
Graylog Documentation

Release 1.0.0

Graylog, Inc.

Aug 04, 2017

1	Architectural considerations	3
1.1	Minimum setup	3
1.2	Bigger production setup	4
1.3	Highly available setup with Graylog Radio	5
2	Installing Graylog	7
2.1	Virtual machine appliances	7
2.2	The manual setup	8
2.3	Operating system packages	15
2.4	Chef, Puppet, Ansible, Vagrant	17
2.5	Amazon Web Services	17
2.6	Docker	17
2.7	Microsoft Windows	17
3	Configuring and tuning Elasticsearch	19
3.1	Configuration	19
3.2	Cluster Status explained	21
4	Sending in log data	23
4.1	What are Graylog message inputs?	23
4.2	Content packs	23
4.3	Syslog	23
4.4	GELF / Sending from applications	25
4.5	Microsoft Windows	25
4.6	Heroku	25
4.7	Ruby on Rails	28
4.8	Raw/Plaintext inputs	29
4.9	JSON path from HTTP API input	29
4.10	Reading from files	30
5	The search query language	33
5.1	Syntax	33
5.2	Escaping	35
5.3	Time frame selector	35
5.4	Search result highlighting	35
6	Streams	37

6.1	What are streams?	37
6.2	Alerts	38
6.3	Outputs	41
6.4	Use cases	41
6.5	How are streams processed internally?	42
6.6	Stream Processing Runtime Limits	43
6.7	Programmatic access via the REST API	44
6.8	FAQs	45
7	Dashboards	47
7.1	Why dashboards matter	47
7.2	How to use dashboards	48
7.3	Examples	49
7.4	Widgets from streams	50
7.5	Modifying dashboards	50
7.6	Dashboard permissions	51
8	Extractors	53
8.1	The problem explained	53
8.2	Graylog extractors explained	53
8.3	The extractor directory	54
8.4	Using regular expressions to extract data	55
8.5	Using Grok patterns to extract data	56
8.6	Normalization	57
9	Message rewriting with Drools	61
9.1	Getting Started	61
9.2	Example rules file	61
9.3	Parsing Message and adding fields	62
10	Load balancer integration	65
10.1	Load balancer state	65
10.2	Graceful shutdown	66
11	The Graylog index model explained	67
11.1	Overview	67
11.2	Eviction of indices and messages	69
11.3	Keeping the metadata in synchronisation	69
11.4	Manually cycling the deflector	69
12	Indexer failures and dead letters	71
12.1	Indexer failures	71
12.2	Dead letters	72
12.3	Common indexer failure reasons	73
13	Plugins	75
13.1	General information	75
13.2	Creating a plugin skeleton	75
13.3	Example Alarm Callback plugin	76
13.4	Building plugins	77
13.5	Installing and loading plugins	77
14	External dashboards	79
14.1	CLI stream dashboard	79
14.2	Browser stream dashboard	80

15	The thinking behind the Graylog architecture and why it matters to you	83
15.1	A short history of Graylog	83
15.2	The log management market today	83
15.3	The future	85
16	Changelog	87
16.1	Graylog 1.0.2	87
16.2	Graylog 1.0.1	87
16.3	Graylog 1.0.0	88
16.4	Graylog 1.0.0-rc.4	89
16.5	Graylog 1.0.0-rc.3	89
16.6	Graylog 1.0.0-rc.2	89
16.7	Graylog 1.0.0-rc.1	90
16.8	Graylog 1.0.0-beta.2	91
16.9	Graylog 1.0.0-beta.1	92
16.10	Graylog2 0.92.4	92
16.11	Graylog 1.0.0-beta.1	93
16.12	Graylog2 0.92.3	93
16.13	Graylog2 0.92.1	93
16.14	Graylog2 0.92.0	94
16.15	Graylog2 0.92.0-rc.1	95
16.16	Graylog2 0.91.3	95
16.17	Graylog2 0.91.3	95
16.18	Graylog2 0.92.0-beta.1	96
16.19	Graylog2 0.91.1	96
16.20	Graylog2 0.90.1	96
16.21	Graylog2 0.91.0-rc.1	97
16.22	Graylog2 0.90.0	97
16.23	Graylog2 0.20.3	97
16.24	Graylog2 0.20.2	98

NOTE: There are multiple options for reading this documentation. See link to the lower left.

Contents:

Architectural considerations

There are a few rules of thumb when scaling resources for Graylog:

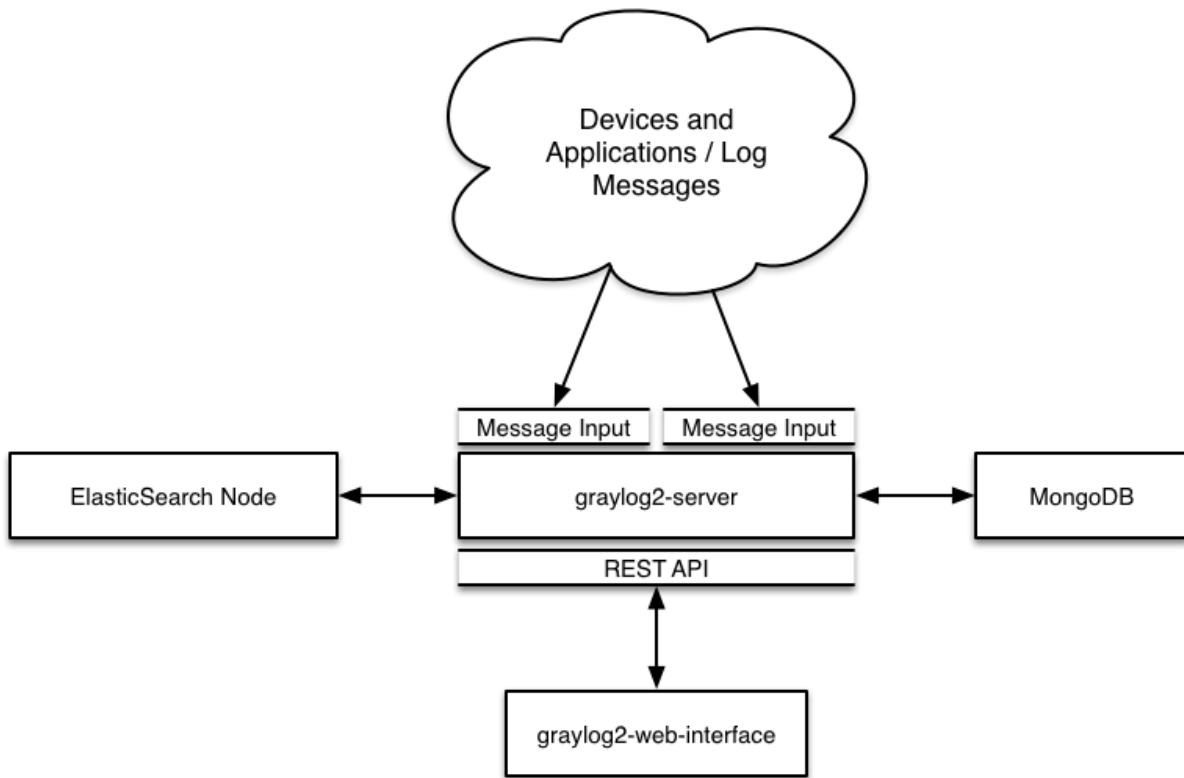
- `graylog-server` nodes should have a focus on CPU power.
- Elasticsearch nodes should have as much RAM as possible and the fastest disks you can get. Everything depends on I/O speed here.
- MongoDB is only being used to store configuration and the dead letter messages, and can be sized fairly small.
- `graylog-web-interface` nodes are mostly waiting for HTTP answers of the rest of the system and can also be rather small.
- `graylog-radio` nodes act as workers. They don't know each other and you can shut them down at any point in time without changing the cluster state at all.

Also keep in mind that messages are **only** stored in Elasticsearch. If you have data loss on Elasticsearch, the messages are gone - except if you have created backups of the indices.

MongoDB is only storing meta information and will be abstracted with a general database layer in future versions. This will allow you to use other databases like MySQL instead.

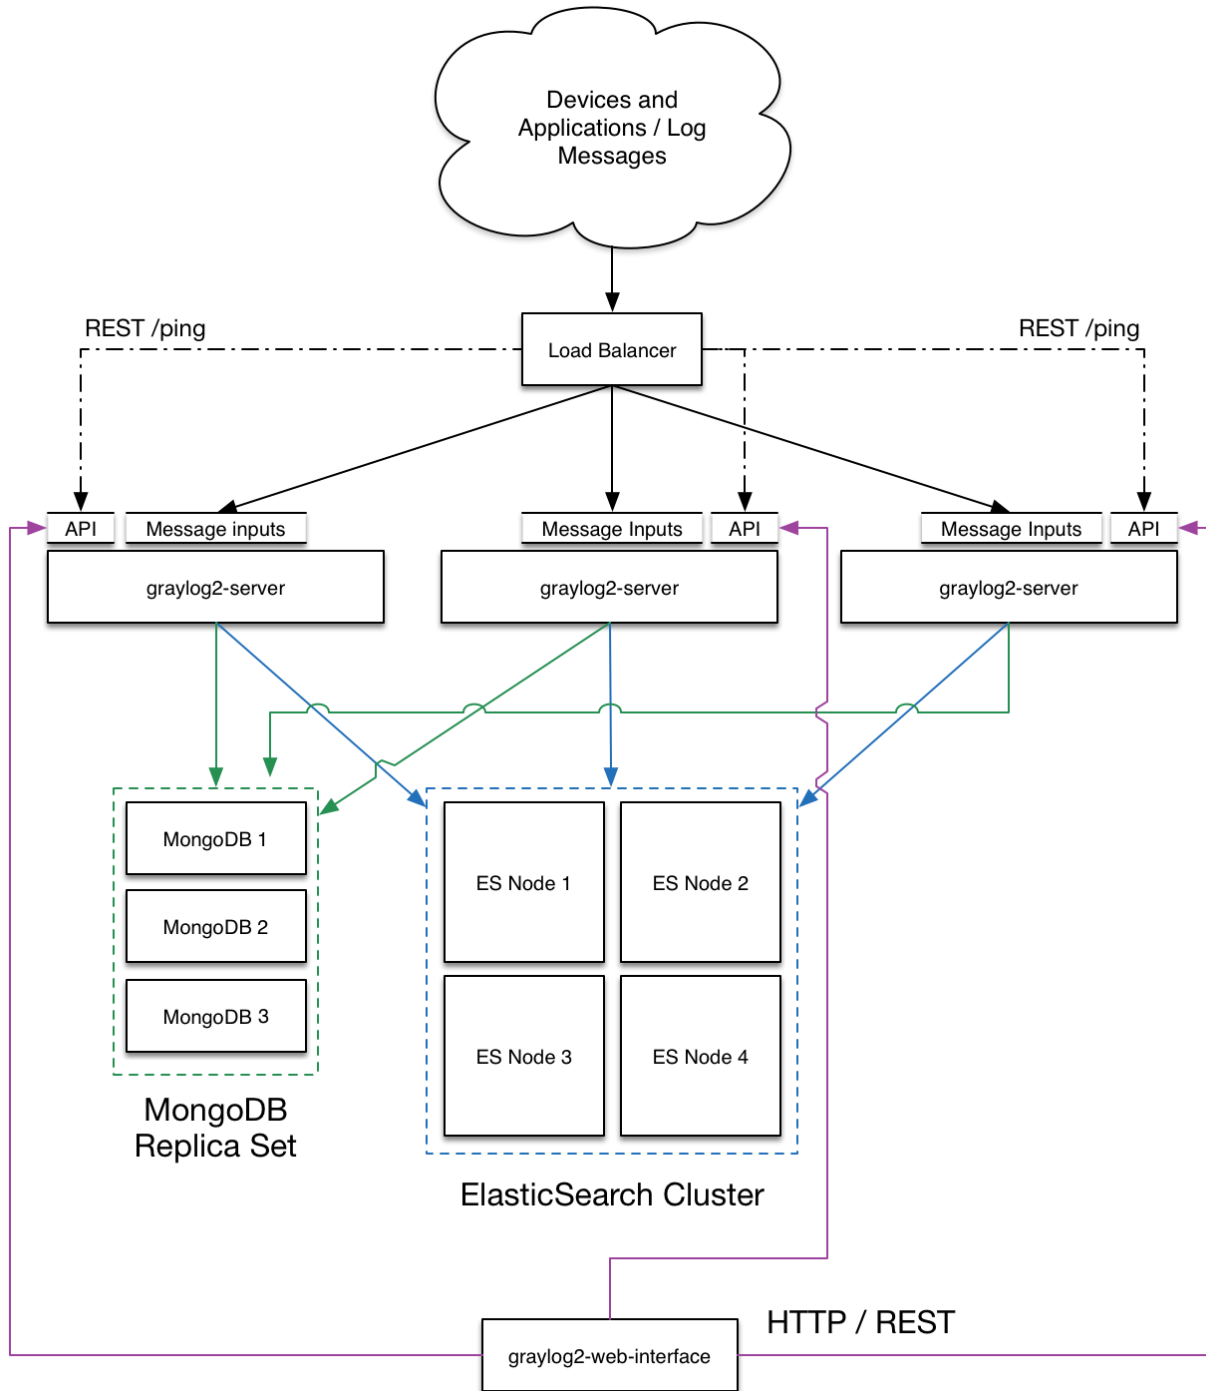
Minimum setup

This is a minimum Graylog setup that can be used for smaller, non-critical, or test setups. None of the components is redundant but it is easy and quick to setup.



