this is done the changes will be merged and the updates applied to the RelEng Jenkins production silo. Any changes to files under *releng/builder/packer* will be validated and images would be built triggered by verify-packer and merge-packer jobs.

Please note that the combination of a Packer definitions from *vars*, *templates* and the *provision* scripts is what defines a given minion. For instance, a minion may be defined as *centos7-builder* which is a combination of Packer OS image definitions from *vars/centos.json*, Packer template definitions from *templates/builder.json* and spinup scripts from *provision/builder.sh*. This combination provides the full definition of the realized minion.

Jenkins starts a minion using the latest image which is built and linked into the Jenkins configuration. Once the base instance is online Jenkins checks out the RelEng/Builder repo on it and executes two scripts. The first is *provision/baseline.sh*, which is a baseline for all of the minions.

The second is the specialized script, which handles any system updates, new software installs or extra environment tweaks that don't make sense in a snapshot. Examples could include installing new package or setting up a virtual environment. Its imperative to ensure modifications to these spinup scripts have considered time taken to install the packages, as this could increase the build time for every job which runs on the image. After all of these scripts have executed Jenkins will finally attach the minion as an actual minion and start handling jobs on it.

Flavors

Performance flavors come with dedicated CPUs and are not shared with other accounts in the cloud so should ensure consistent performance.

Table 1: Flavors

Instance Type	CPUs	Memory
odl-standard-1	1	4
odl-standard-2	2	8
odl-standard-4	4	16
odl-standard-8	8	32
odl-standard-16	16	64
odl-highcpu-2	2	2
odl-highcpu-4	4	4
odl-highcpu-8	8	8

Pool: ODLVEX

Pool: ODLVEX - HOT (Heat Orchestration Templates)

HOT integration enables to spin up integration labs servers for CSIT jobs using heat, rathar than using jclouds (deprecated). Image names are updated on the project specific job templates using the variable *{odl,docker,openstack,tools}_system_image* followed by image name in the format <*platform*> - <*template*> - <*date-stamp*>.

Following are the list of published images available to be used with Jenkins jobs.

- ZZCI - CentOS 7 - autorelease - 20180125-2240

- ZZCI - CentOS 7 - builder - 20180109-0417
- ZZCI - CentOS 7 - builder - 20180110-1659
- ZZCI - CentOS 7 - builder - 20180201-2139
- ZZCI - CentOS 7 - devstack - 20171208-1648
- ZZCI - CentOS 7 - devstack-ocata - 20171208-1649
- ZZCI - CentOS 7 - devstack-pike - 20171208-1649
- ZZCI - CentOS 7 - docker - 20171209-0317
- ZZCI - CentOS 7 - docker - 20180109-0346
- ZZCI - CentOS 7 - docker - 20180110-1659
- ZZCI - CentOS 7 - docker - 20180417-0311
- ZZCI - CentOS 7 - java-builder - 20171206-1842
- ZZCI - CentOS 7 - java-builder - 20171209-0032
- ZZCI - CentOS 7 - robot - 20171207-1911
- ZZCI - Ubuntu 14.04 - gbp - 20171208-2336
- ZZCI - Ubuntu 16.04 - gbp - 20171213-2018
- ZZCI - Ubuntu 16.04 - mininet-ovs-25 - 20171208-1847
- ZZCI - Ubuntu 16.04 - mininet-ovs-26 - 20171208-1847
- ZZCI - Ubuntu 16.04 - mininet-ovs-28 - 20180301-1041

Creating Jenkins Jobs

Jenkins Job Builder takes simple descriptions of Jenkins jobs in YAML format and uses them to configure Jenkins.

- Jenkins Job Builder (JJB) documentation
- RelEng/Builder Gerrit
- RelEng/Builder Git repository

Getting Jenkins Job Builder

OpenDaylight uses Jenkins Job Builder to translate our in-repo YAML job configuration into job descriptions suitable for consumption by Jenkins. When testing new Jenkins Jobs in the *Jenkins Sandbox*, you'll need to use the *jenkins-jobs* executable to translate a set of jobs into their XML descriptions and upload them to the sandbox Jenkins server.

We document *installing* *jenkins-jobs* below.

Installing Jenkins Job Builder

We recommend using *pip* to assist with JJB installs, but we also document *installing from a git repository manually*. For both, we recommend using Python *Virtual Environments* to isolate JJB and its dependencies.

The builder/jjb/requirements.txt file contains the currently recommended JJB version. Because JJB is fairly unstable, it may be necessary to debug things by installing different versions. This is documented for both *pip-assisted* and *manual* installs.

Virtual Environments

For both *pip-assisted* and *manual* JJB installs, we recommend using Python Virtual Environments to manage JJB and its Python dependencies. The python-virtualenvwrapper tool can help you do so.

Documentation is available for installing python-virtualenvwrapper. On Linux systems with pip (typical), they amount to:

```
sudo pip install virtualenvwrapper
```

A virtual environment is simply a directory that you install Python programs into and then append to the front of your path, causing those copies to be found before any system-wide versions.

Create a new virtual environment for JJB.

```
# Virtaulenvwrapper uses this dir for virtual environments
$ echo $WORKON_HOME
/home/daniel/.virtualenvs
# Make a new virtual environment
$ mkvirtualenv jjb
# A new venv dir was created
(jjb)$ ls -rc $WORKON_HOME | tail -n 1
jjb
# The new venv was added to the front of this shell's path
(jjb)$ echo $PATH
/home/daniel/.virtualenvs/jjb/bin:<my normal path>
# Software installed to venv, like pip, is found before system-wide copies
(jjb)$ command -v pip
/home/daniel/.virtualenvs/jjb/bin/pip
```

With your virtual environment active, you should install JJB. Your install will be isolated to that virtual environment's directory and only visible when the virtual environment is active.

You can easily leave and return to your venv. Make sure you activate it before each use of JJB.

```
(jjb)$ deactivate
$ command -v jenkins-jobs
# No jenkins-jobs executable found
$ workon jjb
(jjb)$ command -v jenkins-jobs
$WORKON_HOME/jjb/bin/jenkins-jobs
```

Installing JJB using pip

The recommended way to install JJB is via pip.

First, clone the latest version of the releng-builder-repo.

```
$ git clone --recursive https://git.opendaylight.org/gerrit/p/releng/builder.git
```

Before actually installing JJB and its dependencies, make sure you've *created and activated* a virtual environment for JJB.

```
$ mkvirtualenv jjb
```

The recommended version of JJB to install is the version specified in the builder/jjb/requirements.txt file.

```
# From the root of the releng/builder repo
(jjb)$ pip install -r jjb/requirements.txt
```

To validate that JJB was successfully installed you can run this command:

```
(jjb)$ jenkins-jobs --version
```

TODO: Explain that only the currently merged jjb/requirements.txt is supported, other options described below are for troubleshooting only.

To change the version of JJB specified by builder/jjb/requirements.txt to install from the latest commit to the master branch of JJB's git repository:

```
$ cat jjb/requirements.txt
-e git+https://git.openstack.org/openstack-infra/jenkins-job-builder#egg=jenkins-job-
↪builder
```

To install from a tag, like 1.4.0:

```
$ cat jjb/requirements.txt
-e git+https://git.openstack.org/openstack-infra/jenkins-job-builder@1.4.0
↪#egg=jenkins-job-builder
```

Updating releng/builder repo or global-jjb

Follow these steps to update the releng/builder repo. The repo uses a submodule from a global-jjb repo so that common source can be shared across different projects. This requires updating the releng/builder repo periodically to pick up the changes. New versions of jjb could also require updating the releng/builder repo. Follow the previous steps earlier for updating jenkins-jobs using the builder/jjb/requirements.txt file. Ensure that the version listed in the file is the currently supported version, otherwise install a different version or simply upgrade using *pip install –upgrade jenkins-job-builder*.

The example below assumes the user has cloned releng/builder to *~/git/releng/builder*. Update the repo, update the submodules and then submit a test to verify it works.

```
cd ~/git/releng/builder
git checkout master
git pull
git submodule update --init --recursive
jenkins-jobs --conf jenkins.ini test jjb/ netvirt-csit-1node-openstack-queens-
↪upstream-stateful-fluorine
```

Installing JJB Manually

This section documents installing JJB from its manually cloned repository.

Note that *installing via pip* is typically simpler.

Checkout the version of JJB's source you'd like to build.

For example, using master:

```
$ git clone https://git.openstack.org/openstack-infra/jenkins-job-builder
```

Using a tag, like 1.4.0:

```
$ git clone https://git.openstack.org/openstack-infra/jenkins-job-builder
$ cd jenkins-job-builder
$ git checkout tags/1.4.0
```

Before actually installing JJB and its dependencies, make sure you've *created and activated* a virtual environment for JJB.

```
$ mkvirtualenv jjb
```

You can then use JJB's requirements.txt file to install its dependencies. Note that we're not using *sudo* to install as root, since we want to make use of the venv we've configured for our current user.

```
# In the cloned JJB repo, with the desired version of the code checked out
(jjb)$ pip install -r requirements.txt
```

Then install JJB from the repo with:

```
(jjb)$ pip install .
```

To validate that JJB was successfully installed you can run this command:

```
(jjb)$ jenkins-jobs --version
```

Jenkins Job Templates

The OpenDaylight RelEng/Builder project provides jjb-templates that can be used to define basic jobs.

The *Gerrit Trigger* listed in the jobs are keywords that can be used to trigger the job to run manually by simply leaving a comment in Gerrit for the patch you wish to trigger against.

All jobs have a default build-timeout value of 360 minutes (6 hrs) but can be overrided via the opendaylight-infra-wrappers' build-timeout property.

TODO: Group jobs into categories: every-patch, after-merge, on-demand, etc. TODO: Reiterate that "remerge" triggers all every-patch jobs at once, because when only a subset of jobs is triggered, Gerrit forgets valid -1 from jobs outside the subset. TODO: Document that only drafts and commit-message-only edits do not trigger every-patch jobs. TODO: Document test-{project}-{feature} and test-{project}-all.

Maven Properties

We provide a properties which your job can take advantage of if you want to do something different depending on the job type that is run. If you create a profile that activates on a property listed blow. The JJB templated jobs will be able to activate the profile during the build to run any custom code you wish to run in your project.

```
-Dmerge   : This flag is passed in our Merge job and is equivalent to the
            Maven property
            <merge>true</merge>.
-Dsonar   : This flag is passed in our Sonar job and is equivalent to the
            Maven property
            <sonar>true</sonar>.
```

Jenkins Sandbox

The jenkins-sandbox instance's purpose is to allow projects to test their JJB setups before merging their code over to the RelEng master silo. It is configured similarly to the master instance, although it cannot publish artifacts or vote in Gerrit.

If your project requires access to the sandbox please open an OpenDaylight Helpdesk ticket (<helpdesk@opendaylight.org>) and provide your ODL ID.

Notes Regarding the Sandbox

- Jobs are automatically deleted every Saturday at 08:00 UTC
- Committers can login and configure Jenkins jobs in the sandbox directly (unlike with the master silo)
- Sandbox configuration mirrors the master silo when possible
- Sandbox jobs can NOT upload artifacts to Nexus
- Sandbox jobs can NOT vote on Gerrit

Configuration

Make sure you have Jenkins Job Builder [properly installed](#jjb_install).

If you do not already have access, open an OpenDaylight Helpdesk ticket (<helpdesk@opendaylight.org>) to request access to ODL's sandbox instance. Integration/Test (integration-test-wiki) committers have access by default.

JJB reads user-specific configuration from a jenkins.ini. An example is provided by releng/builder at example-jenkins.ini.

```
# If you don't have RelEng/Builder's repo, clone it
$ git clone --recursive https://git.opendaylight.org/gerrit/p/releng/builder.git
# Make a copy of the example JJB config file (in the builder/ directory)
$ cp jenkins.ini.example jenkins.ini
# Edit jenkins.ini with your username, API token and ODL's sandbox URL
$ cat jenkins.ini
<snip>
[jenkins]
user=<your ODL username>
password=<your ODL Jenkins sandbox API token>
url=https://jenkins.opendaylight.org/sandbox
<snip>
```

To get your API token, login to the Jenkins **sandbox** instance (*not the main master Jenkins instance, different tokens*), go to your user page (by clicking on your username, for example), click "Configure" and then "Show API Token".

Manual Method

If you *installed JJB locally into a virtual environment*, you should now activate that virtual environment to access the *jenkins-jobs* executable.

```
$ workon jjb
(jjb)$
```

You'll want to work from the root of the RelEng/Builder repo, and you should have your *jenkins.ini* file [properly configured](#sandbox_config).

Testing Jobs

It's good practice to use the *test* command to validate your JJB files before pushing them.

```
jenkins-jobs --conf jenkins.ini test jjb/ <job-name>
```

If the job you'd like to test is a template with variables in its name, it must be manually expanded before use. For example, the commonly used template *{project}-csit-verify-1node-{functionality}* might expand to *ovsdb-csit-verify-1node-netvirt*.

```
jenkins-jobs --conf jenkins.ini test jjb/ ovsdb-csit-verify-1node-netvirt
```

Successful tests output the XML description of the Jenkins job described by the specified JJB job name.

Pushing Jobs

Once you've *configured your 'jenkins.ini'* and *verified your JJB jobs* produce valid XML descriptions of Jenkins jobs you can push them to the Jenkins sandbox.

Important: When pushing with *jenkins-jobs*, a log message with the number of jobs you're pushing will be issued, typically to stdout. **If the number is greater than 1** (or the number of jobs you passed to the command to push) then you are pushing too many jobs and should **'ctrl+c' to cancel the upload**. Else you will flood the system with jobs.

```
INFO:jenkins_jobs.builder:Number of jobs generated:  1
```

Failing to provide the final '<job-name>' param will push all jobs!

```
# Don't push all jobs by omitting the final param! (ctrl+c to abort)
jenkins-jobs --conf jenkins.ini update jjb/ <job-name>
```

Alternatively, you can push a job to the Jenkins sandbox with a special comment in a releng/builder gerrit patch. The job will be based off of the code your patch is based upon. Meaning, if your patch is changing something related to the job you are pushing, those changes will exist in the sandbox job. The format of the comment is:

```
jjb-deploy <job name>
```

Note: Also note that wildcards can be used in <job name> which will expand all jobs that exist for the pattern.

Running Jobs

Once you have your Jenkins job configuration *pushed to the Sandbox* you can trigger it to run.

Find your newly-pushed job on the Sandbox's web UI. Click on its name to see the job's details.

Make sure you're logged in to the Sandbox.

Click "Build with Parameters" and then "Build".

Wait for your job to be scheduled and run. Click on the job number to see details, including console output.

Make changes to your JJB configuration, re-test, re-push and re-run until your job is ready.

3.2.2 Release Workflow

This page documents the workflow for releasing for projects that are not built and released via the Autorelease project.

Sections:

- *Workflow*
- *Release Job*

Workflow

OpenDaylight uses Nexus as it's artifact repository for releasing artifacts to the world. The workflow involves using Nexus to produce a staging repository which can be tested and reviewed before being approved to copy to the final destination opendaylight.release repo. The workflow in general is as follows:

1. Project create release tag and push to Gerrit
2. Project will contact helpdesk@opendaylight.org with project name and build tag to produce a release candidate / staging repo
3. Helpdesk will run a build and notify project of staging repo location
4. Project tests staging repo and notifies Helpdesk with go ahead to release
5. Helpdesk clicks Release repo button in Nexus
6. (optional) Helpdesk runs Jenkins job to push update-site.zip to p2repos sites repo

Step 6 is only necessary for Eclipse projects that need to additionally deploy an update site to a webserver.

Release Job

There is a JJB template release job which should be used for a project if the project needs to produce a staging repo for release. The supported Job types are listed below, use the one relevant to your project.

Maven|Java {name}-release-java – this job type will produce a staging repo in Nexus for Maven projects.

P2 Publisher {name}-publish-p2repo – this job type is useful for projects that produce a p2 repo that needs to be published to a special URL.

