
DEEP-Hybrid-DataCloud releases Documentation

DEEP-Hybrid-DataCloud consortium

Jan 25, 2019

Contents

1	DEEP-Hybrid-DataCloud releases	3
1.1	DEEP-1 Genesis	3
2	Release repositories	23
2.1	Source Code repositories	23
2.2	Artefacts repositories	23
3	Release schedule	25
3.1	Support Model	25
4	Supported platforms	27
4.1	Supported artifacts & packaging formats	27

You'll find here useful information regarding the DEEP-HybridDataCloud services and components **releases**, their schedules, documentation and support.

DEEP-Hybrid-DataCloud releases

1.1 DEEP-1 Genesis

The [DEEP - HybridDataCloud](#) project is pleased to announce the general availability of its **first public software release**, codenamed **Genesis**

1.1.1 Included components

Alien4Cloud - DEEP plugin

Table of Contents

- *Alien4Cloud - DEEP plugin*
 - *Release Notes*
 - *Documentation*
 - *Support*

The **Alien4Cloud - DEEP plugin** is the first step towards the adoption of a visual composition approach of TOSCA templates (based on Alien4Cloud), in order to facilitate the widespread adoption of the standard, and its integration with the INDIGO-DataCloud Orchestrator, which is already part of the EOSC-HUB service catalogue.

- With this approach, the user can visually compose complex applications that involve, for example, the dynamic deployment of a container orchestration platform on an IaaS Cloud site that executes a highly-available Docker-based application to facilitate application delivery.
- The users can also deploy an Apache Mesos cluster with GPU support that contains a deep learning application for the recognition of certain plant species, offered as a service to a community of users. This introduces unprecedented flexibility, from visual composition, to the automated application delivery, using a graphical

interface that is already integrated with an Orchestrator layer that performs resource provision from multiple Clouds and application configuration.

- The integration of easy-to-use graphical interfaces builds a bridge between the users and the orchestration services. It also represents a step forward to foster the adoption of innovative computing services that are hidden from the user, as they can focus on the high-level description of the services requirements and definition, instead of working on their technical implementation.

Release Notes

v 1.1-r0

What's new

This is the first release of the Alien4Cloud - DEEP plugin providing:

- Support for visual composition of TOSCA templates
- Support for the PaaS Orchestrator

List of RfCs

Most important features and improvements:

- [DPD-179] - Support for users in Alien4Cloud
- [DPD-180] - Support for user on the fly
- [DPD-182] - Indigo Tosca Error simple compute node
- [DPD-183] - Repository in the indigo tosca types not handled by A4C
- [DPD-184] - Add a README
- [DPD-185] - Description of IndigoDC data types
- [DPD-186] - Get status for deployed infrastructure
- [DPD-187] - Rm non-required props with null
- [DPD-188] - Generate and write Javadoc
- [DPD-189] - Expose the imports section indigo types
- [DPD-192] - Missing Properties in the Configuration of the Orchestrator Plugin
- [DPD-197] - All TOSCA methods should be strings
- [DPD-198] - Adhere to Google Java Style
- [DPD-270] - /usr/bin/env: bad interpreter: Text file busy
- [DPD-280] - Missing users after restart
- [DPD-281] - Dockerfile custom config
- [DPD-282] - Add support for orchestrator instance confi
- [DPD-283] - De-activate logging of sensitive data
- [DPD-285] - Test coverage $\geq 70\%$
- [DPD-288] - Error loading certificate

- [DPD-294] - README - Obtain Certificate
- [DPD-309] - Change envs name for certificates name holders
- [DPD-311] - Change IDs used to identify deployment
- [DPD-335] - [README] Add how to use REST API
- [DPD-336] - [README] Add documentation versioning
- [DPD-341] - A4C-DEEP security dependency issue
- [DPD-343] - A4C-DEEP documentation
- [DPD-359] - Release notes for Alien4Cloud-deep v. 1.1-r0
- [DPD-365] - Alien4Cloud-deep v. 1.1-r0 - test reports