Bigger production setup

This is a setup for bigger production environments. It has several `graylog-server` nodes behind a load balancer that share the processing load. The load balancer can ping the `graylog-server` nodes via REST/HTTP to check if they are alive and take dead nodes out of the cluster.



Highly available setup with Graylog Radio

Beginning with Graylog 1.0 on we do no longer recommend running Graylog Radio because we are now using a high-performant message journal (from the Apache Kafka project) in every *graylog-server* instance which is spooling all incoming messages to disk immediately and is able to buffer load spikes just at least as good as Graylog Radio was, but with less dependencies and maintenance overhead.

If you are running a setup with Graylog Radio we recommend to shut down the Graylog Radio architecture including AMQP or Kafka brokers completely and directly send messages to the *graylog-server* nodes. If you have been using Graylog Radio for load balancing, you should now put a classic load balancer in front of your *graylog-server* nodes.

This approach has been proven to work great in large high-throughput setups of several of our large scale customers and immensely reduced complexity of their setups.

The Kafka and AMQP inputs are still supported and can be used to build a custom setup using message brokers, if you want to keep using that. A reason for this might be that Graylog is not the only subscriber to the messages on the bus. However we would recommend to use Graylog forwarders to either write to a message bus after processing or write to other systems directly.

Installing Graylog

Modern server architectures and configurations are managed in many different ways. Some people still put new software somewhere in `opt` manually for each server while others have already jumped on the configuration management train and fully automated reproducible setups.

Graylog can be installed in many different ways so you can pick whatever works best for you. We recommend to start with the *Virtual machine appliances* for the fastest way to get started and then pick one of the other, more flexible installation methods to build a production ready setup. (Note: The *Virtual machine appliances* are suitable for production usage because they are prepared to scale out to multiple servers when required.)

This chapter is explaining the many ways to install Graylog and aims to help choosing the one that fits your needs.

Virtual machine appliances

The easiest way to get started with a production ready Graylog setup is using our official virtual machine appliances. We offer those for the following platforms:

- [OVA for VMware / Virtualbox](#)
- [OpenStack](#)
- [Amazon Web Services / EC2](#)
- [Docker](#)

Please follow the links for specific instructions and downloads. The next chapters explain general principles of the appliances:

Configuring the appliances

The great thing about the new appliances is the `graylog-ctl` tool that we are shipping with them. We want you to get started with a customised setup as soon as quickly as possible so you can now do things like:

```
graylog-ctl set-email-config <smtp server> [--port=<smtp port> --user=<username> --  
↪password=<password>]  
graylog-ctl set-admin-password <password>  
graylog-ctl set-timezone <zone acronym>  
graylog-ctl reconfigure
```

Assign a static IP

Per default the appliance make use of DHCP to setup the network. If you want to access Graylog under a static IP please edit the file `/etc/network/interfaces` like this (just the important lines):

```
auto eth0  
  iface eth0 inet static  
    address <static IP address>  
    netmask <netmask>  
    gateway <default gateway>  
    pre-up sleep 2
```

Activate the new IP and reconfigure Graylog to make use of it:

```
$ sudo ifdown eth0 && sudo ifup eth0  
$ sudo graylog-ctl reconfigure
```

Wait some time till all services are restarted and running again. Afterwards you should be able to access Graylog with the new IP.

Scaling out

We are also providing an easy way to automatically scale out to more boxes once you grew out of your initial setup. Every appliance is always shipping with all required Graylog components and you can at any time decide which role a specific box should take:

```
graylog-ctl reconfigure-as-server  
graylog-ctl reconfigure-as-webinterface  
graylog-ctl reconfigure-as-datanode
```

The manual setup

We recommend to only run this if you have good reasons not to use one of the other production ready installation methods described in this chapter.

Manual setup: graylog-server on Linux

Prerequisites

You will need to have the following services installed on either the host you are running `graylog-server` on or on dedicated machines:

- [Elasticsearch 1.3.7 or later](#) (Elasticsearch 2.x is currently not supported)
- MongoDB (as recent stable version as possible, **at least v2.0**)

Most standard MongoDB packages of Linux distributions are outdated. Use the [official MongoDB APT repository](#) (available for many distributions and operating systems)

You also **must** install **Java 7** or higher! Java 6 is not compatible with Graylog and will also not receive any more publicly available bug and security fixes by Oracle.

A more detailed guide for installing the dependencies will follow. **The only important thing for Elasticsearch is that you set the exactly same cluster name (e. g. “cluster.name: graylog”) that is being used by Graylog in the Elasticsearch configuration (“conf/elasticsearch.yml”).**

Downloading and extracting the server

Download the tar archive from the [download pages](#) and extract it on your system:

```
~$ tar xvfz graylog-VERSION.tgz
~$ cd graylog-VERSION
```

Configuration

Now copy the example configuration file:

```
~# cp graylog.conf.example /etc/graylog/server/server.conf
```

You can leave most variables as they are for a first start. All of them should be well documented.

Configure at least the following variables in `/etc/graylog/server/server.conf`:

- **is_master = true**
 - Set only one `graylog-server` node as the master. This node will perform periodical and maintenance actions that slave nodes won't. Every slave node will accept messages just as the master nodes. Nodes will fall back to slave mode if there already is a master in the cluster.
- **password_secret**
 - You must set a secret that is used for password encryption and salting here. The server will refuse to start if it's not set. Generate a secret with for example `pwgen -N 1 -s 96`. If you run multiple `graylog-server` nodes, make sure you use the same `password_secret` for all of them!
- **root_password_sha2**
 - A SHA2 hash of a password you will use for your initial login. Set this to a SHA2 hash generated with `echo -n yourpassword | shasum -a 256` and you will be able to log in to the web interface with username `admin` and password `yourpassword`.
- **elasticsearch_max_docs_per_index = 20000000**
 - How many log messages to keep per index. This setting multiplied with `elasticsearch_max_number_of_indices` results in the maximum number of messages in your Graylog setup. It is always better to have several more smaller indices than just a few larger ones.
- **elasticsearch_max_number_of_indices = 20**
 - How many indices to have in total. If this number is reached, the oldest index will be deleted. **Also take a look at the other retention strategies that allow you to automatically delete messages based on their age.**
- **elasticsearch_shards = 4**

- The number of shards for your indices. A good setting here highly depends on the number of nodes in your Elasticsearch cluster. If you have one node, set it to 1. Read more about this in the knowledge base article about *Configuring and tuning Elasticsearch*.
- `elasticsearch_replicas = 0`
 - The number of replicas for your indices. A good setting here highly depends on the number of nodes in your Elasticsearch cluster. If you have one node, set it to 0. Read more about this in the knowledge base article about *Configuring and tuning Elasticsearch*.
- `mongodb_*`
 - Enter your MongoDB connection and authentication information here. Make sure that you connect the web interface to the same database. You don't need to configure `mongodb_user` and `mongodb_password` if `mongodb_useauth` is set to `false`.

Starting the server

You need to have Java installed. Running the OpenJDK is totally fine and should be available on all platforms. For example on Debian it is:

```
~$ apt-get install openjdk-7-jre
```

You need at least Java 7 as Java 6 has reached EOL.

Start the server:

```
~$ cd bin/  
~$ ./graylogctl start
```

The server will try to write a `node_id` to the `graylog-server-node-id` file. It won't start if it can't write there because of for example missing permissions.

See the startup parameters description below to learn more about available startup parameters. Note that you might have to be `root` to bind to the popular port 514 for syslog inputs.

You should see a line like this in the debug output of `graylog-server` successfully connected to your Elasticsearch cluster:

```
2013-10-01 12:13:22,382 DEBUG: org.elasticsearch.transport.netty - [graylog-server]_  
↪connected to node [[Unuscione, Angelo][thN_gIBkQDm2ab7k-2Zaaw][inet[/10.37.160.  
↪227:9300]]]
```

You can find the `graylog-server` logs in the directory `logs/`.

Important: All `graylog-server` instances must have synchronised time. We strongly recommend to use `NTP` or similar mechanisms on all machines of your Graylog infrastructure.

Supplying external logging configuration

The `graylog-server` uses Log4j for its internal logging and ships with a *default log configuration file* <<https://github.com/Graylog2/graylog2-server/blob/1.0.0/graylog2-server/src/main/resources/log4j.xml>> which is embedded within the shipped JAR.

In case you need to overwrite the configuration `graylog-server` uses, you can supply a Java system property specifying the path to the configuration file in your `graylogctl` script. Append this before the `-jar` parameter:


```
-Dlog4j.configuration=file:///tmp/logj4.xml
```

Substitute the actual path to the file for the `/tmp/log4j.xml` in the example.

In case you do not have a log rotation system already in place, you can also configure Graylog to rotate logs based on their size to prevent its logs to grow without bounds.

One such example `log4j.xml` configuration is shown below. Graylog includes the `log4j-extras` companion classes to support time based and size based log rotation. This is the example:

```
<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE log4j:configuration PUBLIC "-//APACHE//DTD LOG4J 1.2//EN" "log4j.dtd">
<log4j:configuration xmlns:log4j="http://jakarta.apache.org/log4j/">

  <appender name="FILE" class="org.apache.log4j.rolling.RollingFileAppender">
    <rollingPolicy class="org.apache.log4j.rolling.FixedWindowRollingPolicy" >
      <param name="activeFileName" value="/tmp/server.log" /> <!-- ADAPT -->
      <param name="fileNamePattern" value="/tmp/server.%i.log" /> <!-- ADAPT -->
      <param name="minIndex" value="1" /> <!-- ADAPT -->
      <param name="maxIndex" value="10" /> <!-- ADAPT -->
    </rollingPolicy>
    <triggeringPolicy class="org.apache.log4j.rolling.SizeBasedTriggeringPolicy">
      <param name="maxFileSize" value="5767168" /> <!-- ADAPT: For example 5.
↪5MB in bytes -->
    </triggeringPolicy>
    <layout class="org.apache.log4j.PatternLayout">
      <param name="ConversionPattern" value="%d %-5p: %c - %m%n"/>
    </layout>
  </appender>

  <!-- Application Loggers -->
  <logger name="org.graylog2">
    <level value="info"/>
  </logger>
  <!-- this emits a harmless warning for ActiveDirectory every time which we can't
↪work around :( -->
  <logger name="org.apache.directory.api.ldap.model.message.BindRequestImpl">
    <level value="error"/>
  </logger>
  <!-- Root Logger -->
  <root>
    <priority value="info"/>
    <appender-ref ref="FILE"/>
  </root>
</log4j:configuration>
```

Command line (CLI) parameters

There are a number of CLI parameters you can pass to the call in your `graylogctl` script:

- `-h, --help`: Show help message
- `-f CONFIGFILE, --configfile CONFIGFILE`: Use configuration file *CONFIGFILE* for Graylog; default: `/etc/graylog/server/server.conf`
- `-t, --configtest`: Validate the Graylog configuration and exit with exit code 0 if the configuration file is syntactically correct, exit code 1 and a description of the error otherwise

- `-d, --debug`: Run in debug mode
- `-l, --local`: Run in local mode. Automatically invoked if in debug mode. Will not send system statistics, even if enabled and allowed. Only interesting for development and testing purposes.
- `-s, --statistics`: Print utilization statistics to STDOUT
- `-r, --no-retention`: Do not automatically delete old/outdated indices
- `-p PIDFILE, --pidfile PIDFILE`: Set the file containing the PID of graylog to *PIDFILE*; default: */tmp/graylog.pid*
- `-np, --no-pid-file`: Do not write PID file (overrides *-p/pidfile*)
- `--version`: Show version of Graylog and exit

Problems with IPv6 vs. IPv4?