3.3 Integration Testing Guide

The Integration Testing Guide provides details on how to contribute test code to OpenDaylight.

Contents:

3.3.1 System Test Guide

Introduction

This step by step guide aims to help projects with the task of creating a System Test job that runs in Continuous Integration.

A System Test job will normally install a controller distribution in one or more VMs and will run a functionality test using some test tool (e.g. mininet). This job will run periodically, tipically once or twice a day.

All projects defining top-level features (essential functionality) and that have decided to use the OpenDaylight CI for system test must create system test jobs.

System test jobs rely on Robot Framework, this is because Robot FW provides:

- Structure for test creation and execution (e.g. test suites, test cases that PASS/FAIL).
- Easy test debug (real time logs, etc. . .).
- Test reports in Jenkins.

For those projects creating system test, Integration group will provide:

- Robot Framework support and assistance.
- Review of system test code. The code will be pushed to integration/test git (csit/suites/$project/).
- JJB templates to install controller and execute a robot test to verify a project functionality (releng/builder git, jjb/integration/).

Create basic system test

Download Integration/Test Repository:

```
git clone ssh://${USERNAME}@git.opendaylight.org:29418/integration/test.git
cd test
```

Follow the instructions in pulling-and-pushing-the-code to know more about pulling and pushing code.

Create a folder for your project robot test:

```
mkdir test/csit/suites/$project
cd test/csit/suites/$project
```

Replace $project with your project name.

Move your robot suites (test folders) into the project folder:

If you do not have any robot test yet, copy integration basic folder suite into your folder. You can later improve this suite or replace it by your own suites:

```
cp -R test/csit/suites/integration/basic basic
```

This suite will verify Restconf is operational.

Create a test plan

A test plan is a text file indicating which robot test suites (including integration repo path) will be executed to test a project functionality:

```
vim test/csit/testplans/$project-$functionality.txt
```

Replace $project with your project name and $functionality with the functionality you want to test.

If you took the basic test from integration, the test plan file should look like this:

```
# Place the suites in run order:
integration/test/csit/suites/$project/basic
```

Save the changes and exit editor.

Optional: Version specific test plan

Integration/Test is not part of the simultaneous release, so the same suites are used for testing all supported ODL versions. There may be API changes between different releases of ODL, which may require different logic in your Robot tests. If the difference is small, it is recommended to act upon value of ODL_STREAM variable (e.g. "beryllium", "boron", "carbon", etc).

If the difference is big, you may want to use different list of suites in testplan. One way is to define separate jobs with different functionality names. But the more convenient way is to define stream-specific testplan. For example:

```
vim test/csit/testplans/$project-$functionality-boron.txt
```

would contain a list of suites for testing Boron, while $project-$functionality.txt would still contain the default list (used for streams without stream specific testplans).

Optional: Create a script or config plan

Sometimes the environment prepared by scripts in releng/builder is not suitable as is, and there are changes to be done before controller is installed (script plan) or before it is started (config plan). You may create as many bash scripts as you need in test/csit/scripts/ and then list them in the scriplans or configplans folder:

```
vim test/csit/scriptplans/$project-$functionality.txt
```

Save and push Test changes

Add the changes and push them in the integration/test repo:

```
git add -A
git commit -s
git push
```

Create system test job

Download RelEng Builder repository:

```
git clone ssh://${USERNAME}@git.opendaylight.org:29418/releng/builder
cd builder
```

Follow the instructions in pulling-and-pushing-the-code to know more about pulling and pushing code.

Create a new file and modify the values according to your project:

```
vim jjb/$project/$project-csit-$functionality.yaml
```

For a Managed project it should look like this:

```
---
- project:
    name: openflowplugin-csit-flow-services
    jobs:
      - inttest-csit-1node

    # The project name
    project: 'openflowplugin'

    # The functionality under test
    functionality:
      - flow-services
      - gate-flow-services

    # Project branches
    stream:
      - fluorine:
          branch: 'master'
      - oxygen:
          branch: 'stable/oxygen'
      - nitrogen:
          branch: 'stable/nitrogen'
      - carbon:
          branch: 'stable/carbon'
          karaf-version: 'karaf3'

    install:
      - all:
          scope: 'all'

    # Features to install
    install-features: >
        odl-openflowplugin-flow-services-rest,
        odl-openflowplugin-app-table-miss-enforcer,
        odl-openflowplugin-nxm-extensions

    # Robot custom options
    robot-options: ''
```

Explanation:

- name: give some name like $project-csit-$functionality.
- jobs: replace 1node by 3node if your test is develop for 3node cluster.
- project: set your your project name here (e.g. openflowplugin).
- functionality: set the functionality you want to test (e.g. flow-services). Note this has also to match the robot test plan name you defined in the earlier section *create a test plan* (e.g. openflowplugin-flow-services.txt)
- stream: list the project branches you are going to generate system test. Only last branch if the project is new.
- install: this specifies controller installation, 'only' means only features in install-features will be installed, 'all' means all compatible features will be installed on top (multi-project features test).
- install-features: list of features you want to install in controller separated by comma.

- robot-options: robot option you want to pass to the test separated by space.

For Self-Managed project, we need 2 extra parameters:

- trigger-jobs: Self-Managed CSIT will run after succesful project merge, so just fill with '{project}-merge-{stream}'.
- repo-url: Self-Managed project feature repository maven URL (see example below).

So in this case it should look like this:

```
---
- project:
    name: usc-csit-channel
    jobs:
      - inttest-csit-1node

    # The project name
    project: 'usc'

    # The functionality under test
    functionality: 'channel'

    # Project branches
    stream:
      - fluorine:
          branch: 'master'
          trigger-jobs: '{project}-merge-{stream}'
          # yamllint disable-line rule:line-length
          repo-url: 'mvn:org.opendaylight.usc/usc-features/1.6.0-SNAPSHOT/xml/features
↪'

    install:
      - all:
          scope: 'all'

    # Features to install
    install-features: 'odl-restconf,odl-mdsal-apidocs,odl-usc-channel-ui'

    # Robot custom options
    robot-options: ''
```

Save the changes and exit editor.

Optional: Change default tools image

By default a system test spins a tools VM that can be used to run some test tool like mininet, netconf tool, BGP simulator, etc. The default values are listed below and you only need to specify them if you are changing something, for example "tools_system_count: 0" will skip the tools VM if you do not need it. For a list of available images see images-list:

```
---
- project:
    name: openflowplugin-csit-flow-services
    jobs:
      - inttest-csit-1node

    # The project name
```

(continues on next page)

(continued from previous page)

```
    project: 'openflowplugin'

    # The functionality under test
    functionality:
      - flow-services
      - gate-flow-services

    # Project branches
    stream:
      - fluorine:
          branch: 'master'
      - oxygen:
          branch: 'stable/oxygen'
      - nitrogen:
          branch: 'stable/nitrogen'
      - carbon:
          branch: 'stable/carbon'
          karaf-version: 'karaf3'

    install:
      - all:
          scope: 'all'

    # Job images
    tools_system_image: 'ZZCI - Ubuntu 16.04 - mininet-ovs-28 - 20180301-1041'

    # Features to install
    install-features: >
        odl-openflowplugin-flow-services-rest,
        odl-openflowplugin-app-table-miss-enforcer,
        odl-openflowplugin-nxm-extensions

    # Robot custom options
    robot-options: ''
```

Optional: Plot a graph from your job

Scalability and peformance tests not only PASS/FAIL but most important they provide a number or value we want to plot in a graph and track over different builds.

For that you can add the plot configuration like in this example below:

```
---
- project:
    name: openflowplugin-csit-cbench
    jobs:
      - inttest-csit-1node

    # The project name
    project: 'openflowplugin'

    # The functionality under test
    functionality: 'cbench'

    # Project branches
```

(continues on next page)

(continued from previous page)

```
    stream:
      - fluorine:
          branch: 'master'
      - oxygen:
          branch: 'stable/oxygen'
      - nitrogen:
          branch: 'stable/nitrogen'
      - carbon:
          branch: 'stable/carbon'
          karaf-version: 'karaf3'

    install:
      - only:
          scope: 'only'

    # Job images
    tools_system_image: 'ZZCI - Ubuntu 16.04 - mininet-ovs-28 - 20180301-1041'

    # Features to install
    install-features: 'odl-openflowplugin-flow-services-rest,odl-openflowplugin-drop-
↪test'

    # Robot custom options
    robot-options: '-v duration_in_secs:60 -v throughput_threshold:20000 -v latency_
↪threshold:5000'

    # Plot Info
    01-plot-title: 'Throughput Mode'
    01-plot-yaxis: 'flow_mods/sec'
    01-plot-group: 'Cbench Performance'
    01-plot-data-file: 'throughput.csv'
    02-plot-title: 'Latency Mode'
    02-plot-yaxis: 'flow_mods/sec'
    02-plot-group: 'Cbench Performance'
    02-plot-data-file: 'latency.csv'
```

Explanation:

- There are up to 10 plots per job and every plot can track different values, for example max, min, average recorded in a csv file. In the example above you can skip the 02-* lines if you do not use second plot.
- plot-title: title for your plot.
- plot-yaxis: your measurement (xaxis is build # so no need to fill).
- plot-group: just a label, use the same in case you have 2 plots.
- plot-data-file: this is the csv file generated by robot framework and contains the values to plot. Examples can be found in openflow-performance.

Optional: Add Patch Test Job to verify project patches

With the steps above your new csit job will run daily on latest generated distribution. There is one more extra and optional step if you also want to run your system test to verify patches in your project.

The patch test is triggered in gerrit using the keyword:

```
test-$project-$feature
```

The job will:

- Build the gerrit patch.
- Create a distribution containing the patch.
- Trigger some system test (csit) that already exists and you specify with the $feature definition below.

Create $project-patch-test.yaml file in your jjb folder:

```
vim jjb/$project/$project-patch-test-jobs.yaml
```

Fill the information as below:

```
---
- project:
    name: openflowplugin-patch-test
    jobs:
      - inttest-patch-test

    # The project name
    project: 'openflowplugin'

    # Project branches
    stream:
      - fluorine:
          branch: 'master'
          os-branch: 'queens'
      - oxygen:
          branch: 'stable/oxygen'
          os-branch: 'queens'
      - nitrogen:
          branch: 'stable/nitrogen'
          os-branch: 'pike'
      - carbon:
          branch: 'stable/carbon'
          os-branch: 'ocata'
          karaf-version: 'karaf3'

    jdk: 'openjdk8'

    feature:
      - core:
          csit-list: >
              openflowplugin-csit-1node-gate-flow-services-all-{stream},
              openflowplugin-csit-1node-gate-scale-only-{stream},
              openflowplugin-csit-1node-gate-perf-stats-collection-only-{stream},
              openflowplugin-csit-1node-gate-perf-bulkomatic-only-{stream},
              openflowplugin-csit-3node-gate-clustering-only-{stream},
              openflowplugin-csit-3node-gate-clustering-bulkomatic-only-{stream},
              openflowplugin-csit-3node-gate-clustering-perf-bulkomatic-only-{stream}

      - netvirt:
          csit-list: >
              netvirt-csit-1node-openstack-{os-branch}-gate-stateful-{stream}

      - cluster-netvirt:
```

(continues on next page)

(continued from previous page)

```
        csit-list: >
            netvirt-csit-3node-openstack-{os-branch}-gate-stateful-{stream}
```

Explanation:

- name: give some name like $project-patch-test.
- project: set your your project name here (e.g. openflowplugin).
- stream: list the project branches you are going to generate system test. Only last branch if the project is new.
- feature: you can group system tests in features. Note there is a predefined feature -all- that triggers all features together.
- Fill the csit-list with all the system test jobs you want to run to verify a feature.

Debug System Test

Before pushing your system test job into jenkins-releng, it is recommended to debug the job as well as the you system test code in the sandbox. To do that:

- Set up sandbox access using jenkins-sandbox-install instruction.
- Push your new csit job to sandbox:

  Method 1:

  you can write a comment in a releng/builder gerrit patch to have the job automatically created in the sandbox. The format of the comment is:

  ```
  jjb-deploy <job name>
  ```

  Method 2:

  ```
  jenkins-jobs --conf jenkins.ini update jjb/ $project-csit-1node-$functionality-
  ↪only-$branch
  ```

- Open your job in jenkins-sandbox and start a build replacing the PATCHREFSPEC parameter by your int/test patch REFSPEC (e.g. refs/changes/85/23185/1). you can find this info in gerrit top right corner 'Download' button.
- Update the PATCHREFSPEC parameter every time you push a new patchset in the int/test repository.

Optional: Debug VM issues in sandbox

In case of problems with the test VMs, you can easily debug these issues in the sandbox by adding the following lines in a Jenkins shell window:

```
cat > ${WORKSPACE}/debug-script.sh <<EOF

<<put your debug shell script here>>

EOF
scp ${WORKSPACE}/debug-script.sh ${TOOLS_SYSTEM_IP}:/tmp
ssh ${TOOLS_SYSTEM_IP} 'sudo bash /tmp/debug-script.sh'
```

Note this will run a self-made debug script with sudo access in a VM of your choice. In the example above you debug on the tools VM (TOOLS_SYSTEM_IP), use ODL_SYSTEM_IP to debug in controller VM.

Save and push JJB changes

Once you are happy with your system test, save the changes and push them in the releng builder repo:

```
git add -A
git commit -s
git push
```

Important: If this is your first system test job, it is recommended to add the int/test patch (gerrit link) in the commit message so that committers can merge both the int/test and the releng/builder patches at the same time.

Check system test jobs in Jenkins

Once your patches are merged your system test can be browsed in jenkins-releng:

- $project-csit-1node-$functionality-only-$branch -> The single-feature test.
- $project-csit-1node-$functionality-all-$branch -> The multi-project test.
- $yourproject-patch-test-$feature-$branch -> Patch test job.

Note that jobs in jenkins-releng cannot be reconfigured, only jobs in jenkins-sandbox can, that is why it is so important for testers to get access to sandbox.

Support

Integration people are happy to support with questions and recommendations:

- Integration IRC: OpenDaylight channel 'opendaylight-integration
- Integration Mail: OpenDaylight list 'integration-dev@lists.opendaylight.org'

3.3.2 Cluster testing

Contents:

Carbon cluster testing

Contents:

Description of test scenarios

This is a test plan written around M1 of Carbon cycle.

During the cycle several limitations were found, which resulted in tests which implement the scenarios is ways different from what is described here.

For list of limitations and differences, see caveats page. For more detailed descriptions of test cases as implemented, see test description page.

Controller Cluster Service Functional Tests

The purpose of functional tests is to establish a known baseline behavior for basic services exposed to application plugins when the cluster member nodes encounter problems.

Isolation Mechanics Three-node scenarios executed in tests below need to be repeated for three distinct modes of isolation:

1. JVM freeze, initiated by 'kill -STOP <pid>' on the JVM process, followed by a 'kill -CONT <pid>' after three minutes. This simulates a long-running garbage collection cycle, VM suspension or similar, after which the JVM recovers without losing state and scheduled timers going off simultaneously.
2. Network-level isolation via firewalling. Simulates a connectivity issue between member nodes, while all nodes continue to work as usual. This should be done by firewalling all traffic to and from the target node.
3. JVM restart. This simulates a hard error, such as JVM error, VM reboot, and similar. The JVM loses its state and the scenario tests whether the failed node is able to result its operations as a member of the cluster.

Leader Shutdown The Shard implementation allows a leader to be shut down at run time, which is expected to perform a clean hand over to a new leader, elected from the remaining shard members.

DOMDataBroker

Also known as 'the datastore', provides MVCC transaction and data change notifications.

Leader Stability

The goal is to ensure that a single-established shard does not flap, i.e. does not trigger leader movement by causing crashes or timeouts. This is performed by having the BGP load generator run injection of 1 million prefixes, followed by their removal.

This test is executed in three scenarios:

- Single node
- Three-node, with shard leader being local
- Three-node, with shard leader being remote

Success criteria are:

- Both injection and removal succeed
- No transaction errors reported to the generator
- No leader movement on the backend

Clean Leader Shutdown

The goal is to ensure that applications do not observe disruption when a shard leader is shut down cleanly. This is performed by having a steady-stream producer execute operations against the shard and then initiate leader shard shutdown, then the producer is shut down cleanly.