Complete list of solved issues is available at: [Github](#)

Known Issues

Please see the list of issues on [Github](#)

Dependencies

The Alien4Cloud-DEEP plugin v. 1.1-r0 is known to work with the following:

- Alien4Cloud custom version: 2.0.0-UPV-1.0.0
- Plugin version: 1.0.4-SNAPSHOT
- Alien4Cloud Settings Manager version: 1.0.1
- TOSCA custom-types version: 1.0

List of Artifacts

- **Docker Container:**
 - [indigodatacloud/alien4cloud-deep:DEEP-1](#) (signed)
 - [indigodatacloud/alien4cloud-deep:1.1-r0](#)

Documentation

Detailed documentation is available at:

- [GitHub README](#)

Support

- GitHub issues: <https://github.com/indigo-dc/alien4cloud-deep/issues>
- DEEP project's internal support ticketing system: <http://jira.deep-hybrid-datacloud.eu/>

CloudInfoProvider-DEEP

Table of Contents

- *CloudInfoProvider-DEEP*
 - *Release Notes*
 - *Documentation*
 - *Support*

The **Cloud Information Provider** generates a representation of cloud resources, that can be published inside a BDII (using the provided LDIF templates for a GlueSchema v2 representation) or any other component like the INDIGO Configuration Management Database (CMDB) (Using specific templates).

The generated representation is described using a [Mako](#) template having access to the cloud middleware information.

The current document describe the main features implemented for the Cloud Information Provider under the DEEP-Hybrid-Datacloud project.

Main documentation for the product is maintained [upstream](#), including usage and contribution (for developers) guidelines. The documentation for the features pending to be merged upstream and/or specific to DEEP-Hybrid-Datacloud will be added below, associated with the release version.

Release Notes

v 0.10.4

What's new

This versions adds InfiniBand support for OpenStack in LDIF format:

- cloud-info-provider shall provides now information, through GLUE2ResourceID, when a given OpenStack flavor provides InfiniBand.
- The convention is to use the extra_specs parameter from the flavor's metadata. This parameter is a list of key:value pairs, where the both key's and value's names are not fixed. The provider should allow to define the key's name and the value string to match via command-line options. The InfiniBand support will be published through the GLUE2ExecutionEnvironmentNetworkInfo attribute.

List of RfCs

- DPD-284 - Add InfiniBand support for OpenStack in LDIF format

Known Issues

- None

List of Artifacts

- CentOS-7 RPMS

- cloud-info-provider-deep-0.10.4-1.el7.noarch.rpm
- cloud-info-provider-deep-openstack-0.10.4-1.el7.noarch.rpm
- **Ubuntu 16.04 DEBS**
 - python-cloud-info-provider-deep-openstack_0.10.4-1_all.deb
 - python-cloud-info-provider-deep_0.10.4-1_all.deb

Documentation

Detailed documentation is available at:

- <https://github.com/EGI-Foundation/cloud-info-provider>

Documentation specific for the DEEP-HDC version is available at:

- <https://github.com/indigo-dc/cloud-info-provider-deep/blob/DEEP/README-DEEP.md>

Support

- GitHub issues: <https://github.com/EGI-Foundation/cloud-info-provider/issues?q=is%3Aissue+is%3Aclosed>
- DEEP project's internal support ticketing system: <http://jira.deep-hybrid-datacloud.eu/>
- EGI GGUS Support Unit: https://wiki.egi.eu/wiki/GGUS:EGI_Cloud_Information_Discovery_FAQ

DEEPaaS API

Table of Contents

- *DEEPaaS API*
 - *Release Notes*
 - *Documentation*
 - *Support*

DEEP as a Service (DEEPaaS) is a REST API that is focused on providing access to machine learning models. By using DEEPaaS users can easily run a REST API in from of their model, thus accessing its functionality via HTTP calls.

Release Notes

v 0.1.0

What's new

This is the initial release of DEEPaaS API (deepaas) exposing basic model functionality.