If your *graylog-server* instance refuses to listen on IPv4 addresses and always chooses for example a *rest_listen_address* like `:::12900` you can tell the JVM to prefer the IPv4 stack.

Add the *java.net.preferIPv4Stack* flag in your *graylogctl* script or from wherever you are calling the *graylog.jar*:

```
~$ sudo -u graylog java -Djava.net.preferIPv4Stack=true -jar graylog.jar
```

Manual setup: graylog-web-interface on Linux

Prerequisites

The only thing you need is at least one compatible *graylog-server* node. Please use the same version number to make sure that it is compatible.

You also **must** use **Java 7!** Java 6 is not compatible with Graylog and will also not receive any more publicly available bug and security fixes by Oracle.

Downloading and extracting the web-interface

Download the package from the [download pages](#).

Extract the archive:

```
~$ tar xvfz graylog-web-interface-VERSION.tgz
~$ cd graylog-web-interface-VERSION
```

Configuring the web interface

Open `conf/graylog-web-interface.conf` and set the two following variables:

- `graylog2-server.uris="http://127.0.0.1:12900/"`: This is the list of *graylog-server* nodes the web interface will try to use. You can configure one or multiple, separated by commas. Use the *rest_listen_uri* (configured in *graylog.conf*) of your *graylog-server* instances here.
- `application.secret=""`: A secret for encryption. Use a long, randomly generated string here. (for example generated using `pwgen -N 1 -s 96`)

Starting the web interface

You need to have Java installed. Running the OpenJDK is totally fine and should be available on all platforms. For example on Debian it is:

```
~$ apt-get install openjdk-7-jre
```

You need at least Java 7 as Java 6 has reached EOL.

Now start the web interface:

```
~$ bin/graylog-web-interface
Play server process ID is 5723
[info] play - Application started (Prod)
[info] play - Listening for HTTP on /0:0:0:0:0:0:0:9000
```

The web interface will listen on port 9000. You should see a login screen right away after pointing your browser to it. Log in with username `admin` and the password you configured at `root_password_sha2` in the `graylog.conf` of your `graylog-server`.

Changing the listen port and address works like this:

```
~$ bin/graylog-web-interface -Dhttp.port=1234 -Dhttp.address=127.0.0.1
```

Java generally prefers to bind to an IPv6 address if that is supported by your system, while you might want to prefer IPv4. To change Java's default preference you can pass `-Djava.net.preferIPv4Stack=true` to the startup script:

```
~$ bin/graylog-web-interface -Djava.net.preferIPv4Stack=true
```

All those `-D` settings can also be added to the `JAVA_OPTS` environment variable which is being read by the startup script, too.

You can start the web interface in background for example like this:

```
~$ nohup bin/graylog-web-interface &
```


Custom configuration file path

You can put the configuration file into another directory like this:

```
~$ bin/graylog-web-interface -Dconfig.file=/etc/graylog-web-interface.conf
```

Create a message input and send a first message

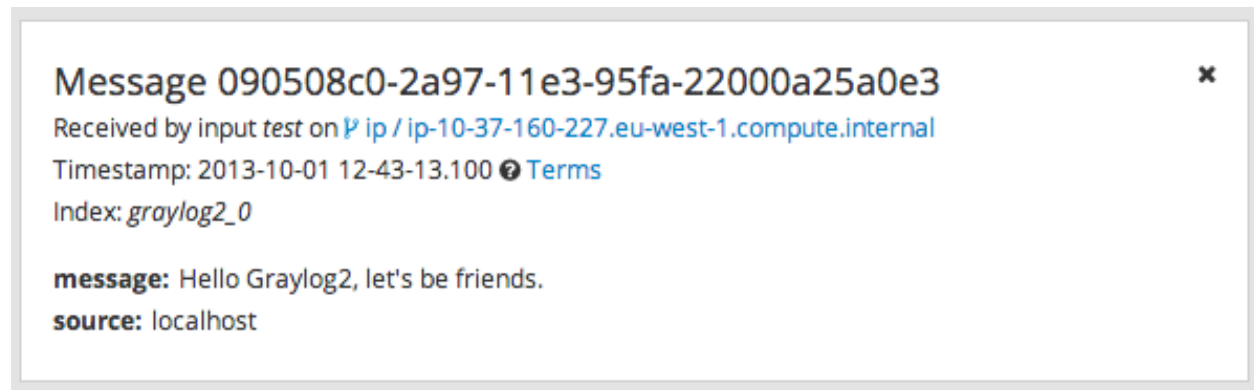
Log in to the web interface and navigate to *System -> Nodes*. Select your `graylog-server` node there and click on *Manage inputs*.



Launch a new *Raw/Plaintext UDP* input, listening on port 9099 and listening on 127.0.0.1. No need to configure anything else for now. The list of running inputs on that node should show you your new input right away. Let's send a message in:

```
echo "Hello Graylog, let's be friends." | nc -w 1 -u 127.0.0.1 9099
```

This has sent a short string to the raw UDP input you just opened. Now search for *friends* using the searchbar on the top and you should already see the message you just sent in. Click on it in the table and see it in detail:



You have just sent your first message to Graylog! Why not spawn a syslog input and point some of your servers to it? You could also create some user accounts for your colleagues.

HTTPS

Enabling HTTPS is easy. Just start the web interface like this:

```
bin/graylog-web-interface -Dhttps.port=443
```

This will generate self-signed certificate. To use proper certificates you must configure a Java key store. Most signing authorities provide instructions on how to create a Java keystore and the official keystore utility docs can be found [here](#).

- `https.keyStore` The path to the keystore containing the private key and certificate, if not provided generates a keystore for you
- `https.keyStoreType` The key store type, defaults to JKS
- `https.keyStorePassword` The password, defaults to a blank password
- `https.keyStoreAlgorithm` The key store algorithm, defaults to the platforms default algorithm

To disable HTTP without SSL completely and enforce HTTPS, use this parameter:

```
-Dhttp.port=disabled
```

Configuring logging

The default setting of the web interface is to write its own logs to STDOUT. You can take control of the logging by specifying an own [Logback](#) configuration file to use:

```
bin/graylog-web-interface -Dlogger.file=/etc/graylog-web-interface-log.xml
```

This is an example Logback configuration file that has a disabled STDOUT appender and an enabled appender that writes to a file (`/var/log/graylog/web/graylog-web-interface.log`), keeps 30 days of logs in total and creates a new log file if a file should have reached a size of 100MB:

```
<configuration>
  <!--
  <appender name="STDOUT" class="ch.qos.logback.core.ConsoleAppender">
    <encoder>
      <pattern>%date %-5level [%thread] - [%logger]- %msg%n</pattern>
    </encoder>
  </appender>
  -->

  <appender name="ROLLING_FILE" class="ch.qos.logback.core.rolling.
  ↪RollingFileAppender">
    <file>/var/log/graylog/web/graylog-web-interface.log</file>
    <rollingPolicy class="ch.qos.logback.core.rolling.TimeBasedRollingPolicy">
      <FileNamePattern>/var/log/graylog/web/graylog-web-interface.log.%d{yyyy-
  ↪MM-dd}.%i.log.gz</FileNamePattern>
      <MaxHistory>30</MaxHistory>
      <timeBasedFileNamingAndTriggeringPolicy class="ch.qos.logback.core.
  ↪rolling.SizeAndTimeBasedFNATP">
        <maxFileSize>100MB</maxFileSize>
        </timeBasedFileNamingAndTriggeringPolicy>
      </rollingPolicy>
      <encoder class="ch.qos.logback.classic.encoder.PatternLayoutEncoder">
        <pattern>%date [%thread] %-5level %logger{36} - %msg%n</pattern>
      </encoder>
    </appender>

  <root level="INFO">
    <!--<appender-ref ref="STDOUT" />-->
    <appender-ref ref="ROLLING_FILE" />
  </root>
</configuration>
```

Operating system packages

Until configuration management systems made their way into broader markets and many datacenters, one of the most common ways to install software on Linux servers was to use operating system packages. Debian has DEB, Red Hat has RPM and many other distributions are based on those or come with own package formats. Online repositories of

software packages and corresponding package managers make installing and configuring new software a matter of a single command and a few minutes of time.

Graylog offers official DEB and RPM package repositories for Ubuntu 12.04, Ubuntu 14.04, Debian 7 and CentOS 6.

The repositories can be setup by installing a single package. Once that's done the Graylog packages can be installed via `apt-get` or `yum`. The packages can also be downloaded with a web browser at <https://packages.graylog2.org/> if needed.

Make sure to install and configure MongoDB and Elasticsearch before starting the Graylog services.

Ubuntu 14.04

Download and install `graylog-1.0-repository-ubuntu14.04_latest.deb` via `dpkg(1)` and also make sure that the `apt-transport-https` package is installed:

```
$ wget https://packages.graylog2.org/repo/packages/graylog-1.0-repository-ubuntu14.04_
↳latest.deb
$ sudo dpkg -i graylog-1.0-repository-ubuntu14.04_latest.deb
$ sudo apt-get install apt-transport-https
$ sudo apt-get update
$ sudo apt-get install graylog-server graylog-web
```

Ubuntu 12.04

Download and install `graylog-1.0-repository-ubuntu12.04_latest.deb` via `dpkg(1)` and also make sure that the `apt-transport-https` package is installed:

```
$ wget https://packages.graylog2.org/repo/packages/graylog-1.0-repository-ubuntu12.04_
↳latest.deb
$ sudo dpkg -i graylog-1.0-repository-ubuntu12.04_latest.deb
$ sudo apt-get install apt-transport-https
$ sudo apt-get update
$ sudo apt-get install graylog-server graylog-web
```

Debian 7

Download and install `graylog-1.0-repository-debian7_latest.deb` via `dpkg(1)` and also make sure that the `apt-transport-https` package is installed:

```
$ wget https://packages.graylog2.org/repo/packages/graylog-1.0-repository-debian7_
↳latest.deb
$ sudo dpkg -i graylog-1.0-repository-debian7_latest.deb
$ sudo apt-get install apt-transport-https
$ sudo apt-get update
$ sudo apt-get install graylog-server graylog-web
```

CentOS 6

Download and install `graylog-1.0-repository-el6_latest.rpm` via `rpm(8)`:

```
$ sudo rpm -Uvh https://packages.graylog2.org/repo/packages/graylog-1.0-repository-  
→el6_latest.rpm  
$ yum install graylog-server graylog-web
```

Please open an [issue](#) in the [Github repository](#) if you run into any packaging related issues. **Thank you!**

Chef, Puppet, Ansible, Vagrant

The DevOps movement turbocharged market adoption of the newest generation of configuration management and orchestration tools like [Chef](#), [Puppet](#) or [Ansible](#). Graylog offers official scripts for all three of them:

- <https://supermarket.chef.io/cookbooks/graylog2>
- <https://forge.puppetlabs.com/graylog2/graylog2>
- <https://galaxy.ansible.com/list#/roles/3162>

There are also official [Vagrant](#) images if you want to spin up a local virtual machine quickly. (Note that the pre-built *Virtual machine appliances* are a preferred way to run Graylog in production)

- <https://github.com/Graylog2/graylog2-images/tree/master/vagrant>

Amazon Web Services

The *Virtual machine appliances* are supporting Amazon Web Services EC2 AMIs as platform.