This test is executed in two scenarios:

- Three-node, with shard leader being local
- Three-node, with shard leader being remote

Success criteria are:

- No transaction errors occur
- Producer shuts down cleanly (i.e. all transactions complete successfully)

Test tool: *test-transaction-producer*, running at 1K tps

- Steady, configurable producer started with:
- A transaction chain
- Single transactions (note: these cannot overlap)
- Configurable transaction rate (i.e. transactions-per-second)
- Single-operation transactions
- Random mix across 1M entries

Explicit Leader Movement

The goal is to ensure that applications do not observe disruption when a shard leader is moved as the result of explicit application request. This is performed by having a steady-stream producer execute operations against the shard and then initiate shard leader shutdown, then the producer is shut down cleanly.

This test is executed in three scenarios:

- Three-node, with shard leader being local and becoming remote
- Three-node, with shard leader being remote and remaining remote
- Three-node, with shard leader being remote and becoming local

Success criteria are:

- No transaction errors occur
- Producer shuts down cleanly (i.e. all transactions complete successfully)

Test tool: test-transaction-producer, running at 1K tps Test tool: *test-leader-mover*

- Uses cds-dom-api to request shard movement

Leader Isolation

The goal is to ensure the datastore succeeds in basic isolation/rejoin scenario, simulating either a network partition, or a prolonged GC pause.

This test is executed in the following two scenarios:

- Three-node, partition heals within TRANSACTION_TIMEOUT
- Three-node, partition heals after 2*TRANSACTION_TIMEOUT

Using following steps:

1. Start test-transaction producer, running at 1K tps, non-overlapping, from all nodes to a single shard
2. Isolate leader
3. Wait for followers to initiate election
4. Un-isolate leader

5. Wait for partition to heal
6. Restart failed producer

Success criteria:

- Followers win election in 3
- No transaction failures occur if the partition is healed within TRANSACTION_TIMEOUT
- Producer on old leader works normally after step 6)

Test tool: test-transaction-producer

Client Isolation

The purpose of this test is to ascertain that the failure modes of cds-access-client work as expected. This is performed by having a steady stream of transactions flowing from the frontend and isolating the node hosting the frontend from the rest of the cluster.

This test is executed in one scenario:

- Three node, test-transaction-producer running on a non-leader
- Three node, test-transaction-producer running on the leader

Success criteria:

- After TRANSACTION_TIMEOUT failures occur
- After HARD_TIMEOUT client aborts

Test tool: test-transaction-producer

Listener Isolation

The goal is to ensure listeners do no observe disruption when the leader moves. This is performed by having a steady stream of transactions being observed by the listeners and having the leader move.

This test is executed in two scenarios:

- Three node, test-transaction-listener running on the leader
- Three node, test-transaction-listener running on a non-leader

Using these steps:

- Start the listener on target node
- Start test-transaction-producer on each node, with 1K tps, non-overlapping data
- Trigger shard movement by shutting down shard leader
- Stop producers without erasing data
- Stop listener

Success criteria:

- Listener-internal data tree has to match data stored in the data tree

Test tool: *test-transaction-listener*

- Subscribes a DTCL to multiple subtrees (as specified)

- DTCL applies reported changes to an internal DataTree

DOMRpcBroker

Responsible for routing RPC requests to their implementations and routing responses back to the caller.

RPC Provider Precedence

The aim is to establish that remote RPC implementations have lower priority than local ones, which is to say that any movement of RPCs on remote nodes does not affect routing as long as a local implementation is available.

Test is executed only in a three-node scenario, using the following steps:

1. Register an RPC implementation on each node
2. Invoke RPC on each node
3. Unregister implementation on one node
4. Invoke RPC on that node
5. Re-register implementation on than node
6. Invoke RPC on that node

Success criteria:

- Invocation in steps 2) and 6) results in a response from local node
- Invocation in step 4) results in a response from one of the other two nodes

RPC Provider Partition and Heal

This tests establishes that the RPC service operates correctly when faced with node failures.

Test is executed only in a three-node scenario, using the following steps:

1. Register an RPC implementation on two nodes
2. Invoke RPC on each node
3. Isolate one of the nodes where RPC is registered
4. Invoke RPC on each node
5. Un-isolate the node
6. Invoke RPC on all nodes

Success criteria:

- Step 2) routes the RPC the node nearest node (local or remote)
- Step 4) works, routing the RPC request to the implementation in the same partition
- Step 6) routes the RPC the node nearest node (local or remote)

Action Provider Precedence

The aim is to establish that remote action implementations have lower priority than local ones, which is to say that any movement of actions on remote nodes does not affect routing as long as a local implementation is available.

Test is executed only in a three-node scenario, using the following steps:

1. Register an action implementation on each node
2. Invoke action on each node
3. Unregister implementation on one node
4. Invoke action on that node
5. Re-register implementation on than node
6. Invoke action on that node

Success criteria:

- Invocation in steps 2) and 6) results in a response from local node
- Invocation in step 4) results in a response from one of the other two nodes

Action Provider Partition and Heal

This tests establishes that the RPC service for actions operates correctly when faced with node failures.

Test is executed only in a three-node scenario, using the following steps:

1. Register an action implementation on two nodes
2. Invoke action on each node
3. Isolate one of the nodes where RPC is registered
4. Invoke action on each node
5. Un-isolate the node
6. Invoke action on all nodes

Success criteria:

- Step 2) routes the action request the node nearest node (local or remote)
- Step 4) works, routing the action request to the implementation in the same partition
- Step 6) routes the RPC the node nearest node (local or remote)

DOMNotificationBroker

Provides routing of YANG notifications from publishers to subscribers.

No-loss rate

The purpose of this test is to determine the broker can forward messages without loss. We do this on a single-node setup by incrementally adding publishers and subscribers.

This test is executed in one scenario:

- Single-node

Steps:

- Start test-notification-subscriber
- Start test-notification-publisher at 5K notifications/sec
- Run for 5 minutes, verify no notifications lost
- Add another pair of publisher/subscriber, repeat for rate of 60K notifications/sec

Success criteria:

- No notifications lost at rate of 60K notifications/sec

Test tool: *test-notification-publisher*

- Publishes notifications containing instance id and sequence number
- Configurable rate (i.e. notifications-per-second)

Test tool: *test-notification-subscriber*

- Subscribes to specified notifications from publisher
- Verifies notification sequence numbers
- Records total number of notifications received and number of sequence errors

Cluster Singleton

Cluster Singleton service is designed to ensure that only one instance of an application is registered globally in the cluster.

Master Stability

The goal is to establish the service operates correctly in face of application registration changing without moving the active instance.

The test is performed in a three-node cluster using following steps:

1. Register candidate on each node
2. Wait for master activation
3. Remove non-master candidate,
4. Wait one minute
5. Restore the removed candidate

Success criteria:

- After step 2) there is exactly one master in the cluster
- The master does not move to a different node for the duration of the test

Partition and Heal

The goal is to establish the service operates correctly in face of node failures.

The test is performed in a three-node cluster using following steps:

1. Register candidate on each node
2. Wait for master activation
3. Isolate master node
4. Wait two minutes
5. Un-isolate (former) master node
6. Wait one minute

Success criteria:

- After step 3), master instance is brought down on isolated node
- During step 4) majority partition elects a new master
- Until 5) occurs, old master remains deactivated
- After 6) old master remains deactivated

Chasing the Leader

This test aims to establish the service operates correctly when faced with rapid application transitions without having a stabilized application.

This test is performed in a three-node setup using the following steps:

1. Register a candidate on each node
2. Wait for master activation
3. Newly activated master unregisters itself
4. Repeat 2

Success criteria:

- No failures occur for 5 minutes
- Transition speed is at least 100 movements per second

Controller Cluster Services Longevity Tests

1. Run No-Loss Rate test for 24 hours. No message loss, instability or memory leaks may occur.
2. Repeat Leader Stability test for 24 hours. No transaction failures, instability, leader movement or memory leaks may occur.
3. Repeat Explicit Leader Movement test for 24 hours. No transaction failures, instability, leader movement or memory leaks may occur.
4. Repeat RPC Provider Precedence test for 24 hours. No failures or memory leaks may occur.
5. Repeat RPC partition and Heal test for 24 hours. No failures or memory leaks may occur.
6. Repeat Chasing the Leader test for 24 hours. No memory leaks or failures may occur.

7. Repeat Partition and Heal test for 24 hours. No memory leaks or failures may occur.

NETCONF System Tests

Netconf is an MD-SAL application, which listens to config datastore changes, registers a singleton for every configured device, instantiated singleton is updating device connection data in operational datastore, maintaining a mount point and handling access to the mounted device.

Basic configuration and mount point access

No disruptions, ordinary netconf operation with restconf calls to different cluster members.

Test is executed in a three-node scenario, using the following steps:

1. Configure connection to test device on member-1.
2. Create, update and delete data on the device using calls to member-2.
3. Each state change confirmed by reading device data on member-3.
4. De-configure the device connection.

Success criteria:

- All reads confirm data operations are applied correctly.

Device owner killed

Killing current device owner leads to electing new owner. Operations are still applied.

The test is performed in a three-node cluster using following steps:

1. Configure connection to test device on member-1.
2. Create data on the device using a call to member-2.
3. Locate and kill the device owner member.
4. Wait for a new owner to get elected.
5. Update data on the device using a call to one of the surviving members.
6. Restart the killed member.
7. Update the data again using a call to the restarted member.

Success criteria:

- Each operation (including restart) is confirmed by reads on all members currently up.

Rolling restarts

Each member is restarted (start is waiting for cluster sync) in succession, this is to guarantee each Leader is affected.

The test is performed in a three-node cluster using following steps:

1. Configure connection to test device on member-1.
2. Kill member-1.

3. Create data on the device using a call to member-2.
4. Start member-1.
5. Kill member-2.
6. Update data on the device using a call to member-3.
7. Start member-2.
8. Kill member-3.
9. Delete data on the device using a call to member-1.
10. Start member-3.

Success criteria:

- After every operation, reads on both living members confirm it was applied.
- After every start, a read on the started node confirms it sees the device data from the previous operation.

Caveats

This sub-page describes ways the test implementation (or results) differs from the original specification and which information motivates the difference.

Jenkins job structure

- Information

At the start of test implementation, all the Controller 3node test cases were added into an existing Jenkins job.

During test development it was become clear, that adding all possible tests would make the job to run too long.

Dividing the job into several smaller ones is possible, but most likely the history would be lost, unless Linux Foundation admins figure out a way to create multiple job clones with history copied.

- Testing consequence

Even with number of test cases reduced (see below), the job duration is around three and half hours.

- How to fix

After Carbon SR2 release, the jobs can be split, as there will be enough time to generate new history till Carbon SR3.

Akka bugs

These are bugs which need either a fix in Akka codebase, or a workaround which would be too time-consuming to implement in ODL.

Both bugs manifest as UnreachableMember event (without intentional isolation).

Slow heartbeats

- Information

Akka sends periodic heartbeats in order to detect when the other member is being unresponsive.

The heartbeats are being serialized into the same TCP channel as ordinary data, which means if ODL is processing big amount of data, the heartbeats can spend a long time in TCP (or other) buffers before being processed. When this time exceeds a specific value (currently 6 seconds), the peer memeber is declared unreachable, generally leading to leader movement.

This affects BGP test results on 3node setup, as ODL is processing BGP data as quickly as possible, but the current BGP implementation does not handle rib owner movement gracefully (and leader movement is explicitly checked by the test, as the scenario dictates it should not happen). This does not affect other data broker tests, 1000 transactions per second do not generate critical throughput.

- Testing consequence

Three test cases are failing due to Bug 8318.

- How to fix

Possibly, a different akka configuration could be applied to separate akka cluster status messages into a different TCP stream than ordinary data stream.

Otherwise, a contribution to Akka project would be needed.

Reachability gossip

- Information

Akka uses a gossip protocol to advertize one member's reachability to other members. There is a logic which allows for faster detection of unreachable members, when a member can declare its peer unreachable if it got information from another peer which is considered more up-to-date.

Ocassionally, this logic results in undesired behavior. This is when the supposedly up-to-date peer has been isolated and now it is rejoining. Depending on timing, this can introduce additional leader movement, or a very brief moment when a member "forgets" RPC registrations from other member.

This is causing bugs 8420 and 8430.

- Testing consequence

This affects "partition and heal" scenarios in singleton testing. In functional tests, the failures are infrequent enough to consider the test mostly stable overall, but the corresponding longevity jobs are failing consistently.

The tests for "partition and heal" scenarios in RPC testing have been changed to tolerate wrong RPC results for 10 seconds to work around this Akka bug.

- How to fix

This does not seem fixable on ODL level, contribution to Akka project is needed.

Missing features

Cluster yang notifications

- Information

Yang notifications are not delivered to peer members. Bug 2139 is only fixed for data change notifications, not Yang notifications.

Bug 2140 tracks adding this missing functionality.

- Testing consequence

Notification suites are running on 1-node setup only.

- How to fix

After the funtionality is added, it will be straightforward to add 3node tests.

New features

Tell-based protocol

- Information

Tell-based protocol is an alternative to ask-based protocol from Boron. Which protocol to use is decided by a line in a configuration file (org.opendaylight.controller.cluster.datastore.cfg).

Some scenarios are expected to fail due to known limitations of ask-based protocol. More specifically, if a shard leader moves while a transaction is open in ask-based protocol, the transaction will fail (AskTimeoutException).

This affects only data broker tests, not RPC calls.

- Testing consequence

In principle, this doubles the number of configurations to be tested, but see below.

- How to fix

It is planned for tell-based protocol to become the default setting after Carbon SR2. After that, tests for ask-based protocol can be converted or removed.

Prefix-based shards

- Information

Tell-based shards are an alternative to module-based shards from Boron. Tell-based shards can be only created dynamically (as opposed to being read from a configuration file at startup). It is possible to use both types of shards, but data writes and reads use different API, so any Mdsal application needs to know which API to use.

The implementation of prefix-based shards is hardwired to tell-based protocol (even if ask-based protocol is configured as the default).

- Testing consequence

This doubles the number of configurations to be tested, for tests related to data droker (RPCs are unaffected).

- How to fix

ODL contains great many applications which use APIs for module-based shards. It is expected that multiple releases would still need both types of tests cases. Module-based shards will be deprecated and removed eventually.

Producer options

- Information

Data producers for module-based shards can produce either chained transactions or standalone transactions. Data producers for prefix-based shards can produce either non-isolated transactions (change notifications can combine several transactions together) or isolated transactions.

- Testing consequence

In principle, this results in multiple Robot test cases for the same documented scenario case, but see below.

- How to fix

All test cases will be needed in forseeable future. Instead, more negative test cases may need be added to verify different options lead to different behavior.

Initial leader placement

- Information

Some scenarios do not specify initial locations of relevant shard leaders. Test results can depend on it in presence of bugs.

This is mostly relevant to BGP test, which has three relevant members: Rib owner, default operation shard leader and topology operational shard leader.

- Testing consequence

Two test cases are tested. The two shard leaders are always together, rib owner is either co-located or not. This is done by suite moving shard leaders after detecting rib owner location.

- How to fix

Even more placements can be tested when job duration stops being the limiting factor.

Reduced BGP scaling

- Information

Rib owner maintains de-duplicated data structures. Other members get serialized copies and they do not de-duplicate.

Even single node strugless to fit into 6GB heap with tell-based protocol, see Bug 8649.

- Testing consequence

Scale from reported tests reduced from 1 million prefixes to 300 thousand prefixes.

- How to fix

Other members should be able to perform de-duplication, but developing that takes effort.

In the meantime, Linux Foundation could be convinced to allow for bigger VMs, currently limited by infrastructure available.

Increased timeouts

RequestTimeoutException

- Information

With tell-based protocol, restconf requests might stay open up to 120 seconds before returning an error. Even shard state reads using Jolokia can take long time if the shard actor is busy processing other messages.