List of RfCs

- Task implemented/partially implemented - Expose model functionality to end users - Expose functionality via API - Provide access to application - Provide resources to efficiently train a model - Facilitate the deployment of a trained model for inference - Provide a easy way to access a deployed application by end users - Provide a way to submit input data to model

Known Issues

- None

List of Artifacts

- **Docker Container:**
 - indigodatacloud/deepaas:DEEP-1 (signed)
 - indigodatacloud/deepaas:0.1.0

v 0.2.0

What's new

This new version allows the user to specify which arguments the underlying model accepts. In this release users can implement a “get_train_args” function in the module entry point, allowing to expose those arguments through the API.

List of RfCs

- Ability to specify parameters for “train”

Known Issues

- None

List of Artifacts

- **Docker Container:**
 - indigodatacloud/deepaas:DEEP-1 (signed)
 - indigodatacloud/deepaas:0.2.0

Documentation

Detailed documentation can be found at:

- [DEEPaaS API ReadTheDocs](#)

Support

- GitHub issues: <https://github.com/indigo-dc/DEEPaaS/issues>
- DEEP project's internal support ticketing system: <http://jira.deep-hybrid-datacloud.eu/>

Infrastructure Manager

Table of Contents

- *Infrastructure Manager*
 - *Release Notes*
 - *Documentation*
 - *Support*

IM is a tool that deploys complex and customized virtual infrastructures on IaaS Cloud deployments (such as AWS, OpenStack, etc.). It eases the access and the usability of IaaS clouds by automating the VMI (Virtual Machine Image) selection, deployment, configuration, software installation, monitoring and update of the virtual infrastructure. It supports APIs from a large number of virtual platforms, making user applications cloud-agnostic. In addition it integrates a contextualization system to enable the installation and configuration of all the user required applications providing the user with a fully functional infrastructure.

This version evolved in the projects **INDIGO-Datacloud** and **DEEP – Hybrid DataCloud**. It is used by the **INDIGO Orchestrator** to contact Cloud sites to finally deploy the VMs/containers.

Release Notes

v1.7.5

What's new

This new version adds a number of new features and improvements, like the support for hybrid deployments, as well as some bug fixes, detailed bellow

List of RfCs

New features and improvements

- [DPD-167] - Enable hybrid deployment
- [DPD-172] - Add option to disable XML-RPC API
- [DPD-173] - Create a connector to Apache CloudStack
- [DPD-174] - Add volume_id to TOSCA BlockStorage type
- [DPD-175] - Add support to disk.*.image_url in OpenStack connector
- [DPD-240] - Facilitate the deployment of a trained model for inference
- [DPD-263] - Enable to specify spot to EC2 instances in TOSCA

- [DPD-264] - Get EC2 instances info from ec2instances.info

Bugs fixed

- [DPD-168] - Problem installing specific branch of Ansible galaxy roles
- [DPD-169] - OCCI server 2.0.4 authentication error
- [DPD-170] - OCCI infrastructure: Volume deleting returns forbidden and takes too long
- [DPD-171] - Fix error reporting memory units in OCCI connector
- [DPD-265] - Error getting the subnets of the default VPC on EC2 conn
- [DPD-271] - Error detecting ansible process end

Known Issues

- None

List of Artifacts

- **CentOS-7 RPMS**
 - IM-1.7.5-1.el7.noarch.rpm
- **Ubuntu 16.04 DEBS**
 - python-im_1.7.5-1_all.deb
- **Docker Container:**
 - indigodatacloud/im:DEEP-1 (signed)
 - indigodatacloud/im:1.7.5

Documentation

Detailed documentation can be found at:

- [GitHub README](#)
- [IM ReadTheDocs](#)
- [Ansible Role for IM](#)

Support

- GitHub issues: <https://github.com/indigo-dc/im/issues>
- DEEP project's internal support ticketing system: <http://jira.deep-hybrid-datacloud.eu/>

Orchent

Table of Contents

- *Orchent*
 - *Release Notes*
 - *Documentation*
 - *Support*

Orchent is a command line application to manage deployments and their resources through the PaaS Orchestrator in a fast and easy way.