Docker

The *Virtual machine appliances* are supporting Docker as runtime.

Microsoft Windows

Unfortunately there is no officially supported way to run Graylog on Microsoft Windows operating systems even though all parts run on the Java Virtual Machine. We recommend to run the *Virtual machine appliances* on a Windows host. It should be technically possible to run Graylog on Windows but it is most probably not worth the time to work your way around the cliffs.

Should you require running Graylog on Windows, you need to disable the message journal in `graylog-server` by changing the following setting in the `graylog.conf`:

```
message_journal_enabled = false
```

Due to restrictions of how Windows handles file locking the journal will not work correctly. This will be improved in future versions.

Please note that this impacts Graylog's ability to buffer messages, so we strongly recommend running the Linux-based OVAs on Windows.

Configuring and tuning Elasticsearch

We strongly recommend to use a dedicated Elasticsearch cluster for your Graylog setup. If you are using a shared Elasticsearch setup, a problem with indices unrelated to Graylog might turn the cluster status to yellow or red and impact the availability and performance of your Graylog setup.

Configuration

Configuration of graylog-server nodes

The most important settings to make a successful connection are the Elasticsearch cluster name and the discovery mode. Graylog is able to discover the Elasticsearch nodes using multicast. This is great for development and proof of concepts but we recommend to use classic unicast discovery in production.

Cluster Name

You need to tell `graylog-server` which Elasticsearch cluster to join. The Elasticsearch cluster default name is `elasticsearch` and configured for every Elasticsearch node in its `elasticsearch.yml` configuration file with the setting `cluster.name`. Configure the same name in every `graylog.conf` as `elasticsearch_cluster_name`. We recommend to call the cluster `graylog-production` and not `elasticsearch`.

The `elasticsearch.yml` file is typically located in `/etc/elasticsearch/`.

Discovery mode

The default discovery mode is multicast. Graylog will try to find other Elasticsearch nodes automatically. This usually works fine when everything is running on the same system but gets problematic quickly when running in a bigger network topology. We recommend to use unicast for production setups. Configure Zen unicast discovery in Graylog with the following lines in your configuration file:

```
# Disable multicast
elasticsearch_discovery_zen_ping_multicast_enabled = false
# List of Elasticsearch nodes to connect to
elasticsearch_discovery_zen_ping_unicast_hosts = es-node-1.example.org:9300,es-node-2.
↳example.org:9300
```

Also make sure to configure **Zen unicast discovery** in the Elasticsearch configuration file by adding the `discovery.zen.ping.multicast.enabled` and `discovery.zen.ping.unicast.hosts` setting with the list of Elasticsearch nodes to `elasticsearch.yml`:

```
discovery.zen.ping.multicast.enabled: false
discovery.zen.ping.unicast.hosts: ["es-node-1.example.org:9300" , "es-node-2.example.
↳org:9300"]
```

The Elasticsearch default communication port is *9300/tcp* (not to be confused with the HTTP interface running on port *9200/tcp* by default). The communication port can be changed in the Elasticsearch configuration file (`elasticsearch.yml`) with the configuration setting `transport.tcp.port`. Make sure that Elasticsearch binds to a network interface that Graylog can connect to (see `network.host`).

Configuration of Elasticsearch nodes

Disable dynamic scripting

Elasticsearch prior to version 1.2 had an insecure default configuration which could lead to a remote code execution. (see [here](#) and [here](#) for details)

Make sure to add `script.disable_dynamic: true` to the `elasticsearch.yml` file to disable the dynamic scripting feature and prevent possible remote code executions.

Control access to Elasticsearch ports

Since Elasticsearch has no authentication mechanism at time of this writing, make sure to restrict access to the Elasticsearch ports (default: *9200/tcp* and *9300/tcp*). Otherwise the data is readable by anyone who has access to the machine over network.

Open file limits

Because Elasticsearch has to keep a lot of files open simultaneously it requires a higher open file limit that the usual operating system defaults allow. **Set it to at least 64000 open file descriptors.**

Graylog will show a notification in the web interface when there is a node in the Elasticsearch cluster which has a too low open file limit.

Read about how to raise the open file limit in the corresponding [Elasticsearch documentation page](#).

Heap size

It is strongly recommended to raise the standard size of heap memory allocated to Elasticsearch. Just set the `ES_HEAP_SIZE` environment variable to for example `24g` to allocate 24GB. We recommend to use around 50% of the available system memory for Elasticsearch (when running on a dedicated host) to leave enough space for the system caches that Elasticsearch uses a lot.

Tuning Elasticsearch

Graylog is already setting specific configuration per index it creates. This is enough tuning for a lot of use cases and setups. A more detailed guide on deeper tuning of Elasticsearch is following.

Cluster Status explained

Elasticsearch provides a classification for the cluster health:

RED

The red status indicates that some or all of the primary shards are not available. In this state, no searches can be performed until all primary shards are restored.

YELLOW

The yellow status means that all of the primary shards are available but some or all shard replicas are not.

With only one Elasticsearch node, the cluster state cannot become green because shard replicas cannot be assigned. This can be solved by adding another Elasticsearch node to the cluster.

If the cluster is supposed to have only one node it is okay to be in the yellow state.

GREEN

The cluster is fully operational. All primary and replica shards are available.

Sending in log data

A Graylog setup is pretty worthless without any data in it. This page explains the basic principles of getting your data into the system and also explains common fallacies.

What are Graylog message inputs?

Message inputs are the Graylog parts responsible for accepting log messages. They are launched from the web interface (or the REST API) in the *System -> Inputs* section and are launched and configured without the need to restart any part of the system.

Content packs

Content packs are bundles of Graylog input, extractor, stream, dashboard, and output configurations that can provide full support for a data source. Some content packs are shipped with Graylog by default and some are available from the website. Content packs that were downloaded from [here](#) can be imported using the Graylog web interface.

You can load and even create own content packs from the *System -> Content Packs* section of your Graylog web interface.

Syslog

Graylog is able to accept and parse [RFC 5424](#) and [RFC 3164](#) compliant syslog messages and supports TCP transport with both the octet counting or termination character methods. UDP is also supported and the recommended way to send log messages in most architectures.

Many devices, especially routers and firewalls, do not send RFC compliant syslog messages. This might result in wrong or completely failing parsing. In that case you might have to go with a combination of *raw/plaintext* message inputs that do not attempt to do any parsing and *Extractors*.

Rule of thumb is that messages forwarded by `rsyslog` or `syslog-ng` are usually parsed flawlessly.

Sending syslog from Linux hosts

rsyslog

Forwarding syslog messages with rsyslog is easy. The only important thing to get the most out of your logs is following [RFC 5424](#). The following examples configures your rsyslog daemon to send RFC 5424 date to Graylog syslog inputs:

UDP:

```
$template GRAYLOGRFC5424, "<%PRI%>%PROTOCOL-VERSION% %TIMESTAMP:::date-rfc3339%
->%HOSTNAME% %APP-NAME% %PROCID% %MSGID% %STRUCTURED-DATA% %msg%\n"
*. * @graylog.example.org:514;GRAYLOGRFC5424
```

TCP:

```
$template GRAYLOGRFC5424, "<%PRI%>%PROTOCOL-VERSION% %TIMESTAMP:::date-rfc3339%
->%HOSTNAME% %APP-NAME% %PROCID% %MSGID% %STRUCTURED-DATA% %msg%\n"
*. * @@graylog.example.org:514;GRAYLOGRFC5424
```

(The difference between UDP and TCP is using @ instead of @@ as target descriptor.)

Alternatively, the rsyslog built-in template `RSYSLOG_SyslogProtocol23Format` sends log messages in the same format as above. This exists in rsyslog versions of at least 5.10 or later.

The UDP examples above becomes:

```
*. * @graylog.example.org:514;RSYSLOG_SyslogProtocol23Format
```

syslog-ng

Configuring syslog-ng to send syslog to Graylog is equally simple. Use the `syslog` function to send RFC 5424 formatted syslog messages via TCP to the remote Graylog host:

```
# Define TCP syslog destination.
destination d_net {
    syslog("graylog.example.org" port(514));
};
# Tell syslog-ng to send data from source s_src to the newly defined syslog_
->destination.
log {
    source(s_src); # Defined in the default syslog-ng configuration.
    destination(d_net);
};
```

Sending syslog from MacOS X hosts

Sending log messages from MacOS X syslog daemons is easy. Just define a `graylog-server` instance as UDP log target by adding this line in your `/etc/syslog.conf`:

```
*. * @graylog.example.org:514
```

Now restart `syslogd`:

```
$ sudo launchctl unload /System/Library/LaunchDaemons/com.apple.syslogd.plist
$ sudo launchctl load /System/Library/LaunchDaemons/com.apple.syslogd.plist
```

Important: If syslogd was running as another user you might end up with multiple syslogd instances and strange behaviour of the whole system. Please check that only one syslogd process is running:

```
$ ps aux | grep syslog
lennart      58775   0.0   0.0   2432768    592  s004  S+   6:10PM   0:00.00  grep_
↳ syslog
root         58759   0.0   0.0   2478772   1020  ??   Ss   6:09PM   0:00.01  /usr/
↳ sbin/syslogd
```

That's it! Your MacOS X syslog messages should now appear in your Graylog system.

GELF / Sending from applications

The Graylog Extended Log Format (GELF) is a log format that avoids the shortcomings of classic plain syslog and is perfect to logging from your application layer. It comes with optional compression, chunking and most importantly a clearly defined structure. There are [dozens of GELF libraries](#) for many frameworks and programming languages to get you started.

Read more about GELF [on graylog.org](http://graylog.org).

GELF via HTTP

You can send in all GELF types via HTTP, including uncompressed GELF that is just a plain JSON string.

After launching a GELF HTTP input you can use the following endpoints to send messages:

```
http://graylog.example.org:[port]/gelf (POST)
```

Try sending an example message using curl:

```
curl -XPOST http://graylog.example.org:12202/gelf -p0 -d '{"short_message":"Hello_
↳ there", "host":"example.org", "facility":"test", "_foo":"bar"}'
```

Both keep-alive and compression are supported via the common HTTP headers. The server will return a 202 Accepted when the message was accepted for processing.

Microsoft Windows

Our recommended way to forward Windows log data (for example EventLog) to Graylog is to use the open source [nxlog community edition](#). It comes with a native Graylog GELF output that nicely structures your log messages.

Heroku

Heroku allows you to forward the logs of your application to a custom syslog server by creating a so called [Syslog drain](#). The drain sends all logs to the configured server(s) via TCP. Following example shows you how to configure Graylog to receive the Heroku logs and extract the different fields into a structured log message.

Creating a Graylog input for Heroku log messages

Create a new **RAW/Plaintext TCP** input as shown below.

The screenshot shows the Graylog web interface. At the top, it says 'System / Inputs'. Below that is a section titled 'Inputs in Cluster' with a sub-header 'Inputs in Cluster' and a description: 'Graylog2 nodes accept data via inputs. Launch or terminate as many inputs as you want here.' There is a dropdown menu showing 'Raw/Plaintext TCP' and a green button labeled 'Launch new input'.

The 'Launch new input' dialog box is open, titled 'Launch new input: Raw/Plaintext TCP'. It contains the following fields and options:

- Node(s) to spawn input on:** Select the node you want to spawn this input on. A dropdown menu shows '3c1749a2 /'.
- or:** Global input (started on all nodes)
- Title:** Select a name of your new input that describes it. A text input field contains 'Heroku Drain'.
- Port:** Port to listen on. A text input field contains '5556'.
- Override source (optional):** The source is a hostname derived from the received packet by default. Set this if you want to override it with a custom string.

At the bottom of the dialog box, there are two buttons: 'Close' and 'Launch'.

The Graylog `Extractor` library contains a set of extractors to parse the Heroku log format. You can import that set into the newly created input so all parts of the log messages will be extracted into separate fields:

Open the extractor management for the input.