- Testing consequence

This increases duration for tests which need to verify transaction errors do happen after sufficiently long isolation. Also, duration is increased if a test fails on a read which is otherwise quick.

- How to fix

This involves a trade-off between stability and responsiveness. As MD-SAL applications rarely tolerate transaction failures, users would prefer stability. That means relatively longer timeouts are there to stay, which means test case duration will stay high in negative (or failing positive) tests.

Client abort timeout

- Information

Client abort timeout is currently set to 15 minutes. The operational consequence is just an inability to start another data producer on a member isolated for that long. This test has too long duration compared to its usefulness.

- Testing consequence

This test case has never been implemented.

Instead a test with isolation shorter than 120 seconds is implemented, the test verifies the data producer continues its operation without RequestTimeoutException.

- How to fix

It is straighforward to add the missing test cases when job duration stops being a limiting factor.

No shard shutdown

- Common information.

There are multiple RPCs offering different "severity" of shard shutdown. For technical details see comments on change 58580.

If tests perform rigorous teardown, the shard replica should be re-activated, which is an operation not every RPC supports.

Listener stability suite

- Information

Current implementation of data listeners relies on a shard replica to be active on a member which is to receive the notification. Until that is imroved, Bug 8629 prevents this scenario from being tested as described.

- Testing consequence

The suite uses become-leader RPC instead. This has an added benefit of test case being able to pick which member is to become the new leader (adding one more test case when the old leader was not co-located with the listener).

Also, no teardown step is needed, the final cluster state is not missing any shard replica.

- How to fix

The original test can be implemented when listener implementation changes. But the test which uses become-leader might be better overall.

Clean leader shutdown suite

- Information

Some implementations of shutdown RPCs have a side effect of also shutting down shard state notifier. For details see Bug 8794.

The remove-shard-replica RPC does not have this downside, but it changes shard configuration, which was not intended by the original scenario definition.

- Testing consequence

Test cases for this scenario were switched to use remove-shard-replica.

- How to fix

There is an open debate on whether "shard shutdown" RPC with less operations (compared to remove-shard-replica) is something user wants and should be given access to.

If yes, tests can be switched to such an RPC, assuming the shard notifier issue is also fixed.

Hard reboots between test cases

- Information

Timing errors in Robot code lead to Robot being unable to restore original state without restarts.

During development, we started without any hard reboots, and that was finding bugs in teardown steps of scenarios. But test independence was more important at that time, so current tests are less sensitive to teardown failures.

- Testing consequence

Around 115 second per ODL reboot, this time is added to every test case running time. Together with increased timeouts, this motivates leaving out some test cases to allow faster change verification.

- How to fix

Ideally, we would want both jobs with hard resets and jobs without them. The jobs without resets can be added gradually after splitting the current single job.

Isolation mechanics

- Information

During development, it was found that freeze and kill mechanics affect the co-located java test driver without exposing any new bugs.

Turns out AAA functionality attempts to read from datastore, so isolated member returns http status code 401.

- Testing consequence

Only iptables filtering is used in order to reduce test job duration.

Isolated members are never queried directly. A leader member is considered isolated when other members elect a lew leader. A member is considered rejoined when it responds reporting itself as a follower.

- How to fix

It is straightforward to add test cases for kill and freeze where appropriate, but once again this can be done gradually when job duration is not a limiting factor.

Reduced number of combinations

- Information

Prefix-based shards always use tell-based protocol, so suites which test them with ask-based protocol configuration can be skipped.

Ask-based protocol is known to fail on AskTimeoutException on leader movement, so suites which produce transactions constantly can be skipped.

Most test cases are not sensitive to data producer options.

- Testing consequence

BGP tests and singleton tests use module-based shards only, both protocols. Other suites related to data broker are testing only tell-based protocol, both shard types. Netconf tests and RPC tests use module-based shards with ask-based protocol only. Only client isolaton suite tests different producer options.

- How to fix

More ests can be added gradually (see above).

Possibly, not every combination is worth the duration it takes, but that could be alleviated if Linux Foundation infrastructure grows in size significantly.

Reduced performance

- Information

In order to reduce test job duration, suites wait for minimal functionality (jolokia reporting shards are in sync) after restarting ODL. That means unrelated karaf features might still being installed whet test is in progress. This should not affect functional tests, but it can reduce performance observed.

The only suite observing strong enough performance inpact is chasing the leader.

- Testing consequence

Functional tests for chasing the leader suite tolerate frequencies higher than 50 un-registrations per second. Longevity suite still requires full 100 unregistrations per second.

- How to fix

Suite can wait for better symptom of ODL being ready, for example by requiring CPU usage to become less that a chosen threshold.

Missing logs

- Information

Robot VM has only 2GB of RAM and longevity jobs tend to produce large output.xml files.

Ocasionally, a job can create karaf.log files so large they fail to download, in extreme cases filling ODL VM disk and causing failures.

This affects mostly longevity jobs (and runs with verbose logging) if they pass.

- Testing consequence

Robot data stored is reduced to avoid this issue, sometimes leading to less details available. This issue is still not fully resolved, so ocassionally Robot log or karaf log is still missing if the job in question fails in an unexpected way.

- How to fix

It is possible for Robot test to put additional data into separate files. Unnecessarily verbose logs could be fixed where needed.

As this limitation only hurts in newly occuring bugs, it is not really possible to entirely avoid this.

Weekend outages

- Information

Linux foundation ifrastructure teem occasionally needs to perform changes which affect running jobs. To reduce this impact, such changes are usually done over weekend.

Cluster testing currently contains seve longevity jobs which block resources for 23 hours. As that is a significant portion of available resources, the longevity jobs are only run on weekend where the impact on frequency of other job is less critical.

- Testing consequence

Sometimes, the longevity jobs are affected by infrastructure team activities, leading to lost results or spurious failures. One such symptom is tracked as Bug 8959.

- How to fix

It might be possible to spread longevity jobs over work days. As distributing jobs manually is not a scalable option, a considerable work would be needed to create an automatic way.

Infrastructure changes are not very frequent, and having jobs run at the same predictable time is convenient from reporting point of view, so perhaps it is okay to keep the current setup.

List of test cases

Each test case has a shorter code, tables with results use that code. In result tables, the code is a link to this document, due to coala ReST requirements, the codes are (self-pointing) links also in this document.

Other links point to scenario definitions ao caveat items.

- DOMDataBroker: Producers make 1000 transactions per second, except BGP which works full speed.
  - Leader stability: BGP inject benchmark (thus module shards only), 300k prefixes, 1 Python peer. Progress tracked by counting prefixes in example-ipv4-topology.
    - Ask-based protocol:
      - Single member: *bgp-1n-300k-a*
    - Tell-based protocol:
      - Single member: *bgp-1n-300k-t*
      - Three members:
        - Leaders local: *bgp-3n-300k-ll-t*
        - Leaders remote: *bgp-3n-300k-lr-t*
        - Longevity: *bgp-3n-300k-t-long*
  - Clean leader shutdown, Tell-based protocol:
    - Module-based shards:
      - Shard leader local to producer: *ddb-cls-ms-ll-t*

        - Shard leader remote to producer: *ddb-cls-ms-lr-t*
    - Prefix-based shards:
        - Shard leader local to producer: *ddb-cls-ps-ll-t*
        - Shard leader remote to producer: *ddb-cls-ps-lr-t*
- Explicit leader movement, Tell-based protocol:
    - Module-based shards:
        - Local leader to remote: *ddb-elm-ms-lr-t*
        - Remote leader to other remote: *ddb-elm-ms-rr-t*
        - Remote leader to local: *ddb-elm-ms-rl-t*
        - Longevity (randomized direction): *ddb-elm-mc-t-long*
    - Prefix-based shards:
        - Local leader to remote: *ddb-elm-ps-lr-t*
        - Remote leader to other remote: *ddb-elm-ps-rr-t*
        - Remote leader to local: *ddb-elm-ps-rl-t*
- Leader isolation (network partition only), Tell-based protocol:
    - Module-based shards:
        - Heal within transaction timeout: *ddb-li-ms-st-t*
        - Heal after transaction timeout: *ddb-li-ms-dt-t*
    - Prefix-based shards:
        - Heal within transaction timeout: *ddb-li-ps-st-t*
        - Heal after transaction timeout: *ddb-li-ps-dt-t*
- Client isolation, Tell-based protocol:
    - Module-based shards:
        - Leader local:
            - Simple transactions: *ddb-ci-ms-ll-st-t*
            - Transaction chain: *ddb-ci-ms-ll-ct-t*
        - Leader remote:
            - Simple transactions: *ddb-ci-ms-lr-st-t*
            - Transaction chain: *ddb-ci-ms-lr-ct-t*
    - Prefix-based shards:
        - Leader local:
            - Isolated transactions: *ddb-ci-ps-ll-it-t*
            - Non-isolated transactions: *ddb-ci-ps-ll-nt-t*
        - Leader remote:
            - Isolated transactions: *ddb-ci-ps-lr-it-t*
            - Non-isolated transactions: *ddb-ci-ps-lr-nt-t*

    - Listener stablity, Tell-based protocol:
        - Module-based shards:
            - Local to remote: *ddb-ls-ms-lr-t*
            - Remote to remote: *ddb-ls-ms-rr-t*
            - Remote to local: *ddb-ls-ms-rl-t*
        - Prefix-based shards:
            - Local to remote: *ddb-ls-ps-lr-t*
            - Remote to remote: *ddb-ls-ps-rr-t*
            - Remote to local: *ddb-ls-ps-rl-t*
- DOMRpcBroker, ask-based protocol:
    - RPC Provider Precedence:
        - Functional: *drb-rpp-ms-a*
        - Longevity: *drb-rpp-ms-a-long*
    - RPC Provider Partition and Heal:
        - Functional: *drb-rph-ms-a*
        - Longevity: *drb-rph-ms-a-long*
    - Action Provider Precedence: *drb-app-ms-a*
    - Action Provider Partition and Heal: *drb-aph-ms-a*
- DOMNotificationBroker: Only for 1 member, ask-based protocol.
    - No-loss rate: Publisher-subscriber pairs, 5k nps per pair.
        - Functional (5 minute tests for 1, 4 and 12 pairs): *dnb-1n-60k-a*
        - Longevity (12 pairs): *dnb-1n-60k-a-long*
- Cluster Singleton:
    - Ask-based protocol:
        - Master Stability: *ss-ms-ms-a*
        - Partition and Heal:
            - Functional: *ss-ph-ms-a*
            - Longevity: *ss-ph-ms-a-long*
        - Chasing the Leader:
            - Functional: *ss-cl-ms-a*
            - Longevity: *ss-cl-ms-a-long*
    - Tell-based protocol:
        - Master Stability: *ss-ms-ms-t*
        - Partition and Heal: *ss-ph-ms-t*
        - Chasing the Leader: *ss-cl-ms-t*
- Netconf system tests (ask-based protocol, module-based shards):

- Basic access: *netconf-ba-ms-a*
- Owner killed: *netconf-ok-ms-a*
- Rolling restarts: *netconf-rr-ms-a*

Permanent draft, inaccessible: Sandbox test report

Test Case Summary

RelEng stability summary.

- tba: Recent failures to be analyzed yet: 0.
- test: Recent failures caused by wrong assumptions in test: 0.
- akka: Recent failures related to pure UnreachableMember: 4.
- tell: Recent failures not clearly caused by UnreachableMember: 6.
- few: Tests passing unless low frequency failure happens: 2 (1 without duplication). (Low frequency means UnreachableMemeber or similar, related to Akka where Controller code has not real control.)
- pass: Tests passing consistently: 41 (39 without duplication).
- Total: 53 (50 without duplication).
- Total minus akka: 49 (46 without duplication).
- Total minus akka passing always or mostly: 43 (40 without duplication).
- Acceptance rate: 43/49=87.75% (40/46=86.95% without duplication).

Table

S017 instead of 2017 means Sandbox run (includes changes not merged to stable/carbon yet).

Last fail is date of last failure not caused by infra (or by a typo in test or by netconf/bgp failing to initialize properly).

"S 17" or "2 17" in Last run means the documented run was superseded by a newer one, but not analyzed yet.

"no sr3" means this test was not run on Sandbox, SR2 result is reported instead. "few" status from SR2 is not inherited (such tests are marked as "pass"). "long ago" means the last real test failue happened somewhere before SR2 release (or never).

TODO: Copy formatting from sr2 page.

Table 2: Releng stability results (pre-SR2)

Scenario name	Type	Last fail	Last run	Bugs	Robot link
bgp-1n-1m-a	pass	no sr3	no sr3		no sr3
bgp-1n-300k-t	pass	no sr3	no sr3		no sr3
bgp-3n-300k-ll-t	akka	no sr3	no sr3	8318	no sr3
bgp-3n-300k-lr-t	akka	no sr3	no sr3	8318	no sr3
ddb-cls-ms-ll-t	pass	long ago	S017-08-24		no fail this week

Continued on next page

Table 2 – continued from previous page

Scenario name	Type	Last fail	Last run	Bugs	Robot link
ddb-cls-ms-lr-t	pass	long ago	S017-08-24		no fail this week
ddb-cls-ps-ll-t	pass	long ago	S017-08-24		no fail this week
ddb-cls-ps-lr-t	pass	long ago	S017-08-24		no fail this week
ddb-elm-ms-lr-t	pass	long ago	S017-08-24		no fail this week
ddb-elm-ms-rr-t	pass	long ago	S017-08-24		no fail this week
ddb-elm-ms-rl-t	pass	long ago	S017-08-24		no fail this week
ddb-elm-ps-lr-t	pass	long ago	S017-08-24		no fail this week
ddb-elm-ps-rr-t	pass	long ago	S017-08-24		no fail this week
ddb-elm-ps-rl-t	pass	long ago	S017-08-24		no fail this week
ddb-li-ms-st-t	pass	long ago	S017-08-24		no fail this week
ddb-li-ms-dt-t	pass	long ago	S017-08-24		no fail this week
ddb-li-ps-st-t	pass	long ago	S017-08-24		no fail this week
ddb-li-ps-dt-t	tell	S017-08-24	S017-08-24	8845	link
ddb-ci-ms-ll-ct-t	pass	long ago	S017-08-24		no fail this week
ddb-ci-ms-ll-st-t	pass	long ago	S017-08-24		no fail this week
ddb-ci-ms-lr-ct-t	pass	long ago	S017-08-24		no fail this week
ddb-ci-ms-lr-st-t	pass	long ago	S017-08-24		no fail this week

Continued on next page

Table 2 – continued from previous page

Scenario name	Type	Last fail	Last run	Bugs	Robot link
ddb-ci-ps-ll-ct-t	pass	long ago	S017-08-24		no fail this week
ddb-ci-ps-ll-st-t	pass	long ago	S017-08-24		no fail this week
ddb-ci-ps-lr-ct-t	pass	long ago	S017-08-24		no fail this week
ddb-ci-ps-lr-st-t	pass	long ago	S017-08-24		no fail this week
ddb-ls-ms-lr-t	pass	long ago	S017-08-24		no fail this week
ddb-ls-ms-rr-t	pass	long ago	S017-08-24		no fail this week
ddb-ls-ms-rl-t	pass	long ago	S017-08-24		no fail this week
ddb-ls-ps-lr-t	tell	S017-08-24	S017-08-24	8733	link
ddb-ls-ps-rr-t	tell	S017-08-24	S017-08-24	8733	link
ddb-ls-ps-rl-t	pass	long ago	S017-08-24		no fail this week
drb-rpp-ms-a	pass	long ago	S017-08-24		no fail this week
drb-rph-ms-a	pass	long ago	S017-08-24		no fail this week
drb-app-ms-a	pass	long ago	S017-08-24		no fail this week
drb-aph-ms-a	pass	long ago	S017-08-24		no fail this week
dnb-1n-60k-a	pass	no sr3	no sr3		no sr3
ss-ms-ms-a	pass	long ago	S017-08-24		no fail this week
ss-ph-ms-a	few	S017-08-24	S017-08-24	8420	link
ss-cl-ms-a	pass	long ago	S017-08-24		no fail this week
ss-ms-ms-t	pass	long ago	S017-08-24		no fail this week

Continued on next page

Table 2 – continued from previous page

Scenario name	Type	Last fail	Last run	Bugs	Robot link
ss-ph-ms-t	few	S017-08-24	S017-08-24	8420	link
ss-cl-ms-t	pass	long ago	S017-08-24		no fail this week
netconf-ba-ms-a	pass	no sr3	no sr3		no fail this week
netconf-ok-ms-a	tell	no sr3	no sr3	9027	no fail this week
netconf-rr-ms-a	tell	no sr3	no sr3	9027	no fail this week
bgp-3n-300k-t-long	akka	no sr3	no sr3	8318	no sr3
ddb-elm-mc-t-long	pass	no sr3	no sr3		no sr3
drb-rpp-ms-a-long	pass	no sr3	no sr3		no sr3
drb-rph-ms-a-long	pass	no sr3	no sr3	8430	no sr3
dnb-1n-60k-a-long	pass	no sr3	no sr3		no sr3
ss-ph-ms-a-long	akka	no sr3	no sr3	8420	no sr3
ss-cl-ms-a-long	tell	S017-08-23	S017-08-23	9054	link

For descriptions of test cases, see description page. Note that the link contains current description, the details might have been implemented differently at SR1 release.