Release Notes

v1.2.2

What's new

The new version provides a number of new features, improvements and bug fixes

List of RfCs

- Bug fixes:
 - Fixed communication with oidc-agent \geq 2.0.0
- New features:
 - Added support for two new flags in deprecate/depupdate operations:
 - * `maxProvidersRetry`: Maximum number of cloud providers to be used in case of failure (Default: UNBOUNDED).
 - * `keepLastAttempt`: In case of failure, keep the resources allocated in the last try (Default: true).
- These flags are supported starting from Orchestrator version 2.1.0

Known Issues

- None

List of Artifacts

- **CentOS-7 RPMS**
 - `orchent-1.2.2.el7.centos.x86_64.rpm`
- **Ubuntu 16.04 DEBS**
 - `orchent-1.2.2-amd64.deb`
- **Docker Container:**
 - `indigodatacloud/orchent:DEEP-1 (signed)`
 - `indigodatacloud/orchent:1.2.2`

Documentation

Detailed documentation is available at:

- [GitHub README](#)
- [GitBook Guides](#)

Support

- GitHub issues: <https://github.com/indigo-dc/orchent/issues>
- DEEP project's internal support ticketing system: <http://jira.deep-hybrid-datacloud.eu/>

PaaS Orchestrator

Table of Contents

- *PaaS Orchestrator*
 - *Release Notes*
 - *Documentation*
 - *Support*

The **INDIGO PaaS Orchestrator** is a component of the PaaS layer that allows to instantiate resources on Cloud Management Frameworks (like [OpenStack](#) and [OpenNebula](#)) and [Mesos](#) clusters.

It takes the deployment requests, expressed through templates written in [TOSCA YAML Simple Profile v1.0](#), and deploys them on the best cloud site available. In order to do that

- it gathers SLAs, monitoring info and other data from other platform services,
- it asks to the cloud provider ranker for a list of the best cloud sites.

Release Notes

v2.1.1-FINAL

What's new

This version provides a number of new features, improvements and bug fixes, like:

- Hybrid deployments on multiple sites
- Support to specifying specialized computing hardware
- Improved support for deployment failures

List of RfCs

- New features:

- Implement Cloud Providers Retry Logic (#288 DPD-121)
- Retrieve of information from CMDB (fixes #290 DPD-124)
- AuthN through OIDC (fixes #291 DPD-125)
- Add support for mesos tasks with GPUs (fixes #292 DPD-126)
- Support IM async infrastructure creation #289
- Allow to import custom CAs in the java truststore #238
- improvements
 - Improve OneData Spaces definition through templates (fixes #293 XSD-28)
- Bug fixed:
 - Return 409 status code for SQLTransientExceptions #285
 - (MySQL) Timestamp columns have implicit default value assigned #283

Known Issues

- None

Dependencies

The PaaS Orchestrator v.2.1.1-FINAL has the following services dependencies

- CMDB release v0.4 (docker image indigodatacloud/cmdb:indigo_2)
- SLAM release v2.0.0 (docker image indigodatacloud/slam:v2.0.0)
- CPR release v0.6.0 (docker image indigodatacloud/cloudproviderranker:indigo_2)
- Monitoring - Zabbix Wrapper 1.0.2 (docker image indigodatacloud/zabbix-wrapper:indigo_2)
- IM release >= 1.7.4 * (Optional) Onedata v18.0.2-rc[11,12]
- toasca-types v3.0.0

List of Artifacts

- **Docker Container:**
 - [indigodatacloud/orchestrator:DEEP-1](#) (signed)
 - [indigodatacloud/orchestrator:1.7.5](#)

Documentation

Detailed documentation is available at:

- [GitHub README](#)
- [GitBook Guides](#)

Support

- GitHub issues: <https://github.com/indigo-dc/orchestrator/issues>
- DEEP project's internal support ticketing system: <http://jira.deep-hybrid-datacloud.eu/>

TOSCA types & templates

Table of Contents

- *TOSCA types & templates*
 - *Release Notes*
 - *Documentation*
 - *Support*

The **TOSCA types** repository shows a YAML description of new types added first in the INDIGO project, and afterwards in the DEEP-HybridDataCloud and eXtreme DataCloud projects, to extend TOSCA Simple Profile in YAML Version 1.0 to add high level entities. In the examples directory there are a set of TOSCA documents using these types that will be supported by the INDIGO, DEEP and XDC components.