Heroku Drain (Raw/Plaintext TCP) **running** Started by Bernd Ahlers on P 3c1749a2 / 2 hours ago **Terminate** Action ▾

Network ID: 0B ← 0B (total: → 9.6kiB ← 0B)

Total connections: 18 (0 active)

```
port: 5556
override_source:
bind_address: 0.0.0.0
recv_buffer_size: 1048576
```

- Manage extractors
- Show metrics
- Add static field
- Messages from this input

Go to the extractor import.

System / Nodes / 3c1749a2 / Input: Heroku Drain / Extractors

Extractors of *Heroku Drain*

Extractors are applied on every message that is received by this input. Use them to extract and transform any text data into fields that allow you easy filtering and analysis later on. Example: Extract the HTTP response code from a log message, transform it to a numeric field and attach it as `http_response_code` to the message.

- Import extractors
- Export extractors

Paste the extractor JSON string into the form and submit.

System / Nodes / 3c1749a2 / Input: Heroku Drain / Extractors / Import

Import extractors to input *Heroku Drain*

Exported extractors can be imported to an input. All you need is the JSON export of extractors from any other Graylog2 setup or from the extractor directory on graylog2.org.

```
{
  "extractors": [
    {
      "condition_type": "none",
      "condition_value": "",
      "converters": [
        {
          "config": {},
          "type": "syslog_pri_level"
        }
      ],
      "cursor_strategy": "copy",
      ...
    }
  ]
}
```

That is all that is needed on the Graylog side. Make sure your firewall setup allows incoming connections on the inputs port!

System / Nodes / 3c1749a2 / Input: Heroku Drain / Extractors Actions ▾

Extractors of *Heroku Drain*

Extractors are applied on every message that is received by this input. Use them to extract and transform any text data into fields that allow you easy filtering and analysis later on. Example: Extract the HTTP response code from a log message, transform it to a numeric field and attach it as `http_response_code` to the message.

➔ Add extractor

Start by loading a message so we have an example to work on.

Load a message recently received by this input Manually load a message

⚙️ Configured extractors ↕ Change execution order by dragging on the icons.

☰ **Level/Severity** (Regular expression)

Trying to extract data from `message` into `level`, leaving the original intact. Added by Bernd Ahlers Details Remove

☰ **Facility** (Regular expression)

Trying to extract data from `message` into `facility`, leaving the original intact. Added by Bernd Ahlers Details Remove

Configuring Heroku to send data to your Graylog setup

Heroku has a detailed [documentation](#) regarding the Syslog drains feature. The following example shows everything that is needed to setup the drain for you application:

```
$ cd path/to/your/heroku/app
$ heroku drains
No drains for this app
$ heroku drains:add syslog://graylog.example.com:5556
Successfully added drain syslog://graylog.example.com:5556
$ heroku drains
syslog://graylog.example.com:5556 (d.8cf52d32-7d79-4653-baad-8cb72bb23ee1)
```

The [Heroku CLI](#) tool needs to be installed for this to work.

Your Heroku application logs should now show up in the search results of your Graylog instance.

Ruby on Rails

This is easy: You just need to combine a few components.

Log all requests and logger calls into Graylog

The recommended way to send structured information (i.e. HTTP return code, action, controller, ... in additional fields) about every request and explicit `Rails.logger` calls is easily accomplished using the [GELF gem](#) and [lograge](#). Lograge builds one combined log entry for every request (instead of several lines like the standard Rails logger) and has a Graylog output since version 0.2.0.

Start by adding Lograge and the GELF gem to your Gemfile:

```
gem "gelf"  
gem "lograge"
```

Now configure both in your Rails application. Usually `config/environments/production.rb` is a good place for that:

```
config.lograge.enabled = true  
config.lograge.formatter = Lograge::Formatters::Graylog2.new  
config.logger = GELF::Logger.new("graylog.example.org", 12201, "WAN", { :host =>  
  ↳"hostname-of-this-app", :facility => "heroku" })
```

This configuration will also send all explicit `Rails.logger` calls (e.g. `Rails.logger.error "Something went wrong"`) to Graylog.

Log only explicit logger calls into Graylog

If you don't want to log information about every request, but only explicit `Rails.logger` calls, it is enough to only configure the Rails logger.

Add the GELF gem to your Gemfile:

```
gem "gelf"
```

...and configure it in your Rails application. Usually `config/environments/production.rb` is a good place for that:

```
config.logger = GELF::Logger.new("graylog.example.org", 12201, "WAN", { :host =>  
  ↳"hostname-of-this-app", :facility => "heroku" })
```

Heroku

You need to apply a workaround if you want custom logging on Heroku. The reason for this is that Heroku injects an own logger (`rails_log_stdout`), that overwrites your custom one. The workaround is to add a file that makes Heroku think that the logger is already in your application:

```
$ touch vendor/plugins/rails_log_stdout/heroku_fix
```

Raw/Plaintext inputs

The built-in *raw/plaintext* inputs allow you to parse any text that you can send via TCP or UDP. No parsing is applied at all by default until you build your own parser using custom *Extractors*. This is a good way to support any text-based logging format.

You can also write *Plugins* if you need extreme flexibility.

JSON path from HTTP API input

The JSON path from HTTP API input is reading any JSON response of a REST resource and stores a field value of it as a Graylog message.

Example

Let's try to read the download count of a release package stored on GitHub for analysis in Graylog. The call looks like this:

```
$ curl -XGET https://api.github.com/repos/YourAccount/YourRepo/releases/assets/12345
{
  "url": "https://api.github.com/repos/YourAccount/YourRepo/releases/assets/12345",
  "id": 12345,
  "name": "somerelease.tgz",
  "label": "somerelease.tgz",
  "content_type": "application/octet-stream",
  "state": "uploaded",
  "size": 38179285,
  "download_count": 9937,
  "created_at": "2013-09-30T20:05:01Z",
  "updated_at": "2013-09-30T20:05:46Z"
}
```

The attribute we want to extract is `download_count` so we set the JSON path to `$.download_count`.

This will result in a message in Graylog looking like this:



You can use Graylog to analyse your download counts now.

JSONPath

JSONPath can do much more than just selecting a simple known field value. You can for example do this to select the first `download_count` from a list of releases where the field `state` has the value `uploaded`:

```
$.releases[?(@.state == 'uploaded')][0].download_count
```

...or only the first download count at all:

```
$.releases[0].download_count
```

You can [learn more about JSONPath here](#).

Reading from files

Graylog is currently not providing an out-of-the-box way to read log messages from files. We do however recommend two fantastic tools to do that job for you. Both come with native Graylog (GELF) outputs:

- fluentd
- logstash

The search query language

Syntax

The search syntax is very close to the Lucene syntax. By default all message fields are included in the search if you don't specify a message field to search in.

Messages that include the term *ssh*:

```
ssh
```

Messages that include the term *ssh* or *login*:

```
ssh login
```

Messages that include the exact phrase *ssh login*:

```
"ssh login"
```

Messages where the field *type* includes *ssh*:

```
type:ssh
```

Messages where the field *type* includes *ssh* or *login*:

```
type:(ssh login)
```

Messages where the field *type* includes the exact phrase *ssh login*:

```
type:"ssh login"
```

Messages that do not have the field *type*:

```
_missing_:type
```

Messages that have the field *type*:

```
_exists_:type
```

By default all terms or phrases are OR connected so all messages that have at least one hit are returned. You can use **Boolean operators and groups** for control over this:

```
"ssh login" AND source:example.org
("ssh login" AND (source:example.org OR source:another.example.org)) OR _exists_
↪:always_find_me
```

You can also use the NOT operator:

```
"ssh login" AND NOT source:example.org
NOT example.org
```

Wildcards: Use ? to replace a single character or * to replace zero or more characters:

```
source:*.org
source:exam?le.org
source:exam?le.*
```

Note that leading wildcards are disabled to avoid excessive memory consumption! You can enable them in your `graylog-server.conf`: `allow_leading_wildcard_searches = true`

Fuzziness: You can search for similar but not equal terms:

```
ssh logni~
source:exmaple.org~
```

This is using the [Damerau-Levenshtein distance](#) with a default distance of 2. You can change the distance like this:

```
source:exmaple.org~1
```

You can also use the fuzzyness operator to do a **proximity** search where the terms in a phrase can have different/fuzzy distances from each other and don't have to be in the defined order:

```
"foo bar"~5
```

Numeric fields support **range queries**. Ranges in square brackets are inclusive, curly brackets are exclusive and can even be combined:

```
http_response_code:[500 TO 504]
http_response_code:{400 TO 404}
bytes:{0 TO 64}
http_response_code:[0 TO 64}
```

You can also do searches with one side unbounded:

```
http_response_code:>400
http_response_code:<400
http_response_code:>=400
http_response_code:<=400
```

It is also possible to combine unbounded range operators:

```
http_response_code:(>=400 AND <500)
```


Escaping

The following characters must be escaped with a backslash:

```
& | : \ / + - ! ( ) { } [ ] ^ " ~ * ?
```

Example:

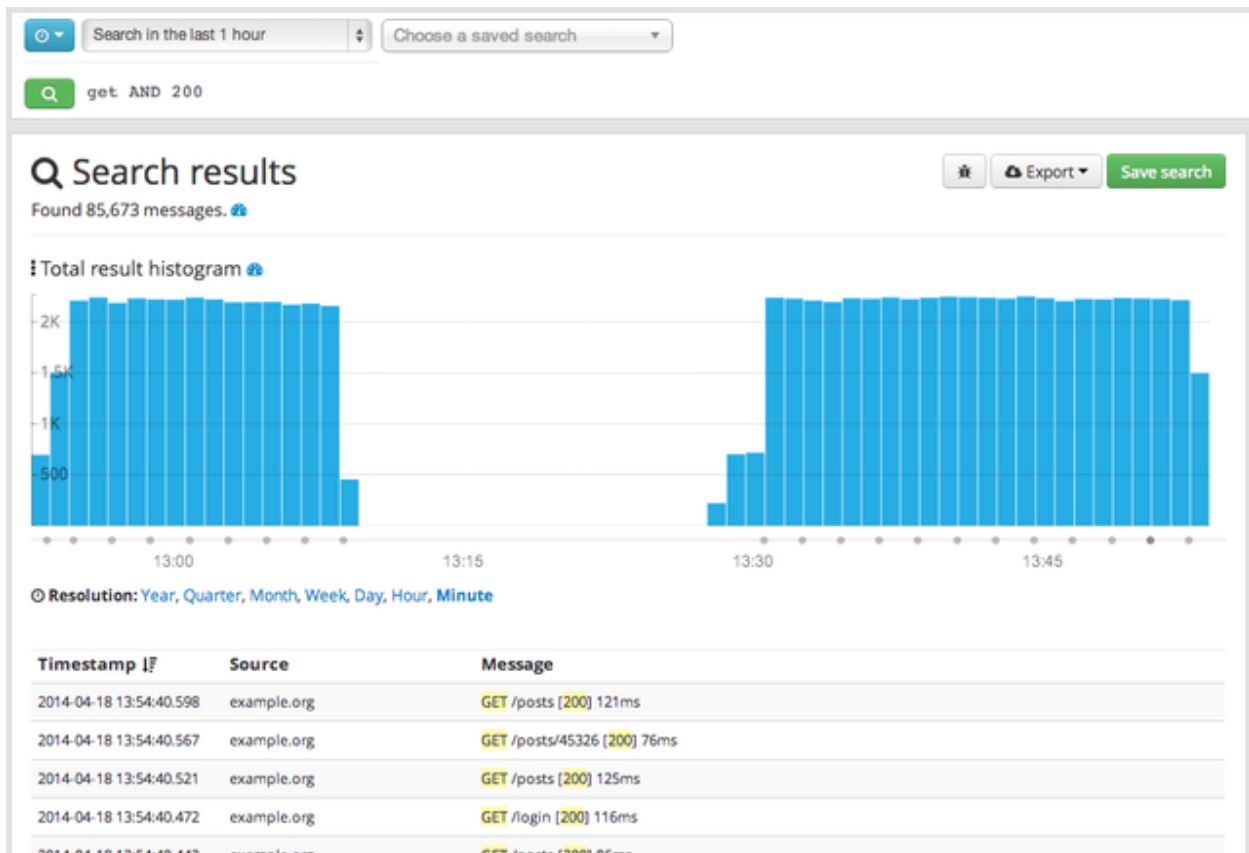
```
resource:\posts\45326
```

Time frame selector

The time frame selector defines in what time range to search in. It offers three different ways of selecting a time range and is vital for search speed: If you know you are only interested in messages of the last hour, only search in that time frame. This will make Graylog search in relevant indices only and greatly reduce system load and required resources. You can read more about this here: *The Graylog index model explained*

Search result highlighting

Graylog supports search result highlighting since v0.20.2:



Enabling/Disabling search result highlighting

Using search result highlighting will result in slightly higher resource consumption of searches. You can enable and disable it using a configuration parameter in the `graylog.conf` of your `graylog-server` nodes:

```
allow_highlighting = true
```

What are streams?