Draft, outdated: Carbon release test report

Table

Table 3: Test results (pre-release)

Scenario name	Run date	Bug numbers	Result
bgp-1n-1m-a	2017-05-23		PASS
bgp-1n-1m-t	2017-05-23		PASS
bgp-3n-300k-ll-t	2017-05-23	8318	FAIL
bgp-3n-300k-lr-t	2017-05-23	8318	FAIL
ddb-cls-ms-ll-t	2017-05-23	8403	FAIL
ddb-cls-ms-lr-t	2017-05-23		PASS
ddb-cls-ps-ll-t	2017-05-23	8403	FAIL
ddb-cls-ps-lr-t	2017-05-23		PASS
ddb-elm-ms-lr-t	2017-05-23	8403	FAIL
ddb-elm-ms-rr-t	2017-05-23		PASS
ddb-elm-ms-rl-t	2017-05-23	8403	FAIL
ddb-elm-ps-lr-t	2017-05-23		PASS
ddb-elm-ps-rr-t	2017-05-23		PASS
ddb-elm-ps-rl-t	2017-05-23	8403	FAIL
ddb-li-ms-st-t	2017-05-23	8445	FAIL
ddb-li-ms-dt-t	2017-05-23	8494	FAIL
ddb-li-ps-st-t	2017-05-23	8371	FAIL

Continued on next page

Table 3 – continued from previous page

Scenario name	Run date	Bug numbers	Result
ddb-li-ps-dt-t	2017-05-23	8371	FAIL
ddb-ci-ms-ll-ct-t	2017-05-23	8494	FAIL
ddb-ci-ms-ll-st-t	2017-05-23	8494	FAIL
ddb-ci-ms-lr-ct-t	2017-05-23		PASS
ddb-ci-ms-lr-st-t	2017-05-23		PASS
ddb-ci-ps-ll-ct-t	2017-05-23	8494	FAIL
ddb-ci-ps-ll-st-t	2017-05-23	8494	FAIL
ddb-ci-ps-lr-ct-t	2017-05-23		PASS
ddb-ci-ps-lr-st-t	2017-05-23		PASS
ddb-ls-ms-ll-t	2017-05-23	8524	FAIL
ddb-ls-ms-lr-t	2017-05-23	8534	FAIL
ddb-ls-ps-ll-t	2017-05-23	8524	FAIL
ddb-ls-ps-lr-t	2017-05-23	8524	FAIL
drb-rpp-ms-a	2017-05-23		PASS
drb-rph-ms-a	2017-05-23		PASS
drb-app-ms-a	2017-05-23		PASS
drb-aph-ms-a	2017-05-23		PASS
dnb-1n-60k-a	2017-05-23		PASS
ss-ms-ms-a	2017-05-23		PASS
ss-ph-ms-a	2017-05-23		PASS
ss-cl-ms-a	2017-05-23		PASS
ss-ms-ms-t	2017-05-23		PASS
ss-ph-ms-t	2017-05-23		PASS
ss-cl-ms-t	2017-05-23		PASS
netconf-ba-ms-a	2017-05-23		PASS
netconf-ok-ms-a	2017-05-23		PASS
netconf-rr-ms-a	2017-05-23		PASS
bgp-3n-300k-t-long	2017-05-14	8443	FAIL
ddb-elm-mc-a-long	2017-05-14	8434	FAIL
drb-rpp-ms-a-long	2017-05-14		PASS
drb-rph-ms-a-long	2017-05-14		PASS
dnb-1n-60k-a-long	2017-05-14		PASS
ss-ph-ms-a-long	2017-05-14	8420	FAIL
ss-cl-ms-a-long	2017-05-14		PASS

For descriptions of test cases, see description page. Note that the link contains current description, the details might have been implemented differently at SR1 release.

Draft, outdated: Carbon SR1 test report

Test Case Summary

RelEng stability summary.

- tba: Recent failures to be analyzed yet: 0.
- test: Recent failures caused by wrong assumptions in test: 0.
- akka: Recent failures related to pure UnreachableMember: 5.
- tell: Recent failures not clearly caused by UnreachableMember: 9.

- few: Tests passing unless low frequency failure happens: 22 (21 without duplication). (Low frequency means UnreachableMemeber or "Message was not delivered, dead letters encountered", both are related to Akka where Controller code has not real control.)
- pass: Tests passing consistently: 17 (15 without duplication).
- Total: 53 (50 without duplication).
- Total minus akka: 48 (45 without duplication).
- Total minus akka passing always or mostly: 39 (36 without duplication).
- Acceptance rate: 39/48=81.25% (36/45=80.00% without duplication).

Table

S017 instead of 2017 means Sandbox run (includes changes not merged to stable/carbon yet).

Last fail is date of last failure not caused by infra (or by a typo in test or by netconf/bgp failing to initialize properly).

"S 17" or "2 17" in Last run means the documented run was superseded by a newer one, but not analyzed yet.

"long ago" means the last real test failue happened before around 2017-05-19, or never.

Table 4: Releng stability results (pre-SR1)

Scenario name	Type	Last fail	Last run	Bugs	Robot link
bgp-1n-1m-a	pass	long ago	2017-07-14		link
bgp-1n-300k-t	pass	long ago	2017-07-14		link
bgp-3n-300k-ll-t	akka	2017-07-14	2017-07-14	8318	link
bgp-3n-300k-lr-t	akka	2017-07-13	2017-07-14	8318	link
ddb-cls-ms-ll-t	few	2017-07-04	2017-07-15	8794	link
ddb-cls-ms-lr-t	few	2017-07-08	2017-07-15	8618	link
ddb-cls-ps-ll-t	few	2017-07-09	2017-07-15	8794	link
ddb-cls-ps-lr-t	pass	long ago	2017-07-15		link
ddb-elm-ms-lr-t	few	2017-06-13	2017-07-15	8618	link
ddb-elm-ms-rr-t	few	2017-06-10	2017-07-15	8618	link
ddb-elm-ms-rl-t	few	2017-06-27	2017-07-15	8749	link
ddb-elm-ps-lr-t	few	2017-06-11	2017-07-15	8664	link
ddb-elm-ps-rr-t	pass	long ago	2017-07-15		link
ddb-elm-ps-rl-t	few	2017-06-07	2017-07-15	8403	link
ddb-li-ms-st-t	tell	2017-07-15	2017-07-15	8792	link
ddb-li-ms-dt-t	tell	2017-07-15	2017-07-15	8619	link
ddb-li-ps-st-t	few	2017-06-08	2017-07-15	8371	link
ddb-li-ps-dt-t	tell	2017-07-15	2017-07-15	8845	link
ddb-ci-ms-ll-ct-t	few	2017-06-07	2017-07-15	8494	link
ddb-ci-ms-ll-st-t	tell	2017-07-15	2017-07-15	8494	link
ddb-ci-ms-lr-ct-t	few	2017-06-08	2017-07-15	8636	link
ddb-ci-ms-lr-st-t	tell	2017-07-15	2017-07-15	8494	link
ddb-ci-ps-ll-ct-t	few	2017-06-28	2017-07-15	8494	link
ddb-ci-ps-ll-st-t	few	2017-06-28	2017-07-15	8494	link
ddb-ci-ps-lr-ct-t	few	2017-06-28	2017-07-15	8494	link
ddb-ci-ps-lr-st-t	few	2017-06-28	2017-07-15	8494	link
ddb-ls-ms-lr-t	tell	2017-07-15	2017-07-15	8792	link

Continued on next page

Table 4 – continued from previous page

Scenario name	Type	Last fail	Last run	Bugs	Robot link
ddb-ls-ms-rr-t	tell	2017-07-14	2017-07-15	8792	link
ddb-ls-ms-rl-t	tell	2017-07-12	2017-07-15	8792	link
ddb-ls-ps-lr-t	pass	long ago	2017-07-15		link
ddb-ls-ps-rr-t	few	2017-06-26	2017-07-15	8733	link
ddb-ls-ps-rl-t	pass	long ago	2017-07-15		link
drb-rpp-ms-a	pass	long ago	2017-07-15		link
drb-rph-ms-a	few	2017-06-28	2017-07-15	8430	link
drb-app-ms-a	pass	long ago	2017-07-15		link
drb-aph-ms-a	few	2017-07-02	2017-07-15	8430	link
dnb-1n-60k-a	pass	long ago	2017-07-15		link
ss-ms-ms-a	pass	long ago	2017-07-15		link
ss-ph-ms-a	few	2017-06-29	2017-07-15	8420	link
ss-cl-ms-a	pass	long ago	2017-07-15		link
ss-ms-ms-t	pass	long ago	2017-07-15		link
ss-ph-ms-t	few	2017-07-15	2017-07-15	8420	link
ss-cl-ms-t	pass	long ago	2017-07-15		link
netconf-ba-ms-a	pass	long ago	2017-07-14		link
netconf-ok-ms-a	few	2017-06-18	2017-07-14	8596	link
netconf-rr-ms-a	pass	long ago	2017-07-14		link
bgp-3n-300k-t-long	akka	2017-07-08	2017-07-08	8318	link
ddb-elm-mc-t-long	tell	2017-07-08	2017-07-08	8618	link
drb-rpp-ms-a-long	few	2017-05-07	2017-07-08	8430	link
drb-rph-ms-a-long	akka	2017-07-08	2017-07-08	8430	link
dnb-1n-60k-a-long	pass	long ago	2017-07-08		link
ss-ph-ms-a-long	akka	2017-07-08	2017-07-08	8420	link
ss-cl-ms-a-long	pass	long ago	2017-07-08		link

For descriptions of test cases, see description page. Note that the link contains current description, the details might have been implemented differently at SR1 release.

Carbon SR2 test report

Test Case Summary

RelEng stability summary.

- tba: Recent failures to be analyzed yet: 0.
- test: Recent failures caused by wrong assumptions in test: 0.
- akka: Recent failures related to pure UnreachableMember: 4.
- tell: Recent failures not clearly caused by UnreachableMember: 4.
- few: Tests passing unless low frequency failure happens: 7 (6 without duplication). (Low frequency means infra issues or UnreachableMemeber, related to Akka where Controller code has not real control.)
- pass: Tests passing consistently: 38 (36 without duplication).
- Total: 53 (50 without duplication).
- Total minus akka: 49 (46 without duplication).

- Total minus akka, passing always or mostly: 45 (42 without duplication).
- Acceptance rate: 45/49=91.83% (42/46=91.30% without duplication).

Table

S017 instead of 2017 means Sandbox run (includes changes not merged to stable/carbon yet).

Last fail is date of last failure not caused by infra (or by a typo in test or by netconf/bgp failing to initialize properly).

"S 17" or "2 17" in Last run means the documented run was superseded by a newer one, but not analyzed yet.

"few" status from SR1 is not inherited (such tests are marked as "pass"). "long ago" means the last real test failue happened somewhere around SR1 release (or before that, or never).

If status is a link, it points to the latest relevant robot failure, or a history to see the stability. In case of failure, Bugs field gives the reason of that failure.

Table 5: Releng stability results (post SR1, pre SR2)

Test case	Last fail	Last run	Bugs	Status
bgp-1n-300k-a	long ago	2017-09-18		PASS
bgp-1n-300k-t	long ago	2017-09-18		PASS
bgp-3n-300k-ll-t	2017-09-16	2017-09-18	8318	AKKA
bgp-3n-300k-lr-t	2017-09-16	2017-09-18	8318	AKKA
ddb-cls-ms-ll-t	2017-08-24	2017-09-18		PASS
ddb-cls-ms-lr-t	long ago	2017-09-18		PASS
ddb-cls-ps-ll-t	long ago	2017-09-18		PASS
ddb-cls-ps-lr-t	long ago	2017-09-18		PASS
ddb-elm-ms-lr-t	long ago	2017-09-18		PASS
ddb-elm-ms-rr-t	long ago	2017-09-18		PASS
ddb-elm-ms-rl-t	long ago	2017-09-18		PASS
ddb-elm-ps-lr-t	long ago	2017-09-18		PASS
ddb-elm-ps-rr-t	long ago	2017-09-18		PASS
ddb-elm-ps-rl-t	long ago	2017-09-18		PASS
ddb-li-ms-st-t	2017-08-18	2017-09-18		PASS
ddb-li-ms-dt-t	2017-08-21	2017-09-18		PASS
ddb-li-ps-st-t	2017-09-01	2017-09-18		PASS
ddb-li-ps-dt-t	2017-09-18	2017-09-18	8845	TELL
ddb-ci-ms-ll-ct-t	long ago	2017-09-18		PASS
ddb-ci-ms-ll-st-t	long ago	2017-09-18		PASS
ddb-ci-ms-lr-ct-t	long ago	2017-09-18		PASS
ddb-ci-ms-lr-st-t	long ago	2017-09-18		PASS
ddb-ci-ps-ll-it-t	long ago	2017-09-18		PASS
ddb-ci-ps-ll-nt-t	long ago	2017-09-18		PASS
ddb-ci-ps-lr-it-t	long ago	2017-09-18		PASS
ddb-ci-ps-lr-nt-t	long ago	2017-09-18		PASS
ddb-ls-ms-lr-t	long ago	2017-09-18		PASS
ddb-ls-ms-rr-t	long ago	2017-09-18		PASS
ddb-ls-ms-rl-t	long ago	2017-09-18		PASS
ddb-ls-ps-lr-t	2017-09-18	2017-09-18	8733	TELL
ddb-ls-ps-rr-t	2017-09-18	2017-09-18	8733	TELL
ddb-ls-ps-rl-t	2017-09-18	2017-09-18	8733	FEW
drb-rpp-ms-a	long ago	2017-09-18		PASS

Continued on next page

Table 5 – continued from previous page

Test case	Last fail	Last run	Bugs	Status
drb-rph-ms-a	long ago	2017-09-18		PASS
drb-app-ms-a	long ago	2017-09-18		PASS
drb-aph-ms-a	long ago	2017-09-18		PASS
dnb-1n-60k-a	long ago	2017-09-18		PASS
ss-ms-ms-a	long ago	2017-09-18		PASS
ss-ph-ms-a	2017-09-01	2017-09-18	8420	FEW
ss-cl-ms-a	long ago	2017-09-18		PASS
ss-ms-ms-t	long ago	2017-09-18		PASS
ss-ph-ms-t	2017-09-17	2017-09-18	9177	FEW
ss-cl-ms-t	long ago	2017-09-18		PASS
netconf-ba-ms-a	long ago	2017-09-18		PASS
netconf-ok-ms-a	long ago	2017-09-18		PASS
netconf-rr-ms-a	2017-09-06	2017-09-18	9006	TELL
bgp-3n-300k-t-long	2017-09-16	2017-09-16	8318	AKKA
ddb-elm-mc-t-long	2017-08-06	2017-09-16		FEW
drb-rpp-ms-a-long	long ago	2017-09-16		FEW
drb-rph-ms-a-long	2017-08-12	2017-09-16		PASS
dnb-1n-60k-a-long	long ago	2017-09-16		FEW
ss-ph-ms-a-long	2017-09-16	2017-09-16	8420	AKKA
ss-cl-ms-a-long	2017-08-06	2017-09-16		PASS

Note that release, sr1 and sandbox pages contain data from before test implementation and documentation structure were finalized, so there may be inconsistencies.