The **TOSCA Templates** repository contain templates Supporting the Use Cases for INDIGO-DataCloud, DEEP-HybridDataCloud and eXtreme DataCloud projects.

Release Notes

v.3.0.0

What's new

This version provides a number of new features, improvements and bug fixes

List of RFCs

- Updated example template for deploying a Mesos cluster
- Added example template for deploying a Mesos cluster with GPU support and a tensorflow container on top of it that uses GPU(s)
- Added GPU support for compute nodes and dockerized jobs (chronos) and apps (marathon)
- Added preemptible_instance property to support “spot” instances
- Added example template for launching the generic_deepaas mesos job
- Added new types for describing a Kubernetes cluster; new Ansible roles implemented.
- Added new type for describing a JupyterHub node; new Ansible role implemented.
- Fixing for Galaxy on cloud:
 - Update support to Galaxy release_18.05 in indigo-dc.galaxycloud
 - Fix proftpd in indigo-dc.galaxycloud

- New ansible role for galaxy tools installation, named indigo-dc.galaxycloud-tools
- Update tosca.nodes.indigo.GalaxyShedTool with new ansible role
- Fix CERN-VM FS reference data mount on cluster worker nodes on galaxy artifacts
- Reworked ansible role (indigo-dc.galaxycloud-fastconfig) to reconfigure an image with Galaxy already installed
- Reworked storage encryption script on indigo-dc.galaxycloud-os
- Fix tasks order in Galaxy elastic cluster tosca template: now galaxy user is created before slurm configuration.

Known Issues

- None

List of Artifacts

- **Tarballs:**
 - tosca-types-3.0.0.tar.gz
 - tosca-templates-3.0.0.tar.gz

Documentation

Detailed documentation can be found at:

- [TOSCA types README](#)
- [TOSCA templates README](#)
- [Brief guide to TOSCA](#)
- [How to deploy a TOSCA Template](#)

Support

- **GitHub issues:**
 - <https://github.com/indigo-dc/tosca-types/issues>
 - <https://github.com/indigo-dc/tosca-templates/issues>
- DEEP project's internal support ticketing system: <http://jira.deep-hybrid-datacloud.eu/>

udocker

Table of Contents

- *udocker*
 - *Release Notes*

- *Documentation*
- *Support*

udocker is a basic user tool to execute simple docker containers in user space without requiring root privileges. Enables download and execution of docker containers by non-privileged users in Linux systems where docker is not available. It can be used to pull and execute docker containers in Linux batch systems and interactive clusters that are managed by other entities such as grid infrastructures or externally managed batch or interactive systems.

udocker does not require any type of privileges nor the deployment of services by system administrators. It can be downloaded and executed entirely by the end user.

udocker is a wrapper around several tools to mimic a subset of the docker capabilities including pulling images and running containers with minimal functionality.

Release Notes

v1.1.3

What's new

This version provides a number of new features, improvements and bug fixes

List of RfCs

Tasks

- [DPD-296] - Add support for nvidia-docker like functionality
- [DPD-301] - Add support for runC to v1.0.0-rc5
- [DPD-303] - Command line interface improvements
- [DPD-304] - Improve metadata handling
- [DPD-306] - Add kill-on-exit to proot on Pn modes
- [DPD-307] - Force python2 as interpreter
- [DPD-310] - Directory absolute symlinks in Fn modes not being converted