The Graylog streams are a mechanism to route messages into categories in realtime while they are processed. You define rules that instruct Graylog which message to route into which streams. Imagine sending these three messages to Graylog:

```
message: INSERT failed (out of disk space)
level: 3 (error)
source: database-host-1

message: Added user 'foo'.
level: 6 (informational)
source: database-host-2

message: smtp ERR: remote closed the connection
level: 3 (error)
source: application-x
```

One of the many things that you could do with streams is creating a stream called *Database errors* that is catching every error message from one of your database hosts.

Create a new stream with these rules: (stream rules are AND connected)

- Field `level` must be greater than 4
- Field `source` must match regular expression `^database-host-\d+`

This will route every new message with a `level` higher than *WARN* and a `source` that matches the database host regular expression into the stream.

A message will be routed into every stream that has all its rules matching. This means that a message can be part of many streams and not just one.

The stream is now appearing in the streams list and a click on its title will show you all database errors.

The next parts of this document cover how to be alerted in case of too many errors, some specific error types that should never happen or how to forward the errors to another system or endpoint.

What's the difference to saved searches?

The biggest difference is that streams are processed in realtime. This allows realtime alerting and forwarding to other systems. Imagine forwarding your database errors to another system or writing them to a file by regularly reading them from the message storage. Realtime streams do this much better.

Another difference is that searches for complex stream rule sets are always comparably cheap to perform because a message is *tagged* with stream IDs when processed. A search for Graylog internally always looks like this, no matter how many stream rules you have configured:

```
streams:[STREAM_ID]
```

Building a query with all rules would cause significantly higher load on the message storage.

How do I create a stream?

1. Navigate to the streams section from the top navigation bar
2. Click "Create stream"
3. Save the stream after entering a name and a description. For example *All error messages* and *Catching all error messages from all sources*
4. The stream is now saved but **not yet activated**. Add stream rules in the next dialogue. Try it against some messages by entering a message ID on the same page. Save the rules when the right messages are matched or not matched.
5. The stream is marked as *paused* in the list of streams. Activate the stream by hitting *Resume this stream* in the *Action* dropdown.

Alerts

You can define conditions that trigger alerts. For example whenever the stream *All production exceptions* has more than 50 messages per minute or when the field *milliseconds* had a too high standard deviation in the last five minutes.

Hit *Manage alerts* in the stream *Action* dropdown to see already configured alerts, alerts that were fired in the past or to configure new alert conditions.

Graylog ships with default *alert callbacks* and can be extended with [plugins](#)

What is the difference between alert callbacks and alert receivers?

There are two type of actions to be triggered when an alert is fired: Alert callbacks or an email to a list of alert receivers.

Alert callbacks are single actions that are just called once. For example: The *Email Alert Callback* is triggering an email to exactly one receiver and the *HTTP Alert Callback* is calling a HTTP endpoint once.

The alert receivers in difference will all receive an email about the same alert.

Email Alert Callback

The email alert callback can be used to send an email to the configured alert receivers when the conditions are triggered. Three configuration options are available for the alert callback to customize the email that will be sent.

Create new Email Alert Callback ×

Sender

The sender of sent out mail alerts

E-Mail Body (optional)

```
#####  
Alert Description: ${check_result.resultDescription}  
Date: ${check_result.triggeredAt}  
Stream ID: ${stream.id}  
Stream title: ${stream.title}  
Stream description: ${stream.description}  
${if stream_url}Stream URL: ${stream_url}${end}  
  
Triggered condition: ${check_result.triggeredCondition}  
#####  
  
${if backlog}Last messages accounting for this alert:  
${foreach backlog message}${message}  
  
${end}${else}<No backlog>  
${end}
```

The template to generate the body from

E-Mail Subject

The subject of sent out mail alerts

The *email body* and *email subject* are JMTE templates. JMTE is a minimal template engine that supports variables, loops and conditions. See the [JMTE documentation](#) for a language reference.

We expose the following objects to the templates.

stream The stream this alert belongs to.

- `stream.id` ID of the stream
- `stream.title` title of the stream
- `stream.description` stream description

stream_url A string that contains the HTTP URL to the stream.

check_result The check result object for this stream.

- `check_result.triggeredCondition` string representation of the triggered alert condition
- `check_result.triggeredAt` date when this condition was triggered
- `check_result.resultDescription` text that describes the check result

backlog A list of message objects. Can be used to iterate over the messages via `foreach`.

message (only available via iteration over the backlog object) The message object has several fields with details about the message. When using the message object without accessing any fields, the `toString()` method of the underlying Java object is used to display it.

- `message.id` autogenerated message id
- `message.message` the actual message text
- `message.source` the source of the message
- `message.timestamp` the message timestamp
- `message.fields` map of key value pairs for all the fields defined in the message

The `message.fields` fields can be useful to get access to arbitrary fields that are defined in the message. For example `message.fields.full_message` would return the `full_message` of a GELF message.

Outputs

The stream output system allows you to forward every message that is routed into a stream to other destinations.

Outputs are managed globally (like message inputs) and not for single streams. You can create new outputs and activate them for as many streams as you like. This way you can configure a forwarding destination once and select multiple streams to use it.

Graylog ships with default outputs and can be extended with [plugins](#).

Use cases

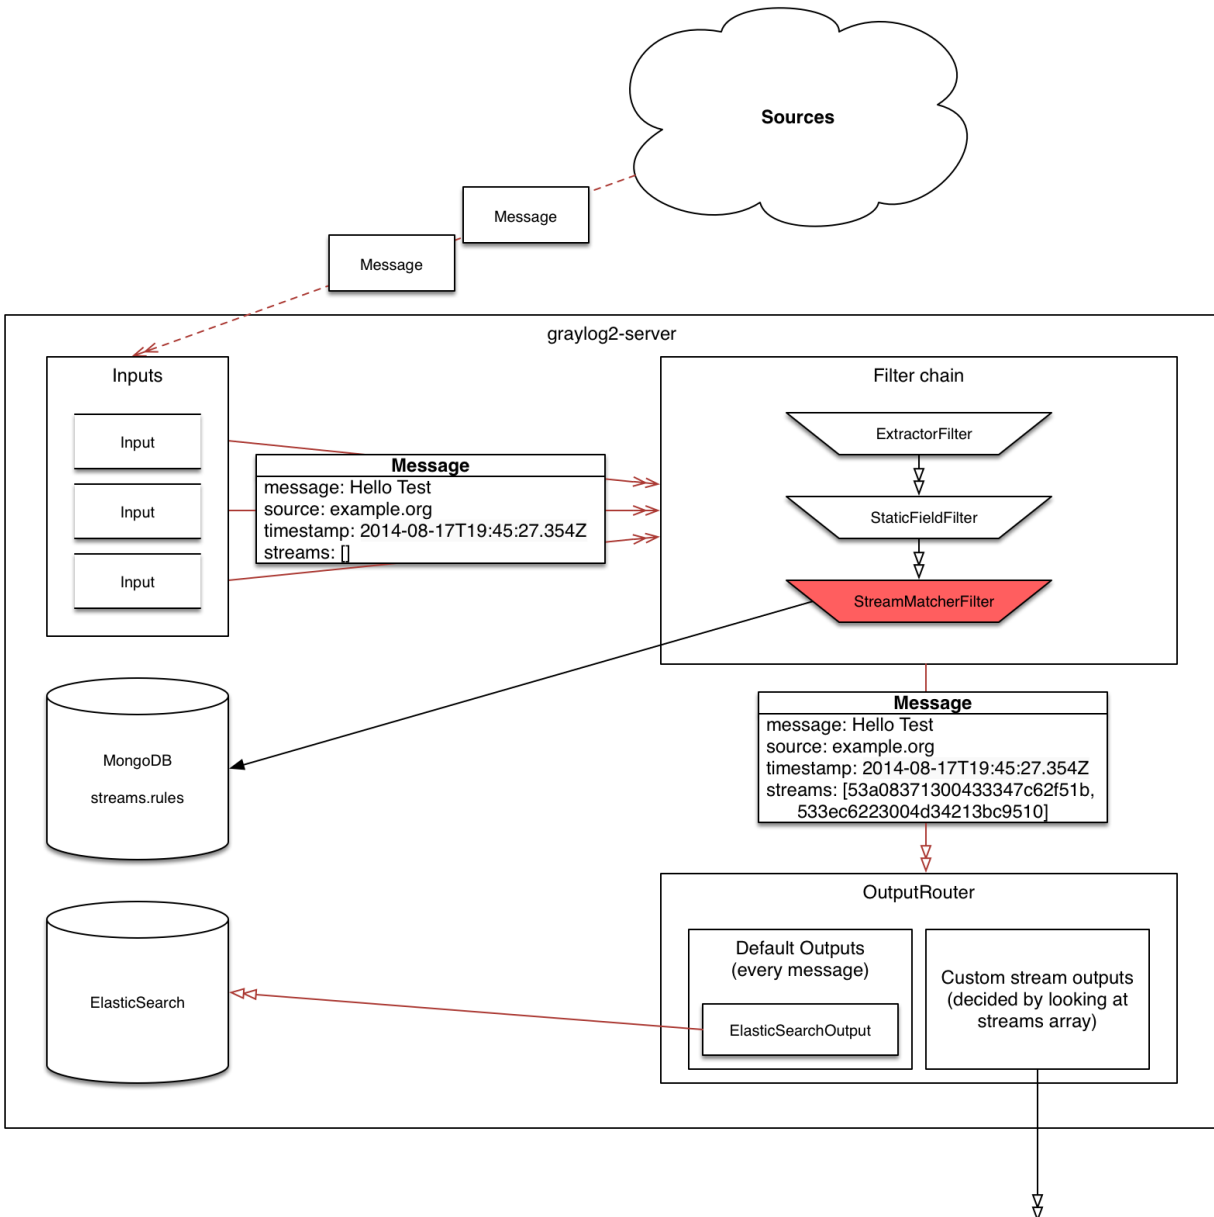
These are a few example use cases for streams:

- Forward a subset of messages to other data analysis or BI systems to reduce their license costs.
- Monitor exception or error rates in your whole environment and broken down per subsystem.
- Get a list of all failed SSH logins and use the *quickvalues* to analyze which user names were affected.
- Catch all HTTP POST requests to `/login` that were answered with a HTTP 302 and route them into a stream called *Successful user logins*. Now get a chart of when users logged in and use the *quickvalues* to get a list of users that performed the most logins in the search time frame.

How are streams processed internally?

The most important thing to know about Graylog stream matching is that there is no duplication of stored messages. Every message that comes in is matched against all rules of a stream. The internal ID of every stream that has *all* rules matching is appended to the `streams` array of the processed message.

All analysis methods and searches that are bound to streams can now easily narrow their operation by searching with a `streams: [STREAM_ID]` limit. This is done automatically by Graylog and does not have to be provided by the user.



Stream Processing Runtime Limits

An important step during the processing of a message is the stream classification. Every message is matched against the user-configured stream rules. If every rule of a stream matches, the message is added to this stream. Applying stream rules is done during the indexing of a message only, so the amount of time spent for the classification of a message is crucial for the overall performance and message throughput the system can handle.