TODO: Re-test Carbon release and SR1 images (with retrofitted tests where needed) so users can see authoritative test results.

External resources:

- Infrastructure Guide
- Running System Tests
- Test Code Guidelines
- Test Case Expectations
- Robot API docs

3.4 Documentation Guide

This guide provides details on how to contribute to the OpenDaylight documentation. OpenDaylight currently uses reStructuredText for documentation and Sphinx to build it as it is widely-used to provide both HTML and pdf documentation that can be easily versioned alongside the code. It also offers similar syntax to Markdown which is familiar to large numbers of people.

Contents

- *Style Guide*
  - *Formatting Preferences*
  - *Key terms*

    - *Common writing style mistakes*
- *reStructuredText-based Documentation*
    - *Directory Structure*
    - *Documentation Layout and Style*
    - *Troubleshooting*
- *Project Documentation Requirements*
    - *Submitting Documentation Outlines (M3)*
    - *Expected Output From Documentation Project*
    - *Boron Project Documentation Requirements*

3.4.1 Style Guide

This section serves two purposes:

1. A guide for those writing documentation to follow.
2. A guide for those reviewing documentation.

That being said, assuming that the content is usable, the bias should be toward merging it rather than blocking on relatively minor edits.

Formatting Preferences

In general, the documentation team has focused on trying to make sure that the instructions are comprehensible, but not being overly pedantic about these things. Along those lines, while we would prefer the following, generally they aren't a reason to -1 in and of themselves:

- No trailing whitespace
- Line wrapping at something reasonable, i.e., 72–100 characters

Key terms

- Functionality: something useful a project provides abstractly
- Feature: a Karaf feature that somebody could install
- Project: a project within OpenDaylight, projects ship features to provide functionality
- OpenDaylight: this refers to the software we release, use this in place of OpenDaylight controller, the OpenDaylight controller, not ODL, not ODC
    - Since there is a controller project within OpenDaylight, using other terms is hard.

Common writing style mistakes

- In per-project user documentation, you should never say *git clone*, but should assume people have downloaded and installed the controller per the getting started guide and start with `feautre:install <something>`
- Avoid statements which are true about part of OpenDaylight, but not generally true.

  - For example: "OpenDaylight is a NETCONF controller." It is, but that is not all it is.
- In general, developer documentation should target external developers to your project so should talk about what APIs you have and how they could use them. It *should not* document how to contribute to your project.

Grammar Preferences

- Avoid contractions: use cannot instead of can't, it is instead of it's, and the like.

Things to get right with spacing and capitalization

Note that all of these apply when using them in text. If they are used as part of URL, class name, or something similar, use the actual capitalization and spacing.

- ACL: not Acl or acl
- API: not api
- ARP: not Arp or arp
- datastore: not data store, Data Store, or DataStore (unless it's a class/object name)
- IPsec, not IPSEC or ipsec
- IPv4 or IPv6: not Ipv4, Ipv6, ipv4, ipv6, IPV4, or IPV6
- Karaf: not karaf
- Linux: not LINUX or linux
- NETCONF: not Netconf or netconf
- Neutron: not neutron
- OSGi: not osgi or OSGI
- Open vSwitch: not OpenvSwitch, OpenVSwitch, or Open V Switch, etc.
- OpenDaylight: not Opendaylight, Open Daylight, or OpenDayLight, etc.
  - also avoid abbreviations like ODL and ODC
- OpenFlow: not Openflow, Open Flow, openflow, etc.
- OpenStack: not Open Stack or Openstack
- QoS: not Qos, QOS, or qos
- RESTCONF: not Restconf or restconf
- RPC not Rpc or rpc
- URL not Url or url
- VM: not Vm or vm
- YANG: not Yang or yang

3.4.2 reStructuredText-based Documentation

When using reStructuredText, we try to follow the python documentation style guide. See: https://docs.python.org/devguide/documenting.html

The best reference for reStrucutedText syntax seems to be the Sphinx Primer on reStructuredText.

To build and review the reStructuredText documentation locally you must have installed locally:

- python
- python-tox

Which both should be available in most distribution's package managers.

Then simply run tox and open the html produced via your favorite web browser as follows:

```
git clone https://git.opendaylight.org/gerrit/docs
cd docs
git submodule update --init
tox -edocs
firefox docs/_build/html/index.html
```

Note: Make sure to run *tox -edocs* and not just *tox*. See *Make sure you run tox -edocs*

Directory Structure

The directory structure for the reStructuredText documentation is rooted in the `docs` directory inside the `docs git` repository.

Below that there are guides hosted directly in the `docs git` repository and there are guides hosted in remote `git` repositories. Usually those are for project-specific information.

For example here is the directory layout on June, 28th 2016:

```
$ tree -L 2
.
├── Makefile
├── conf.py
├── documentation.rst
├── getting-started-guide
│   ├── api.rst
│   ├── concepts_and_tools.rst
│   ├── experimental_features.rst
│   ├── index.rst
│   ├── installing_opendaylight.rst
│   ├── introduction.rst
│   ├── karaf_features.rst
│   ├── other_features.rst
│   ├── overview.rst
│   └── who_should_use.rst
├── index.rst
├── make.bat
├── opendaylight-with-openstack
│   ├── images
│   ├── index.rst
│   ├── openstack-with-gbp.rst
```

(continues on next page)

(continued from previous page)

```
│   ├── openstack-with-ovsdb.rst
│   └── openstack-with-vtn.rst
└── submodules
    └── releng
        └── builder
```

The `getting-started-guide` and `opendaylight-with-openstack` directories correspond to two guides hosted in the `docs` repository, while the `submodules/releng/builder` directory houses documentation for the RelEng/Builder project.

Inside each guide there is usually an `index.rst` file which then includes other files using a `toctree` directive. For example:

```
.. toctree::
   :maxdepth: 1

   getting-started-guide/index
   opendaylight-with-openstack/index
   submodules/releng/builder/docs/index
```

This creates a table of contents on that page where each heading of the table of contents is the root of the files that are included.

Note: When including rst files using `toctree` omit the .rst at the end of the file name.

Adding a submodule

If you want to import a project underneath the documentation project so that the docs can be kept in the separate repo, you can do it using the `git submodule add` command as follows:

```
git submodule add -b master ../integration/packaging docs/submodules/integration/
↪packaging
git commit -s
```

Note: Most projects will not want to use `-b master`, but instead use the branch `.`, which will make track whatever branch of the documentation project you happen to be on.

Unfortunately, `-b .` doesn't work, so you have to manually edit the `.gitmodules` file to add `branch = .` and then commit it. Something like:

```
<edit the .gitmodules file>
git add .gitmodules
git commit --amend
```

When you're done you should have a git commit something like:

```
$ git show
commit 7943ce2cb41cd9d36ce93ee9003510ce3edd7fa9
Author: Daniel Farrell <dfarrell@redhat.com>
Date:   Fri Dec 23 14:45:44 2016 -0500
```

(continues on next page)

(continued from previous page)

```
    Add Int/Pack to git submodules for RTD generation

    Change-Id: I64cd36ca044b8303cb7fc465b2d91470819a9fe6
    Signed-off-by: Daniel Farrell <dfarrell@redhat.com>

diff --git a/.gitmodules b/.gitmodules
index 91201bf6..b56e11c8 100644
--- a/.gitmodules
+++ b/.gitmodules
@@ -38,3 +38,7 @@
        path = docs/submodules/ovsdb
        url = ../ovsdb
        branch = .
+[submodule "docs/submodules/integration/packaging"]
+       path = docs/submodules/integration/packaging
+       url = ../integration/packaging
+       branch = master
diff --git a/docs/submodules/integration/packaging b/docs/submodules/integration/
↪packaging
new file mode 160000
index 00000000..fd5a8185
--- /dev/null
+++ b/docs/submodules/integration/packaging
@@ -0,0 +1 @@
+Subproject commit fd5a81853e71d45945471d0f91bbdac1a1444386
```

As usual, you can push it to Gerrit with `git review`.

Important: It's critical that the Gerrit patch be merged before the git commit hash of the submodule changes. Otherwise, Gerrit won't be able to automatically keep it up-to-date for you.

Documentation Layout and Style

As mentioned previously we try to follow the python documentation style guide which defines a few types of sections:

```
# with overline, for parts
* with overline, for chapters
=, for sections
-, for subsections
^, for subsubsections
", for paragraphs
```

We try to follow the following structure based on that recommendation:

```
docs/index.rst                  -> entry point
docs/____-guide/index.rst       -> part
docs/____-guide/<chapter>.rst   -> chapter
```

In the ____-guide/index.rst we use the # with overline at the very top of the file to determine that it is a part and then within each chapter file we start the document with a section using * with overline to denote that it's the chapter heading and then everything in the rest of the chapter should use:

```
=, for sections
-, for subsections
^, for subsubsections
", for paragraphs
```

Referencing Sections

It's pretty common to want to reference another location in the OpenDaylight documentation and it's pretty easy to do with reStructuredText. This is a quick primer, more information is in the Sphinx section on Cross-referencing arbitrary locations.

Within a single document, you can reference another section simply by:

```
This is a reference to `The title of a section`_
```

Assuming that somewhere else in the same file there a is a section title something like:

```
The title of a section
^^^^^^^^^^^^^^^^^^^^^^
```

It's typically better to use `:ref:` syntax and labels to provide links as they work across files and are resilient to sections being renamed. First, you need to create a label something like:

```
.. _a-label:

The title of a section
^^^^^^^^^^^^^^^^^^^^^^
```

Note: The underscore (_) before the label is required.

Then you can reference the section anywhere by simply doing:

```
This is a reference to :ref:`a-label`
```

or:

```
This is a reference to :ref:`a section I really liked <a-label>`
```

Note: When using `:ref:`-style links, you don't need a trailing underscore (_).

Because the labels have to be unique, it usually makes sense to prefix the labels with the project name to help share the label space, e.g., `sfc-user-guide` instead of just `user-guide`.

Troubleshooting

Nested formatting doesn't work

As stated in the reStructuredText guide, inline markup for bold, italic, and fixed-width can't be nested. Further, it can't be mixed with hyperlinks, so you can't have bold text link somewhere.

This is tracked in a Docutils FAQ question, but there is no clear current plan to fix this.

Make sure you've cloned submodules

If you see an error like this:

```
./build-integration-robot-libdoc.sh: line 6: cd: submodules/integration/test/csit/
↪libraries: No such file or directory
Resource file '*.robot' does not exist.
```

It means that you haven't pulled down the git submodule for the integration/test project. The fastest way to do that is:

```
git submodule update --init
```

In some cases, you might wind up with submodules which are somehow out-of-sync and in that case, the easiest way to fix it is delete the submodules directory and then re-clone the submodules:

```
rm -rf docs/submodules/
git submodule update --init
```

Warning: This will delete any local changes or information you made in the submodules. This should only be the case if you manually edited files in that directory.

Make sure you run tox -edocs

If you see an error like:

```
ERROR:   docs: could not install deps [-rrequirements.txt]; v = InvocationError('/
↪Users/ckd/git-reps/docs/.tox/docs/bin/pip install -rrequirements.txt (see /Users/
↪ckd/git-reps/docs/.tox/docs/log/docs-1.log)', 1)
ERROR:   docs-linkcheck: could not install deps [-rrequirements.txt]; v =␣
↪InvocationError('/Users/ckd/git-reps/docs/.tox/docs-linkcheck/bin/pip install -
↪rrequirements.txt (see /Users/ckd/git-reps/docs/.tox/docs-linkcheck/log/docs-
↪linkcheck-1.log)', 1)
```

It usually means you ran *tox* and not *tox -edocs*, which will result in running jobs inside submodules which aren't supported by the environment defined by the *requirements.txt* file in the documentation tox setup. Just run tox -edocs and it should be fine.

Clear your tox directory and try again

Sometimes, tox will not detect when your `requirements.txt` file has changed and so will try to run things without the correct dependencies. This usually manifests as `No module named X` errors or an `ExtensionError` and can be fixed by deleting the `.tox` directory and building again:

```
rm -rf .tox
tox -edocs
```

Builds on Read the Docs

It appears as though the Read the Docs builds don't automatically clear the file structure between builds and clones. The result is that you may have to clean up the state of old runs of the build script.

As an example, this patch: https://git.opendaylight.org/gerrit/41679

Finally fixed the fact that our builds for failing because they were taking too long by removing directories of generated javadoc that were present from previous runs.

3.4.3 Project Documentation Requirements

Submitting Documentation Outlines (M3)

1. Determine the features your project will have and which ones will be ''user-facing''.
   - In general, a feature is user-facing if it creates functionality that a user would direction interact with.
   - For example, `odl-openflowplugin-flow-services-ui` is likely user-facing since it installs user-facing OpenFlow features, while `odl-openflowplugin-flow-services` is not because it provides only developer-facing features.
2. Determine pieces of documentation you need provide based on the features your project will have and which ones will be user-facing.
   - The kinds of required documentation can be found below in the *Requirements for projects* section.
   - Note that you might need to create multiple different documents for the same kind of documentation. For example, the controller project will likely want to have a developer section for the config subsystem as well as a for the MD-SAL.
3. Clone the docs repo: `git clone https://git.opendaylight.org/gerrit/docs`
4. For each piece of documentation find the corresponding template in the docs repo.
   - For user documentation: `docs.git/docs/templates/template-user-guide.rst`
   - For developer documentation: `ddocs/templates/template-developer-guide.rst`
   - For installation documentation (if any): `docs/templates/template-install-guide.rst`

Note: You can find the rendered templates here:

<Feature> User Guide

Refer to this template to identify the required sections and information that you should provide for a User Guide. The user guide should contain configuration, administration, management, using, and troubleshooting sections for the feature.

Overview

Provide an overview of the feature and the use case. Also include the audience who will use the feature. For example, audience can be the network administrator, cloud administrator, network engineer, system administrators, and so on.

<Feature> Architecture

Provide information about feature components and how they work together. Also include information about how the feature integrates with OpenDaylight. An architecture diagram could help.

Note: Please *do not* include detailed internals that somebody using the feature wouldn't care about. For example, the fact that there are four layers of APIs between a user command and a message being sent to a device is probably not useful to know unless they have some way to influence how those layers work and a reason to do so.

Configuring <feature>

Describe how to configure the feature or the project after installation. Configuration information could include day-one activities for a project such as configuring users, configuring clients/servers and so on.

Administering or Managing <feature>

Include related command reference or operations that you could perform using the feature. For example viewing network statistics, monitoring the network, generating reports, and so on.

For example:

To configure L2switch components perform the following steps.

(a) Step 1:

(b) Step 2:

(c) Step 3:

Tutorials

optional

If there is only one tutorial, you skip the "Tutorials" section and instead just lead with the single tutorial's name. If you do, also increase the header level by one, i.e., replace the carets (^^^) with dashes (- - -) and the dashes with equals signs (===).

<Tutorial Name>

Ensure that the title starts with a gerund. For example using, monitoring, creating, and so on.

Overview

An overview of the use case.

Prerequisites

Provide any prerequisite information, assumed knowledge, or environment required to execute the use case.

Target Environment

Include any topology requirement for the use case. Ideally, provide visual (abstract) layout of network diagrams and any other useful visual aides.

Instructions

Use case could be a set of configuration procedures. Including screenshots to help demonstrate what is happening is especially useful. Ensure that you specify them separately. For example:

Setting up the VM

To set up a VM perform the following steps.