Bugs

- [DPD-297] - Improve parsing of quotes in the command line
- [DPD-298] - Issues in file downloads
- [DPD-299] - Mitigation for upstream limitation in runC without tty
- [DPD-300] - Fix detection of executable with symlinks in container
- [DPD-305] - Improve pathname translation in Fn modes for mounted dirs

Known Issues

- None

List of Artifacts

Tarball:

- `udocker-1.1.3.tar.gz`

Documentation

Detailed documentation can be found at:

- [GitBook Guide](#)
- [GitHub documentation](#)

Support

- **GitHub issues:**
 - <https://github.com/indigo-dc/udocker/issues>
- DEEP project's internal support ticketing system: <http://jira.deep-hybrid-datacloud.eu/>

INDIGO Virtual Router Configurator

Table of Contents

- *INDIGO Virtual Router Configurator*
 - *Release Notes*
 - *Documentation*
 - *Support*

The **INDIGO-DC Virtual Router** does not carry its own code. It consists of open components and all automation is contained in Ansible roles.

The purpose of the INDIGO-DC Virtual Router is establishing an overlay virtual private network to interconnect nodes in a cloud platform deployment even if deployed in multiple, geographically distant sites.

Release Notes

DEEP-v1

What's new

The Ansible roles are intended for automated deployment of the INDIGO-DC Virtual Router through the Infrastructure Manager [2]. They can be deployed manually with Ansible if required.

List of RfCs

- N/A

Known Issues

- None

List of Artifacts

- **Ansible Roles**
 - <https://github.com/indigo-dc/ansible-role-indigovr>

Documentation

Detailed documentation is available at:

- [GitHub README](#)

Support

- [GitHub issues](https://github.com/indigo-dc/ansible-role-indigovr/issues): <https://github.com/indigo-dc/ansible-role-indigovr/issues>
- DEEP project's internal support ticketing system: <http://jira.deep-hybrid-datacloud.eu/>

1.1.2 Highlights

Key technical highlights:

- Alien4Cloud-DEEP plugin
 - Support for visual composition of TOSCA templates
 - PaaS orchestrator support
- cloud-info-provider
 - Support for GPU and Infiniband resources
- DEEPaaS API
 - Support for training a machine learning application
 - Support for performing inferences/analysis/predictions
 - Support only for synchronous requests
 - OpenID Connect support
 - Support for standalone service & OpenWhisk action
- Infrastructure Manager (IM)
 - Improved support for hybrid deployments
 - Support for additional TOSCA types
- PaaS Orchestrator
 - Hybrid deployments on multiple sites
 - Support to specifying specialized computing hardware Improved support for deployment failures

- uDocker
 - Improved support for GPUs and Infiniband
- Visual application topology composition and deployment
 - Graphical composition of complex application topologies
 - Deployment through PaaS orchestrator
- DEEP as a Service
 - Deployment of DEEP Open Catalog components as server-less functions
- DEEP Open Catalog
 - Ready-to-use machine learning and deep learning applications, including:
 - * Machine learning frameworks + JupyterLab
 - * Machine learning ready to use models
 - * Deep learning ready to use models
 - * BigData analytic tools

1.1.3 Release Notes

The DEEP-1/Genesis release consists in 9 Products and a number of technical guides:

- 17 OS packages, for CentOS 7 and Ubuntu 16.04
 - 8 RPMS & SRPMS
 - 6 binary & source DEBS
 - 3 tarballs
- 5 Docker containers

You can find in the later sections the full list of products, with detailed release notes and instructions for their installation/configuration.

1.1.4 Installation Notes

All DEEP - HybridDataCloud products are distributed from standard Operating Systems (OS) repositories and DockerHub registry of the [indigodatacloud](#) organization.