There are certain scenarios when a stream rule takes very long to match. When this happens for a number of messages, message processing can stall, messages waiting for processing accumulate in memory and the whole system could become non-responsive. Messages are lost and manual intervention would be necessary. This is the worst case scenario.

To prevent this, the runtime of stream rule matching is limited. When it is taking longer than the configured runtime limit, the process of matching this exact message against the rules of this specific stream is aborted. Message processing in general and for this specific message continues though. As the runtime limit needs to be configured pretty high (usually a magnitude higher as a regular stream rule match takes), any excess of it is considered a fault and is recorded for this stream. If the number of recorded faults for a single stream is higher than a configured threshold, the stream rule set of this stream is considered faulty and the stream is disabled. This is done to protect the overall stability and performance of message processing. Obviously, this is a tradeoff and based on the assumption, that the total loss of one or more messages is worse than a loss of stream classification for these.

There are scenarios where this might not be applicable or even detrimental. If there is a high fluctuation of the message load including situations where the message load is much higher than the system can handle, overall stream matching can take longer than the configured timeout. If this happens repeatedly, all streams get disabled. This is a clear indicator that your system is overutilized and not able to handle the peak message load.

How to configure the timeout values if the defaults do not match

There are two configuration variables in the configuration file of the server, which influence the behavior of this functionality.

- `stream_processing_timeout` defines the maximum amount of time the rules of a stream are able to spend. When this is exceeded, stream rule matching for this stream is aborted and a fault is recorded. This setting is defined in milliseconds, the default is 2000 (2 seconds).
- `stream_processing_max_faults` is the maximum number of times a single stream can exceed this runtime limit. When it happens more often, the stream is disabled until it is manually reenabled. The default for this setting is 3.

What could cause it?

If a single stream has been disabled and all others are doing well, the chances are high that one or more stream rules are performing bad under certain circumstances. In most cases, this is related to stream rules which are utilizing regular expressions. For most other stream rules types the general runtime is constant, while it varies very much for regular expressions, influenced by the regular expression itself and the input matched against it. In some special cases, the difference between a match and a non-match of a regular expression can be in the order of 100 or even 1000. This is caused by a phenomenon called *catastrophic backtracking*. There are good write-ups about it on the web which will help you understanding it.

Summary: How do I solve it?

1. Check the rules of the stream that is disabled for rules that could take very long (especially regular expressions).
2. Modify or delete those stream rules.

3. Re-enable the stream.

Programmatic access via the REST API

Many organisations already run monitoring infrastructure that are able to alert operations staff when incidents are detected. These systems are often capable of either polling for information on a regular schedule or being pushed new alerts - this article describes how to use the Graylog Stream Alert API to poll for currently active alerts in order to further process them in third party products.

Checking for currently active alert/triggered conditions

Graylog stream alerts can currently be configured to send emails when one or more of the associated alert conditions evaluate to true. While sending email solves many immediate problems when it comes to alerting, it can be helpful to gain programmatic access to the currently active alerts.

Each stream which has alerts configured also has a list of active alerts, which can potentially be empty if there were no alerts so far. Using the stream's ID, one can check the current state of the alert conditions associated with the stream using the authenticated API call:

```
GET /streams/<streamid>/alerts/check
```

It returns a description of the configured conditions as well as a count of how many triggered the alert. This data can be used to for example send SNMP traps in other parts of the monitoring system.

Sample JSON return value:

```
{
  "total_triggered": 0,
  "results": [
    {
      "condition": {
        "id": "984d04d5-1791-4500-a17e-cd9621cc2ea7",
        "in_grace": false,
        "created_at": "2014-06-11T12:42:50.312Z",
        "parameters": {
          "field": "one_minute_rate",
          "grace": 1,
          "time": 1,
          "backlog": 0,
          "threshold_type": "lower",
          "type": "mean",
          "threshold": 1
        },
        "creator_user_id": "admin",
        "type": "field_value"
      },
      "triggered": false
    }
  ],
  "calculated_at": "2014-06-12T13:44:20.704Z"
}
```

Note that the result is cached for 30 seconds.

List of already triggered stream alerts

Checking the current state of a stream's alerts can be useful to trigger alarms in other monitoring systems, but if one wants to send more detailed messages to operations, it can be very helpful to get more information about the current state of the stream, for example the list of all triggered alerts since a certain timestamp.

This information is available per stream using the call:

```
GET /streams/<streamid>/alerts?since=1402460923
```

The since parameter is a unix timestamp value. Its return value could be:

```
{
  "total": 1,
  "alerts": [
    {
      "id": "539878473004e72240a5c829",
      "condition_id": "984d04d5-1791-4500-a17e-cd9621cc2ea7",
      "condition_parameters": {
        "field": "one_minute_rate",
        "grace": 1,
        "time": 1,
        "backlog": 0,
        "threshold_type": "lower",
        "type": "mean",
        "threshold": 1
      },
      "description": "Field one_minute_rate had a mean of 0.0 in the last 1 minutes_
↳with trigger condition lower than 1.0. (Current grace time: 1 minutes)",
      "triggered_at": "2014-06-11T15:39:51.780Z",
      "stream_id": "53984d8630042acb39c79f84"
    }
  ]
}
```

Using this information more detailed messages can be produced, since the response contains more detailed information about the nature of the alert, as well as the number of alerts triggered since the timestamp provided.

Note that currently a maximum of 300 alerts will be returned.

FAQs

Using regular expressions for stream matching

Stream rules support matching field values using regular expressions. Graylog uses the [Java Pattern class](#) to execute regular expressions.

For the individual elements of regular expression syntax, please refer to Oracle's documentation, however the syntax largely follows the familiar regular expression languages in widespread use today and will be familiar to most.

However, one key question that is often raised is matching a string in case insensitive manner. Java regular expressions are case sensitive by default. Certain flags, such as the one to ignore case sensitivity can either be set in the code, or as an inline flag in the regular expression.

To for example route every message that matches the browser name in the following user agent string:

```
Mozilla/5.0 (Macintosh; Intel Mac OS X 10_9_1) AppleWebKit/537.36 (KHTML, like Gecko)
↪Chrome/32.0.1700.107 Safari/537.36
```

the regular expression `.*applewebkit.*` will not match because it is case sensitive. In order to match the expression using any combination of upper- and lowercase characters use the `(?i)` flag as such:

```
(?i).*applewebkit.*
```

Most of the other flags supported by Java are rarely used in the context of matching stream rules or extractors, but if you need them their use is documented on the same Javadoc page by Oracle.

Can I add messages to a stream after they were processed and stored?

No. Currently there is no way to re-process or re-match messages into streams.

Only new messages are routed into the current set of streams.

Can I write own outputs or alert callbacks methods?

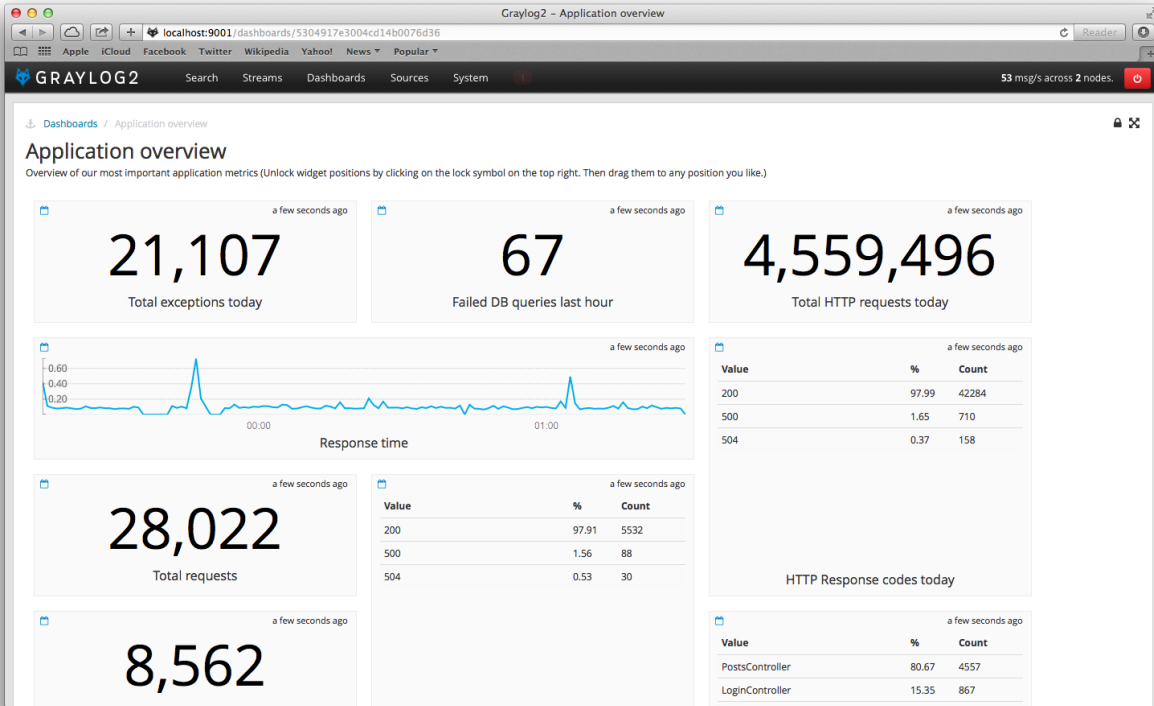
Yes. Please refer to the [plugins](#) documentation page.

Why dashboards matter

Using dashboards allows you to build pre-defined views on your data to always have everything important just one click away.

Sometimes it takes domain knowledge to be able to figure out the search queries to get the correct results for your specific applications. People with the required domain knowledge can define the search query once and then display the results on a dashboard to share them with co-workers, managers, or even sales and marketing departments.

This guide will take you through the process of creating dashboards and storing information on them. At the end you will have a dashboard with automatically updating information that you can share with anybody or just a subset of people based on permissions.




How to use dashboards

Creating an empty dashboard

Navigate to the *Dashboards* section using the link in the top menu bar of your Graylog web interface. The page is listing all dashboards that you are allowed to view. (More on permissions later.) Hit the *Create dashboard* button to create a new empty dashboard.

The only required information is a *title* and a *description* of the new dashboard. Use a specific but not too long title so people can easily see what to expect on the dashboard. The description can be a bit longer and could contain more detailed information about the displayed data or how it is collected.

Hit the *Create* button to create the dashboard. You should now see your new dashboard on the dashboards overview page. Click on the title of your new dashboard to see it. Next, we will be adding widgets to the dashboard we have just created.



Adding widgets

You should have your empty dashboard in front of you. Let's add some widgets! You can add search result information to dashboards with just one click. The following search result types can be added to dashboards:

- Search result counts
- Search result histogram charts
- Field value charts
- Quickvalue results

Once you can see the results of your search, you will see a small blue icon next to the right of the result count and histogram title. Hovering over this will show "Add to dashboard" and clicking the icon will prompt you with a list of dashboards you've created. Select a dashboard to add the widget to it.

Examples

It is strongly recommended to read the getting started guide on basic searches and analysis first. This will make the following examples more obvious for you.

- **Top log sources today**
 - Example search: *, timeframe: Last 24 hours
 - Expand the `source` field in the the sidebar and hit *Quick values*
 - Add quick values to dashboard

- **Number of exceptions in a given app today**
 - Example search: `source:myapp AND Exception`, timeframe: Last 24 hours
 - Add search result count to dashboard
- **Response time chart of a given app**
 - Example search: `source:myapp2`, any timeframe you want
 - Expand a field representing the response time of requests in the sidebar and hit *Generate chart*
 - Add chart to dashboard

Widgets from streams

You can of course also add widgets from stream search results. Every widget added this way will always be bound to streams. If you have a stream that contains every SSH login you can just search for everything (*) in that stream and store the result count as *SSH logins* on a dashboard.

Result

You should now see widgets on your dashboard. You will learn how to modify the dashboard, change cache times and widget positioning in the next chapter.