(a) Step 1

(b) Step 2

(c) Step 3

Installing the feature

To install the feature perform the following steps.

(a) Step 1

(b) Step 2

(c) Step 3

Configuring the environment

To configure the system perform the following steps.

(a) Step 1

(b) Step 2

(c) Step 3

<Feature> Developer Guide

Overview

Provide an overview of the feature, what it logical functionality it provides and why you might use it as a developer. To be clear the target audience for this guide is a developer who will be *using* the feature to build something separate, but *not* somebody who will be developing code for this feature itself.

Note: More so than with user guides, the guide may cover more than one feature. If that is the case, be sure to list all of the features this covers.

<Feature> Architecture

Provide information about feature components and how they work together. Also include information about how the feature integrates with OpenDaylight. An architecture diagram could help. This may be the same as the diagram used in the user guide, but it should likely be less abstract and provide more information that would be applicable to a developer.

Key APIs and Interfaces

Document the key things a user would want to use. For some features, there will only be one logical grouping of APIs. For others there may be more than one grouping.

Assuming the API is MD-SAL- and YANG-based, the APIs will be available both via RESTCONF and via Java APIs. Giving a few examples using each is likely a good idea.

API Group 1

Provide a description of what the API does and some examples of how to use it.

API Group 2

Provide a description of what the API does and some examples of how to use it.

API Reference Documentation

Provide links to JavaDoc, REST API documentation, etc.

<Feature> Installation Guide

Note: *Only* use this template if installation is more complicated than simply installing a feature in the Karaf distribution. Otherwise simply provide the names of all user-facing features in your M3 readout.

This is a template for installing a feature or a project developed in the ODL project. The *feature* could be interfaces, protocol plug-ins, or applications.

Overview

Add overview of the feature. Include Architecture diagram and the positioning of this feature in overall controller architecture. Highlighting the feature in a different color within the overall architecture must help. Include information to describe if the project is within ODL installation package or to be installed separately.

Pre Requisites for Installing <Feature>

- Hardware Requirements
- Software Requirements

Preparing for Installation

Include any pre configuration, database, or other software downloads required to install <feature>.

Installing <Feature>

Include if you have separate procedures for Windows and Linux

Verifying your Installation

Describe how to verify the installation.

Troubleshooting

optional

Text goes here.

Post Installation Configuration

Post Installation Configuration section must include some basic (must-do) procedures if any, to get started.

Mandatory instructions to get started with the product.

- Logging in
- Getting Started
- Integration points with controller

Upgrading From a Previous Release

Text goes here.

Uninstalling <Feature>

Text goes here.

---

5. Copy the template into the appropriate directory for your project.
   - For user documentation: `docs.git/docs/user-guide/${feature-name}-user-guide.rst`
   - For developer documentation: `docs.git/docs/developer-guide/${feature-name}-developer-guide.rst`
   - For installation documentation (if any): `docs.git/docs/getting-started-guide/project-specific-guides/${project-name}.rst`

Note: These naming conventions aren't set in stone, but do help. If you think there's a better name, use it and we'll give feedback on the gerrit patch.

6. Edit the template to fill in the outline of what you will provide using the suggestions in the template. If you feel like a section isn't needed, feel free to omit it.
7. Link the template into the appropriate core rst file
   - For user documentation: `docs.git/docs/user-guide/index.rst`
   - For developer documentation: `docs.git/docs/developer-guide/index.rst`
   - For installation documentation (if any): `docs.git/docs/getting-started-guide/project-specific-guides/index.rst`
   - In each file, it should be pretty clear what line you need to add. In general if you have an rst file `project-name.rst`, you include it by adding a new line `project-name` without the `.rst` at the end.
8. Make sure the documentation project still builds.
   - Run `tox -edocs` from the root of the cloned docs repo.
     - After that, you should be able to find the HTML version of the docs at `docs.git/docs/_build/html/index.html`.
     - See *reStructuredText-based Documentation* for more details about building the docs.
   - The *reStructuredText Troubleshooting* section provides common errors and solutions.
   - If you still have problems e-mail the documentation group at documentation@lists.opendaylight.org
9. Commit and submit the patch
   (a) Commit using:

```
git add --all && git commit -sm "Documentation outline for ${project-
↪shortname}"
```

   (b) Submit using:

```
git review
```

   See the Git-review Workflow page if you don't have git-review installed.
10. Wait for the patch to be merged or to get feedback
    - If you get feedback, make the requested changes and resubmit the patch.
    - When you resubmit the patch, it's helpful if you also post a +0 reply to the gerrit saying what patch set you just submitted and what you fixed in the patch set.
    - The documentation team will also be creating (or asking projects to create) small groups of 2-4 projects that will peer review each other's documentation. Patches which have seen a few cycles of peer review will be prioritized for review and merge by the documentation team.

Expected Output From Documentation Project

The expected output is (at least) 3 PDFs and equivalent web-based documentation:

- User/Operator Guide

- Developer Guide
- Installation Guide

These guides will consist of "front matter" produced by the documentation group and the per-project/per-feature documentation provided by the projects. Note that this is intended to be who is responsible for the documentation and should not be interpreted as preventing people not normally in the documentation group from helping with "front matter" nor preventing people from the documentation group from helping with per-project/per-feature documentation.

Boron Project Documentation Requirements

Kinds of Documentation

These are the expected kinds of documentation and target audiences for each kind.

- **User/Operator:** for people looking to use the feature w/o writing code
  - Should include an overview of the project/feature
  - Should include description of availbe configuration options and what they do
- **Developer:** for people looking to use the feature in code w/o modifying it
  - Should include API documentation, e.g., enunciate for REST, Javadoc for Java, ??? for REST-CONF/models
- **Contributor:** for people looking to extend or modify the feature's source code
- **Installation:** for people looking for instructions to install the feature after they have downloaded the ODL release
  - For most projects, this will be just a list of top-level features and options
    - As an example, l2switch-switch as the top-level feature with the -rest and -ui options
    - We'd also like them to note if the options should be checkboxes (i.e., they can each be turned on/off independently) or a drop down (i.e., at most one can be selected)
    - What other top-level features in the release are incompatible with each feature
    - This will likely be presented as a table in the documentation and the data will likely also be consumed by automated installers/configurators/downloaders
  - For some projects, there is extra installation instructions (for external components) and/or configuration
    - In that case, there will be a (sub)section in the documentation describing this process.
- **HowTo/Tutorial:** walk throughs and examples that are not general-purpose documentation
  - Generally, these should be done as a (sub)section of either user/operator or developer documentation.
  - If they are especially long or complex, they may belong on their own
- **Release Notes:**
  - Release notes are required as part of each project's release review. They must also be translated into reStructuredText for inclusion in the formal documentation.

Requirements for projects

Projects MUST do the following

- Provide reStructuredText documentation including

  - Developer documentation for every feature
    * Most projects will want to logically nest the documentation for individual features under a single project-wide chapter or section
    * This can be provided as a single .rst file or multiple .rst files if the features fall into different groups
    * This should start with ~300 word overview of the project and include references to any automatically-generated API documentation as well as more general developer information (see *Kinds of Documentation*).
  - User/Operator documentation for every every user-facing feature (if any)
    * ''Note: This should be per-feature, not per-project. User's shouldn't have to know which project a feature came from.''
    * Intimately related features, e.g., l2switch-switch, l2switch-switch-rest, and l2switch-switch-ui, can be documented as one noting the differences
    * This can be provided as a single .rst file or multiple .rst files if the features fall into different groups
  - Installation documentation
    * Most projects will simply provide a list of user-facing features and options. See *Kinds of Documentation* above.
  - Release Notes (both on the wiki and reStructuredText) as part of the release review.
- This documentation will be contributed to the docs repo (or possibly imported from the project's own repo with tooling that is under development)
  - Projects MAY be ENCOURGAGED to instead provide this from their own repository if the tooling is developed
  - Projects choosing to meet the requirement this way MUST provide a patch to docs repo to import the project's documentation
- Projects MUST cooperate with the documentation group on edits and enhancements to documentation
  - Note that the documentation team will also be creating (or asking projects to create) small groups of 2-4 projects that will peer review each other's documentation. Patches which have seen a few cycles of peer review will be prioritized for review and merge by the documentation team.

Timeline for Deliverables from Projects

- **M3:** Documentation Started
  - Identified the kinds of documentation that will be provided and for what features
    * Release Notes are not required until release reviews at **RC2**
  - Created the appropriate .rst files in the docs repository (or their own repository if the tooling is available)
  - Have an outline for the expected documentation in those .rst files including the relevant (sub)sections and a sentence or two explaining what will go there
    * Obviously, providing actual documentation in the (sub)sections is encouraged and meets this requirement
  - Milestone readout should include
    1. the list of kinds of documentation
    2. the list of corresponding .rst files and their location, e.g., repo and path

3. the list of commits creating those .rst files
4. the current word counts of those .rst files

- **M4:** Documentation Continues
  - The readout at M4 should include the word counts of all .rst files with links to commits
  - The goal is to have draft documentation complete so that the documentation group can comment on it.
- **M5:** Documentation Complete
  - All (sub)sections in all .rst files have complete, readable, usable content.
  - Ideally, there should have been some interaction with the documentation group about any suggested edits and enhancements
- **RC2:** Release notes
  - Projects must provide release notes as .rst pushed to integration (or locally in the project's repository if the tooling is developed)

3.5 OpenDaylight Release Process Guide

3.5.1 Overview

This guide provides details on various processes related to OpenDaylight's release process and attempts to document the steps used by OpenDaylight Release Engineers to perform release operations.

3.5.2 Processes

Autorelease

The Release Engineering - Autorelease project is targeted at building the artifacts that are used in the release candidates and final full release.

- Open Gerrit Patches
- Jenkins Jobs

Cloning Autorelease

To clone all the autorelease repo including it's submodules simply run the clone command with the ''–recursive'' parameter.

```
git clone --recursive https://git.opendaylight.org/gerrit/releng/autorelease
```

If you forgot to add the –recursive parameter to your git clone you can pull the submodules after with the following commands.

```
git submodule init
git submodule update
```

Creating Autorelease - Release and RC build

An autorelease release build comes from the autorelease-release-<branch> job which can be found on the autorelease tab in the releng master:

- https://jenkins.opendaylight.org/releng/view/autorelease/

For example to create a Boron release candidate build launch a build from the autorelease-release-boron job by clicking the ‘‘Build with Parameters’’ button on the left hand menu:

- https://jenkins.opendaylight.org/releng/view/autorelease/job/autorelease-release-boron/

Note: The only field that needs to be filled in is the ‘‘RELEASE_TAG’’, leave all other fields to their default setting. Set this to Boron, Boron-RC0, Boron-RC1, etc... depending on the build you’d like to create.

Adding Autorelease staging repo to settings.xml

If you are building or testing this release in such a way that requires pulling some of the artifacts from the Nexus repo you may need to modify your settings.xml to include the staging repo URL as this URL is not part of ODL Nexus’ public or snapshot groups. If you’ve already cloned the recommended settings.xml for building ODL you will need to add an additional profile and activate it by adding these sections to the “<profiles>” and “<activeProfiles>” sections (please adjust accordingly).

Note:

- This is an example and you need to “Add” these example sections to your settings.xml do not delete your existing sections.
- The URLs in the <repository> and <pluginRepository> sections will also need to be updated with the staging repo you want to test.

```
<profiles>
  <profile>
    <id>opendaylight-staging</id>
    <repositories>
      <repository>
        <id>opendaylight-staging</id>
        <name>opendaylight-staging</name>
        <url>https://nexus.opendaylight.org/content/repositories/
↪automatedweeklyreleases-1062</url>
        <releases>
          <enabled>true</enabled>
          <updatePolicy>never</updatePolicy>
        </releases>
        <snapshots>
          <enabled>false</enabled>
        </snapshots>
      </repository>
    </repositories>
    <pluginRepositories>
      <pluginRepository>
        <id>opendaylight-staging</id>
        <name>opendaylight-staging</name>
```

(continues on next page)

(continued from previous page)

```
        <url>https://nexus.opendaylight.org/content/repositories/
↪automatedweeklyreleases-1062</url>
        <releases>
          <enabled>true</enabled>
          <updatePolicy>never</updatePolicy>
        </releases>
        <snapshots>
          <enabled>false</enabled>
        </snapshots>
      </pluginRepository>
    </pluginRepositories>
  </profile>
</profiles>

<activeProfiles>
  <activeProfile>opendaylight-staging</activeProfile>
</activeProfiles>
```

Project lifecycle

This page documents the current rules to follow when adding and removing a particular project to Simultaneous Release (SR).

List of states

The state names are short negative phrases describing what is missing to progress to the following state.

- **non-existent** The project is not recognized by Technical Steering Committee (TSC) to be part of OpenDaylight (ODL).
- **non-participating** The project is recognized byt TSC to be an ODL project, but the project has not confirmed participation in SR for given release cycle.
- **non-building** The recognized project is willing to participate, but its current codebase is not passing its own merge job, or the project artifacts are otherwise unavailable in Nexus.
- **not-in-autorelease** Project merge job passes, but the project is not added to autorelease (git submodule, maven module, validate-autorelease job passes).
- **repo-not-in-integration** Project is added do autorelease, but integration/distribution:features-index is not listing all its public feature repositories.
- **distribution-check-not-passing** Project is in autorelease, but its distribution-check job is either not running, or it is failing for any reason.
- **feature-not-in-integration** Feature repositories are referenced, distribution-check job is passing, but some user-facing features are absent from integration/distribution:features-test
- **feature-is-experimental** All user-facing features are in features-test, but at least one of the corresponding functional CSIT jobs does not meet integration/test requirements.
- **ready**

Note: A project may change its state in both directions, this list is to make sure a project is not left in an invalid state, for example distribution referencing feature repositories, but without passing distribution-check job.

Namespaces

Project namespaces in OpenDaylight are used to ensure projects do not have name collisions in code and packages. OpenDaylight enforces namespaces in Nexus using the following patterns:

- ^/org.opendaylight.PROJECT/.*
- ^/org/opendaylight/PROJECT/.*

Where PROJECT is the name of an OpenDaylight project.

In cases where a project has a sub-project we recommend adding an additional level to the path for example *org.opendaylight.integration.test* however no strong enforcement is currently enforced and some projects do this already internally.

This restriction applies to all site repositories in Nexus as well in the event that a project wishes to push a static web site into their allocated site path.

Maven / Java

Maven has a built in namespace routing using <groupId> field in pom files. For example:

```
<project>
  <groupId>org.opendaylight.odlparent</groupId>
  <artifactId>odlparent-lite</artifactId>
  <version>1.8.0-SNAPSHOT</version>
</project>
```

Python

Python projects typically publish to artifacts to PyPi and use their shortname for modules rather than a full path like Java projects do.

setup.py:

```
setup(
    name='spectrometer',
)
```

The structure of a Python project typically determines it's package routing. So a project package spectrometer.reporttool might have a layout like this inside their project root.

```
./  # This is the root of the repository
./setup.py
./spectrometer
./spectrometer/__init__.py
./spectrometer/reporttool
./spectrometer/reporttool/__init__.py
```

Branch Cutting

This page documents the current branch cutting tasks that are needed to be performed at various milestones and which team has the necessary permissions in order to perform the necessary task in Parentheses.

M5 Offset 2

JJB

- Export ${NEXT_RELEASE} and ${CURR_RELEASE} with new and current release names. **(releng/builder committers)**

```
export NEXT_RELEASE="Nitrogen"
export CURR_RELEASE="Carbon"
```

- Change JJB yaml files from *stream:carbon* branch pointer from *master* -> *stable/${CURR_RELEASE,,}* and create new *stream: ${NEXT_RELEASE,,}* branch pointer to branch master. This requires handling two different file formats interspersed with in autorelease projects. **(releng/builder committers)**

```
stream:
  - Nitrogen:
      branch: master
  - Carbon:
      branch: stable/carbon
```

```
- project:
    name: aaa-carbon
    jobs:
      - '{project-name}-verify-{stream}-{maven}-{jdks}'
    stream: nitrogen
    branch: master
```

  - The above manual process of updating individual files is automated with the script. **(releng/builder committers)**

```
cd builder/scripts/branch_cut
./branch_cutter.sh -n $NEXT_RELEASE -c $CURR_RELEASE
```

- Review and submit the changes to releng/builder project. **(releng/builder committers)**

Autorelease

- Block submit permissions for registered users and elevate RE's committer rights on gerrit. **(Helpdesk)**

Note: Enable Exclusive checkbox override any existing persmissions.