The packages repositories have the following structure:

- DEEP-HDC **production** (stable):
 - `deep-hdc/production/{1,2}/centos7/x86_64/{baselupdates}`
 - `deep-hdc/production/{1,2}/ubuntu/dists/xenial/main/{binary-amd64,source}`
 - * containing signed, well tested software components
 - third-party:
 - * `deep-hdc/production/{1,2}/centos7/x86_64/third-party`
 - * `deep-hdc/production/{1,2}/ubuntu/dists/xenial/third-party{binary-amd64,source}`
 - containing packages that are not part of DEEP, or not part of the base OS or EPEL, but used as dependencies by other DEEP components

- DEEP-HDC **testing**: `deep/testing/{1,2}/{centos7,ubuntu}/`
 - containing packages that will become part of the next stable distribution; in the certification and validation phase.
- DEEP-HDC **preview**: `deep/preview/{1,2}/{centos7,ubuntu}/`
 - containing signed packages that will become part of the next stable update, available for technical-previews

All packages in production and preview repositories are signed with the INDIGO - DataCloud gpg key. The public key can be downloaded from [here](#), and the fingerprint from [here](#).

Please import the key BEFORE starting!

- for CentOS7 save the key under `/etc/pki/rpm-gpg/`

```
# rpm --import http://repo.indigo-datacloud.eu/repository/RPM-GPG-KEY-indigodc
```

- for Ubuntu:

```
# wget -q -O - http://repo.indigo-datacloud.eu/repository/RPM-GPG-KEY-indigodc | sudo_↵  
↵ apt-key add -
```

Configuring the use of DEEP-HDC repositories

YUM & APT configuration files are available at:

- CentOS7 - <http://repo.indigo-datacloud.eu/deep-hdc/repos/deep-1.repo>
- Ubuntu 16.04 - http://repo.indigo-datacloud.eu/deep-hdc/repos/deep-1-ubuntu16_04.list

or use the `deephdc-release` package to install DEEP-HDC repositories:

- CentOS7:

```
# wget http://repo.indigo-datacloud.eu/repository/deep-hdc/production/1/centos7/x86_↵  
↵ 64/base/deephdc-release-1.0.0-1.el7.noarch.rpm  
# yum localinstall -y deephdc-release-1.0.0-1.el7.noarch.rpm
```

- Ubuntu 16.04:

```
# wget http://repo.indigo-datacloud.eu/repository/deep-hdc/1/ubuntu/dists/xenial/main/↵  
↵ binary-amd64/deephdc-release_1.0.0-1_amd64.deb  
# dpkg -i deephdc-release_1.0.0-1_amd64.deb
```

These packages will install required dependencies, the INDIGO - DataCloud public key and ensures the precedence of DEEP-HybridDataCloud repositories over EPEL and Ubuntu.

It is strongly recommended the use of the latest version of the **deephdc-release** packages containing the public key and the YUM and APT repository files.

Enable the DEEP-HDC Containers repositories

On the [DockerHub Registry](#), DEEP - HybridDataCloud uses the INDIGO - DataCloud and DEEP-HDC Organizations:

- `indigodatacloud`, for Core Services
- `deephdc`, for DEEP-OC modules

Containers present in those repositories and released in DEEP-1 major release are tagged with “DEEP-1” tag and signed, leveraging the [Docker’s trust features](#) so that users can pull trusted images.

Currently, content trust is disabled by default. You must enable it by setting the `DOCKER_CONTENT_TRUST` environment variable, like below:

```
export DOCKER_CONTENT_TRUST=1
```

For more details regarding the “Content Trust in Docker” please read [Docker’s Documentation](#)