- Setup releng/autorelease repository. **(Release Engineering Team)**

```
git review -s
git submodule foreach 'git review -s'
git checkout master
git submodule foreach 'git checkout master'
git pull --rebase
git submodule foreach 'git pull --rebase'
```

- Create stable/${CURR_RELEASE} branches based on HEAD master. **(Release Engineering Team)**

```
git submodule foreach 'git checkout -b stable/${CURR_RELEASE,,} origin/master'
git push gerrit stable/${CURR_RELEASE,,}
git submodule foreach 'git push gerrit stable/${CURR_RELEASE,,}'
```

- Enable create reference permissions on gerrit for RE's to submit .gitreview patches. **(Helpdesk)**

Note: Enable Exclusive checkbox override any existing persmissions.

- Contribute .gitreview updates to stable/${CURR_RELEASE,,}. **(Release Engineering Team)**

```
git submodule foreach sed -i -e "s#defaultbranch=master#defaultbranch=stable/$
↪{CURR_RELEASE,,}#" .gitreview
git submodule foreach git commit -asm "Update .gitreview to stable/${CURR_RELEASE,
↪,}"
git submodule foreach 'git review -t ${CURR_RELEASE,,}-branch-cut'
sed -i -e "s#defaultbranch=master#defaultbranch=stable/${CURR_RELEASE,,}#" .
↪gitreview
git add .gitreview
git commit -s -v -m "Update .gitreview to stable/${CURR_RELEASE,,}"
git review -t  ${CURR_RELEASE,,}-branch-cut
```

- Merge all .gitreview patches submitted in the above step. **(Release Engineering Team)**
- Remove create reference permissions set on gerrit for RE's. **(Helpdesk)**
- Version bump master by x.(y+1).z. **(Release Engineering Team)**

```
git checkout master
git submodule foreach 'git checkout master'
pip install lftools
lftools version bump ${CURR_RELEASE}
```

- Exclude version bump changes to release notes. **(Release Engineering Team)**

```
git checkout pom.xml scripts/
```

- Push version bump master changes to gerrit. **(Release Engineering Team)**

```
git submodule foreach 'git commit -asm "Bump versions by x.(y+1).z for next dev␣
↪cycle"'
git submodule foreach 'git review -t nitrogen-br-cut'
```

- Merge all version bump patches in the order of dependencies. **(Release Engineering Team)**
- Re-enable submit permissions for registered users and disable elevated RE committer rights on gerrit. **(Helpdesk)**
- Notify release list on branch cutting work completion. **(Release Engineering Team)**

Simultaneous Release

This page explains how the OpenDaylight release process works once the TSC has approved a release.

Preparations

After release candidate is built gpg sign artifacts using odlsign-bulk script in **releng/builder/scripts**.

```
cd scripts/
./odlsign-bulk STAGING_REPO_ID  # eg. autorelease-1367
```

Releasing OpenDaylight

- Block submit permissions for registered users and elevate RE's committer rights on gerrit.

Note: Enable Exclusive checkbox

- Export ${RELEASE} and ${BUILDNUM} with current release name and build number.

```
export RELEASE=Beryllium-SR4
export BRANCH=${RELEASE//-*}
export BUILDNUM=55
```

- Nexus: click release for staging repo **(Helpdesk)**
- Send email to Helpdesk with binary URL to update website **(Helpdesk)**

- Send email to TSC and Release mailing lists announcing release binaries location **(Release Engineering Team)**
- Clone autorelease repository. **(Release Engineering Team)**

```
git clone --recursive https://git.opendaylight.org/gerrit/releng/autorelease
```

- Checkout autorelease and switch to release branch eg stable/carbon **(Release Engineering Team)**

```
git checkout -b stable/${BRANCH,,} origin/stable/${BRANCH,,}
git submodule update --init
git submodule foreach git checkout -b stable/${BRANCH,,} origin/stable/${BRANCH,,}
```

- Make sure your git repo is setup to push (use git-review)

```
git review -s
git submodule foreach 'git review -s'
```

- Download patches (*.bundle files and taglist.log.gz) from log server.

```
mkdir /tmp/patches && cd /tmp/patches
wget https://logs.opendaylight.org/releng/jenkins092/autorelease-release-${BRANCH,
↪,}/${BUILDNUM}/archives/all-bundles.tar.gz
gunzip all-bundles.tar.gz
wget https://logs.opendaylight.org/releng/jenkins092/autorelease-release-${BRANCH,
↪,}/${BUILDNUM}/archives/patches/taglist.log.gz
gunzip taglist.log.gz
```

- Run the following commands for every project in the release, to apply patches to each project directory.

```
pip install lftools
lftools version patch ${RELEASE}
git review -y -t ${RELEASE}
git push gerrit release/${RELEASE,,}
```

- Merge all patches on gerrit in the order generated by merge-order.log
- Tag autorelease too

```
git checkout `cat /tmp/patches/taglist.log | grep autorelease | awk '{print $2}'`
git submodule foreach git checkout release/${RELEASE,,}
git commit -asSm "Release ${RELEASE}"
git tag -asm "OpenDaylight ${RELEASE} release" release/${RELEASE,,}
git push gerrit release/${RELEASE,,}
```

- Re-enable submit permissions for registered users and disable elevated RE committer rights on gerrit.
- Release notes is auto generated by job autorelease-generate-release-notes-${BRANCH,,} triggered at the end of every autorelease build. The release notes file (release_notes.rst) is available under archives.

  Alternatively, release notes can also be manually generated with the script. **(Release Engineering Team)**

```
git checkout stable/${BRANCH,,}
cd scripts/release_notes_management/ && ./build.sh
```

  The output file (release_notes.rst) generated by the build script is available under autorelease/scripts/release_notes_management/projects/.

- Send email to release/tsc/dev notifying tagging and version bump complete **(Release Engineering Team)**

3.6 Spectrometer Documentation

Contents:

3.6.1 Quick Start Guide

The Spectrometer project consists of two sub-projects, the `server` and `web`.

Server side is Python driven and provides the API to collect Git and Gerrit statistics for various OpenDaylight projects.

The web project is NodeJS/React based and provides the visualization by using the APIs provided by the server side.

In order to run the application, you need to install both `server` and `web` sub-projects.

This Quick Started Guide assumes you have Python3 and NodeJS 4.3 installed. To install NodeJS using NVM, see Web > Installation section below.

The Spectrometer project collects data from repositories located locally in your system.

Setup spectrometer-server

Installing spectrometer from pypi is simple and will get you the latest version that is released. Then create a config.py file in /etc/spectrometer/config.py (Example file can be found here)

```
pip install spectrometer
sudo mkdir /etc/spectrometer
sudo vi /etc/spectrometer/config.py
spectrometer server start
```

Verify that spectrometer-server is running by going to **http://localhost:5000**. You should see a Hello World page.

Setup spectrometer-web

Spectrometer Web is still in development so you will need to install it from Git at the time being as there is no package for it yet.

```
git clone https://git.opendaylight.org/gerrit/spectrometer.git
cd spectrometer/web
npm install
npm start
```

Goto **http://localhost:8000**

Testing the setup

By default the OpenDaylight project repositories will be mirrored every 5 minutes (300s), so if this is the first time starting you may have to wait until all repos are mirrored before you can exercise some of the apis.

Once the repos are mirrored you can try a few basic examples to make sure things are working properly:

Examples:

```
http://127.0.0.1:5000/gerrit/branches?project=controller
http://127.0.0.1:5000/gerrit/projects
http://127.0.0.1:5000/git/commits?project=integration/packaging
```

The full Rest APIs are documented here: https://opendaylight-spectrometer.readthedocs.io/en/latest/restapi.html

3.6.2 User Guide

Spectrometer consists of 3 components:

- Spectrometer API Server (backend)
- Spectrometer Web Server (frontend)
- Spectrometer Report Tool

This guide will describe the uses of the 3 systems.

Spectrometer API Server

Production Deployment

When running in production the recommended way is to deploy with gunicorn.

```
gunicorn -b 0.0.0.0:5000 'spectrometer:run_app()'
```

If deploying behind a proxy under a sub-directory additional configuration is necessary for gunicorn application to operate correctly.

example-nginx:

```
location /api {
    proxy_pass         http://127.0.0.1:5000;
    proxy_redirect     http://127.0.0.1:5000/api/ http://$host/api/;

    proxy_set_header   Host             $host;
    proxy_set_header   X-Real-IP        $remote_addr;
    proxy_set_header   X-Forwarded-For  $proxy_add_x_forwarded_for;
    proxy_set_header   SCRIPT_NAME      /api;
}
```

Logging

Spectrometer logs to /var/log/spectrometer by default but that directory must be writeable by the spectrometer user.

```
sudo chown spectrometer /var/log/spectrometer
```

It is possible to override the default log directory by configuring the LOG_DIR parameter in config.py.

```
LOG_DIR = '/path/to/log/directory'
```

Spectrometer Web Server

TODO

Spectrometer Report Tool

The Spectrometer Report Tool can be used to generate reports between 2 reference points in time. Reference points are git commit hashs, branches, or tags. A project like OpenDaylight that tags projects with the same tag name for every release can use this tool to Generate release reports.

```
# spectrometer reporttool full <ref1> <ref2>
spectrometer reporttool --server-url=https://spectrometer.opendaylight.org/api full␣
↪release/beryllium-sr2 release/beryllium-sr1
```

3.6.3 Project Info Specification

Spectrometer supports a PROJECT_INFO.yaml file placed in the root of a project repo. This file is used by spectrometer to parse meta information about the project including things like project description, project contact, committers irc, mailing lists, release names, etc...

```
# This file is used by Spectrometer to determine project meta information
# Please refer to the spec file located here:
# https://opendaylight-spectrometer.readthedocs.io/en/latest/project-info-spec.html

name: spectrometer
display-name: Spectrometer
creation-date: 2015-11-19
termination-date: n/a
description: |
    This is an example summary description of project

    After leaving a blank line in the description we can provide a longer
    more detailed description of the project.

    The details can be as many lines as necessary.
primary-contact: Firstname Lastname <first.last@example.com>
project-lead: Firstname Lastname <first.last@example.com>
categories:
    - application
    - community
    - documentation
    - extensions
    - kernel
    - library
    - protocols
    - services
committers:
    - Firstname Lastname <first.last@example.com>
    - Another Committer <another.committer@example.com>
# When Committers who have made significant contributions to OpenDaylight
# become inactive and thus no longer committers. This key can be used to
# acknowledge their huge contributions by appointing them to Committer
# Emeritus status.
committers-emeritus:
    - Firstname Lastname <first.last@example.com>
contributors:
    - Firstname Lastname <first.last@example.com>
    - Another Contributor <another.contributor@example.com>
wiki: https://wiki.example.org/project
irc: irc://irc.freenode.net/opendaylight-spectrometer
```

(continues on next page)

(continued from previous page)

```
mailing-lists:
    - email: spectrometer-dev@lists.opendaylight.org
      archives: http://lists.opendaylight.org/pipermail/spectrometer-dev/
    - email: spectrometer-users@lists.opendaylight.org
      archives: http://lists.opendaylight.org/pipermail/spectrometer-users/
ci-server: https://jenkins.opendaylight.org
issue-tracker: https://bugs.opendaylight.org
static-analysis: https://sonar.opendaylight.org
repository: https://git.opendaylight.org/gerrit/#/admin/projects/spectrometer
meetings: |
    Free from text field for providing meeting information.
    It can be multiple lines long as necessary.
releases:
    - helium
    - lithium
    - beryllium
    - boron
```

Required fields:

- name
- creation_date
- description
- primary_contact
- project_lead

3.6.4 Documentation Guide

This guide provides details on how to contribute to the documetantion of Spectrometer. The style guide we follow for documentation is the python documentation style guide. See:

https://docs.python.org/devguide/documenting.html

To build and review the documentation locally you can simply run tox and open the html via your favourite web browser.

```
tox -edocs
firefox .tox/docs/tmp/html/index.html
```

3.6.5 Developer Guide

This doc provides details for developers who want to hack on spectrometer. If you have not done so already please refer to the *Quick Start Guide*.

- *Style Guide*
- *Spectrometer Server*
  - *Installing in Dev Mode*
  - *Testing Code*

- *Spectrometer Web*
  - *Installation*
  - *Run spectrometer-web*
  - *UI Technology Stack*
  - *Run spectrometer-web in Production*
  - *Run Test*
  - *Roadmap*
- *Troubleshooting*
  - *Adding new repository*

Style Guide

We follow the Python PEP8 style guide. See: https://www.python.org/dev/peps/pep-0008/

For documentation we follow the Python Documentation Guide. See: https://docs.python.org/devguide/documenting.html

Spectrometer Server

Installing in Dev Mode

In development we want to install spectrometer so that we can modify the code and use it as if in production with changes taking effect immediately. We can achieve this using pip's editable install mode.

```
cd server  # From spectrometer repo root
pip install -e .
spectrometer server -c example-config/config.py start
```

Testing Code

We use tox to manage and run our unit tests. Simply run **tox** in the server directory to initiate the tests. If you don't have tox installed typically it is packaged as **python-tox** in most distros.

```
cd server/  # From spectrometer repo root
tox
```

Spectrometer Web

Installation

To install NodeJS in your system, use the Node Version Manager (NVM), which allows to co-exist multiple NodeJS versions in the same system.

If you already have NodeJS older versions (<= 0.12), it is strongly recommended to completely remove them and reinstall using NVM.

For Linux systems, you can do the following to remove NodeJS:

```
which node # Note down the path
sudo rm -r /path/bin/node /path/bin/npm /path/include/node /path/lib/node_modules ~/.
↪npm
```

Install NVM, NodeJS 4.3.x and NPM:

```
curl -o- https://raw.githubusercontent.com/creationix/nvm/v0.31.0/install.sh | bash
nvm install 4.3.1   # By default this installs npm 2.14.x
npm install npm -g  # This will upgrade npm to 3.7.x
```

Run spectrometer-web

cd web # From the root of the git repo npm install npm start

Goto `http://localhost:8000`

The web project is configured to hot-reload when any changes are made to the code. Most of the time the web browser should auto refresh, if not simply refresh the page.

UI Technology Stack

- NodeJS 4.3 - Bootstrapping and Universal (isomorphic) Javascript execution
- ExpressJS - Web-server-side bootstrap for UI
- ReactJS 0.14 - View Layer
- Redux - Data and State management (Flux pattern)
- Webpack - Build tool
- Babel - Asset compilation, ES6 Transpiler
- FormidableLabs VictoryChart - D3-based React components
- Redux Dev Tools - Tool that allows to track state management

Run spectrometer-web in Production

Production build does not have Devtools and hot reloading middleware. It also minifies scripts and css.

For Production build, execute the following commands:

```
npm run build
npm run start-prod
```

Run Test

Unit Tests are executed using Mocha and Chai assert libraries.

```
npm test
```

Roadmap

1. Dynamic loading of repositories as opposed to loading via config.json

Troubleshooting

Adding new repository

In order to add a new repository to collect statistics, you must make the following changes:

1. Create a soft link in ~/odl-spectrometer to the new repository
2. Edit the server/spectrometer/etc/repositories.yaml and specify the key and path to ~/odl-spectrometer/$repo
3. Edit the web/src/config.json add the project name in the list (this makes it appear in the dropdown)
4. Reload the web page
5. If reload web page does not work, restart python `` `python spectrometer-server` `` and web `` `npm start` ``)

3.6.6 Rest API

Gerrit API

Git API

3.7 Java API Documentation

3.8 NetVirt Contributor Guide