Content trust is associated with the TAG portion of an image. For DEEP-1 (Genesis) release the signed tag is DEEP-1. See example bellow if you want to ensure the correct use of DEEP-HDC images:

```
$ export DOCKER_CONTENT_TRUST=1

$ docker pull indigodatacloud/orchestrator:2.1.1-FINAL
No trust data for 2.1.1-FINAL

$ docker pull indigodatacloud/orchestrator:DEEP-1
Pull (1 of 1): indigodatacloud/orchestrator:DEEP-
↳1@sha256:441c8b037684422ccdf2fdec322c9c09904ed3ce74d9fcc7d2862a9f53ad36be
sha256:441c8b037684422ccdf2fdec322c9c09904ed3ce74d9fcc7d2862a9f53ad36be: Pulling from
↳indigodatacloud/orchestrator
93857f76ae30: Pull complete
[...]
e8c92b40b492: Pull complete
Digest: sha256:441c8b037684422ccdf2fdec322c9c09904ed3ce74d9fcc7d2862a9f53ad36be
Status: Downloaded newer image for indigodatacloud/
↳orchestrator@sha256:441c8b037684422ccdf2fdec322c9c09904ed3ce74d9fcc7d2862a9f53ad36be
Tagging indigodatacloud/
↳orchestrator@sha256:441c8b037684422ccdf2fdec322c9c09904ed3ce74d9fcc7d2862a9f53ad36be
↳as indigodatacloud/orchestrator:indigo_2
$ docker images
REPOSITORY          TAG          IMAGE ID          CREATED
↳          SIZE
indigodatacloud/orchestrator  DEEP-1      bdbe758d9f32     37 hours ago
↳          843MB
```

1.1.5 Documentation

Please find DEEP-1 services and components documentation at [Included components](#).

1.1.6 Support

Most complex software contains bugs, we are not an exception. One of the features of free and open source software is the ability to report bugs, helping to fix or improve the software you use.

DEEP - HybridDataCloud project uses the [INDIGO Catch-All GGUS - Support Unit](#) and the deep-support@listas.csic.eu for general support requests. More details regarding each product support channels are provided in the respective products release pages.

Developers, researchers and IT enthusiasts: feel free to write to deep-info@listas.csic.es to ask for more information on how to use DEEP solutions for your work. For automatic notifications you can register to the DEEP - HybridDataCloud release RSS feed or subscribe to the DEEP - HybridDataCloud Announce Mailing list. You can also socialize with us via [Twitter](#) or by joining our [LinkedIn group](#).

2.1 Source Code repositories

Source code repositories are available on GitHub under the **indigo-dc** and **deephdc** organizations - <http://bit.ly/deep-hdc>

2.2 Artefacts repositories

DEEP-HDC **production** (stable) repositories:

- `deep-hdc/production/{1,2}/centos7/x86_64/{baseupdates}`
- `deep-hdc/production/{1,2}/ubuntu/dists/xenial/main/{binary-amd64,source}`
 - containing signed, well tested software components
- **third-party**:
 - `deep-hdc/production/{1,2}/centos7/x86_64/third-party`
 - `deep-hdc/production/{1,2}/ubuntu/dists/xenial/third-party{binary-amd64,source}`
 - * containing packages that are not part of DEEP-HDC project, or not part of the base OS or EPEL, but used as dependencies by other DEEP-HDC components

YUM & APT **configuration files** are available [here](#) or use the **deephdc-release** package to install DEEP-HDC repositories

Release schedule

- **Time-based releases**
 - **projects' Major releases** - the DEEP-HybridDataCloud project foresees two major releases, distributions, during its lifetime, at around 10 months since the start of the project.
- **As-soon-as-available**
 - **components' Minor/Revision releases** - in a project Major release, Development teams (aka Product Teams) can release updated versions of their components as soon as the DEEP-HDC software quality criteria are met. Thought the project Continuous Integration and Delivery System tests are continuously run giving feedback on the status of the components.

3.1 Support Model

- in a Major Release for each component or service only the latest revision released is supported.
- for each component or service a (major) release is supported at least for the lifetime of the projects' major release in which this version was released the first time.

Supported platforms

- DEEP-HybridDataCloud releases are supported for the following platforms:
 - **CentOS7 & Ubuntu 16.04**
 - * for the products distributed through rpms and deb packages
 - **all platforms supporting Docker containers**
 - * for the products distributed as docker images

4.1 Supported artifacts & packaging formats

- Packages:
 - Binaries: executable packages
 - Sources: when available, package files that contain all of the necessary files to compile/build the respective piece of software
 - Tarballs: clients are distributed as tarballs for all the platforms
- Containers: Docker images are available for almost all project software