Modifying dashboards

You need to *unlock* dashboards to make any changes to them. Hit the lock icon in the top right corner of a dashboard to unlock it. You should now see new icons in the widget appearing.

Unlocked dashboard widgets explained

Unlocked dashboard widgets have three buttons that should be pretty self-explanatory.


- Delete widget
- Change cache time of widget
- Change title of widget

Widget cache times

Widget values are cached in `graylog-server` by default. **This means that the cost of value computation does not grow with every new device or even browser tab displaying a dashboard.** Some widgets might need to show real-time information (set cache time to 1 second) and some widgets might be updated way less often (like *Top SSH users this month*, cache time 10 minutes) to save expensive computation resources.

Repositioning widgets

Just grab a widget with your mouse in unlocked dashboard mode and move it around. Other widgets should adopt and re-position intelligently to make place for the widget you are moving. The positions are automatically saved when dropping a widget.



Dashboard permissions

Graylog users with administrator permissions are always allowed to view and edit all dashboards. Users with *reader* permissions are by default not allowed to view or edit **any** dashboard.

Dashboard Permissions

Application overview ✕

Test Dashboard #1 ✕

Choose dashboards the user can **view**.
Removing read access will remove edit access, too.

Application overview ✕

Choose dashboards the user can **edit**. Values chosen here will enable read access, too.

Navigate to *System* -> *Users* and select a *reader* user you wish to give dashboard permissions. Hit the *edit* button and assign dashboard *view* and *edit* permissions in the edit user dialogue. Don't forget to save the user!

That's it!

Congratulations, you have just gone through the basic principles of Graylog dashboards. Now think about which dashboards to create. We suggest:

- Create dashboards for yourself and your team members
- Create dashboards to share with your manager
- Create dashboards to share with the CIO of your company

Think about which information you need access to frequently. What information could your manager or CIO be interested in? Maybe they want to see how the number of exceptions went down or how your team utilized existing hardware better. The sales team could be interested to see signup rates in realtime and the marketing team will love you for providing insights into low level KPIs that is just a click away.

The problem explained

Syslog ([RFC3164](#), [RFC5424](#)) is the de facto standard logging protocol since the 1980s and was originally developed as part of the sendmail project. It comes with some annoying shortcomings that we tried to improve in [GELF](#) for application logging.

Because syslog has a clear specification in its RFCs it should be possible to parse it relatively easy. Unfortunately there are a lot of devices (especially routers and firewalls) out there that send logs looking like syslog but actually breaking several rules stated in the RFCs. We tried to write a parser that reads all of them as good as possible and failed. Such a loosely defined text message usually breaks the compatibility in the first date field already. Some devices leave out hostnames completely, some use localized timezone names (e. g. “MESZ” instead of “CEST”), and some just omit the current year in the timestamp field.

Then there are devices out there that at least do not claim to send syslog when they don’t but have another completely separate log format that needs to be parsed specifically.

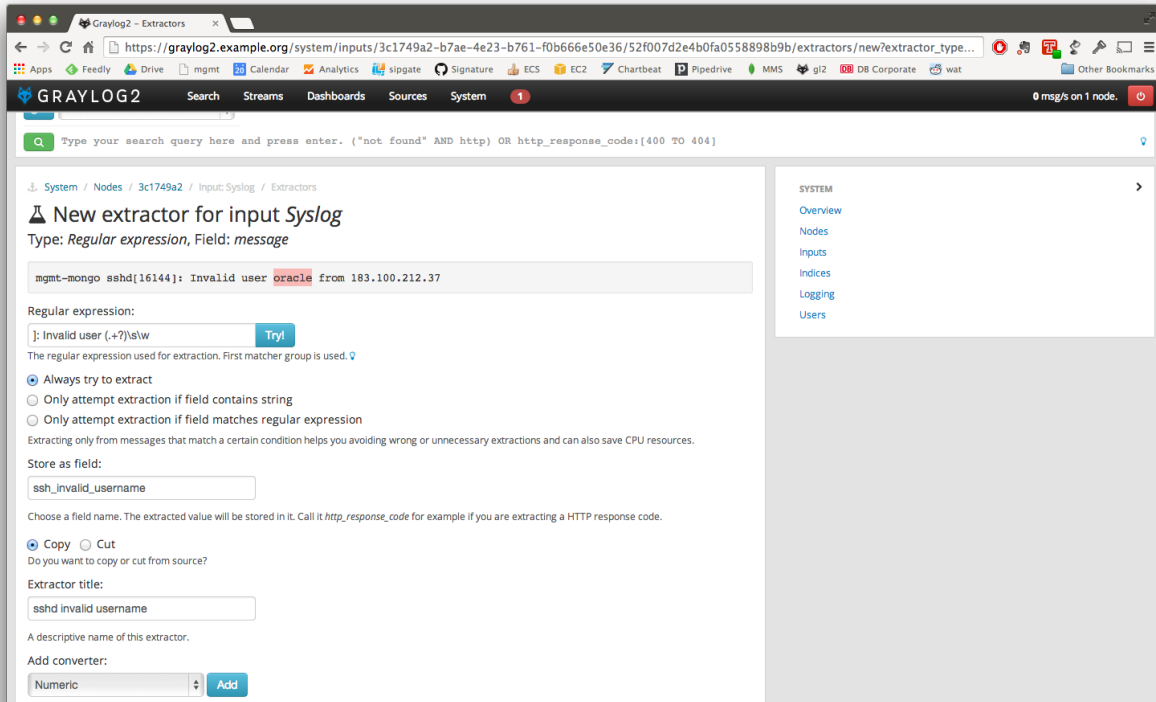
We decided not to write custom message inputs and parsers for all those thousands of devices, formats, firmwares and configuration parameters out there but came up with the concept of *Extractors* introduced the *v0.20.0* series of Graylog.

Graylog extractors explained

The extractors allow you to instruct Graylog nodes about how to extract data from any text in the received message (no matter from which format or if an already extracted field) to message fields. You may already know why structuring data into fields is important if you are using Graylog: There are a lot of analysis possibilities with full text searches but the real power of log analytics unveils when you can run queries like `http_response_code:>=500 AND user_id:9001` to get all internal server errors that were triggered by a specific user.

Wouldn’t it be nice to be able to search for all blocked packages of a given source IP or to get a quickterms analysis of recently failed SSH login usernames? Hard to do when all you have is just a single long text message.

Creating extractors is possible via either Graylog REST API calls or from the web interface using a wizard. Select a message input on the *System -> Inputs* page and hit *Manage extractors* in the actions menu. The wizard allows you to load a message to test your extractor configuration against. You can extract data using for example regular expressions, Grok patterns, substrings, or even by splitting the message into tokens by separator characters. The wizard looks like this and should be pretty intuitive:




You can also choose to apply so called *converters* on the extracted value to for example convert a string consisting of numbers to an integer or double value (important for range searches later), anonymize IP addresses, lower-/uppercase a string, build a hash value, and much more.

The extractor directory

The [data source library](#) provides access to a lot of extractors that you can easily import into your Graylog setup.

Just copy the JSON extractor export into the import dialog of a message input of the fitting type (every extractor set entry in the directory tells you what type of input to spawn, e. g. syslog, GELF, or Raw/plaintext) and you are good to go. The next messages coming in should already include the extracted fields with possibly converted values.

A message sent by Heroku and received by Graylog with the imported *Heroku* extractor set on a plaintext TCP input looks like this: (look at the extracted fields in the message detail view)



Using regular expressions to extract data

Extractors support matching field values using regular expressions. Graylog uses the `Java Pattern` class to evaluate regular expressions.

For the individual elements of regular expression syntax, please refer to Oracle's documentation, however the syntax largely follows the familiar regular expression languages in widespread use today and will be familiar to most.

However, one key question that is often raised is matching a string in case insensitive manner. Java regular expressions are case sensitive by default. Certain flags, such as the one to ignore case sensitivity can either be set in the code, or as an inline flag in the regular expression.

To for example create an extractor that matches the browser name in the following user agent string:

```
Mozilla/5.0 (Macintosh; Intel Mac OS X 10_9_1) AppleWebKit/537.36 (KHTML, like Gecko)
↪Chrome/32.0.1700.107 Safari/537.36
```

the regular expression `(applewebkit)` will not match because it is case sensitive. In order to match the expression using any combination of upper- and lowercase characters use the `(?i)` flag as such:

```
(?i)(applewebkit)
```

Most of the other flags supported by Java are rarely used in the context of matching stream rules or extractors, but if you need them their use is documented on the same Javadoc page by Oracle. One common reason to use regular expression flags in your regular expression is to make use of what is called non-capturing groups. Those are parentheses which only group alternatives, but do not make Graylog extract the data they match and are indicated by `(?:)`.

Using Grok patterns to extract data

Graylog also supports the extracting data using the popular Grok language to allow you to make use of your existing patterns.

Grok is a set of regular expressions that can be combined to more complex patterns, allowing to name different parts of the matched groups.

By using Grok patterns, you can extract multiple fields from a message field in a single extractor, which often simplifies specifying extractors.

Simple regular expressions are often sufficient to extract a single word or number from a log line, but if you know the entire structure of a line beforehand, for example for an access log, or the format of a firewall log, using Grok is advantageous.

For example a firewall log line could contain:

```
len=50824 src=172.17.22.108 sport=829 dst=192.168.70.66 dport=513
```

We can now create the following patterns on the [System/Grok Patterns](#) page in the web interface:

```
BASE10NUM (?<![0-9.+~]) (?>[+-]?(?:[0-9]+(?:\.[0-9]+)?)|(?\.[0-9]+))
NUMBER (?:%{BASE10NUM})
IPV6 ((([0-9A-Fa-f]{1,4}):){7}([0-9A-Fa-f]{1,4}|:)|((([0-9A-Fa-f]{1,4}):){6}(:[0-9A-Fa-f]{1,4}|(25[0-5]|2[0-4]\d|1\d\d|[1-9]?\d)\. (25[0-5]|2[0-4]\d|1\d\d|[1-9]?\d)) {3}|:))|((([0-9A-Fa-f]{1,4}):){5}(((:[0-9A-Fa-f]{1,4}){1,2})|:(25[0-5]|2[0-4]\d|1\d\d|[1-9]?\d)\. (25[0-5]|2[0-4]\d|1\d\d|[1-9]?\d)) {3}|:))|((([0-9A-Fa-f]{1,4}):){4}(((:[0-9A-Fa-f]{1,4}){1,3})|((:[0-9A-Fa-f]{1,4})?: (25[0-5]|2[0-4]\d|1\d\d|[1-9]?\d)\. (25[0-5]|2[0-4]\d|1\d\d|[1-9]?\d)) {3})|:))|((([0-9A-Fa-f]{1,4}):){3}(((:[0-9A-Fa-f]{1,4}){1,4})|((:[0-9A-Fa-f]{1,4}){0,2}:(25[0-5]|2[0-4]\d|1\d\d|[1-9]?\d)\. (25[0-5]|2[0-4]\d|1\d\d|[1-9]?\d)) {3})|:))|((([0-9A-Fa-f]{1,4}):){2}(((:[0-9A-Fa-f]{1,4}){1,5})|((:[0-9A-Fa-f]{1,4}){0,3}:(25[0-5]|2[0-4]\d|1\d\d|[1-9]?\d)\. (25[0-5]|2[0-4]\d|1\d\d|[1-9]?\d)) {3})|:))|((([0-9A-Fa-f]{1,4}):){1}(((:[0-9A-Fa-f]{1,4}){1,6})|((:[0-9A-Fa-f]{1,4}){0,4}:(25[0-5]|2[0-4]\d|1\d\d|[1-9]?\d)\. (25[0-5]|2[0-4]\d|1\d\d|[1-9]?\d)) {3})|:))|((:[0-9A-Fa-f]{1,4}){1,7})|((:[0-9A-Fa-f]{1,4}){0,5}:(25[0-5]|2[0-4]\d|1\d\d|[1-9]?\d)\. (25[0-5]|2[0-4]\d|1\d\d|[1-9]?\d)) {3}|:))|:)))(%.+)?
IPV4 (?<![0-9]) (?: (?:25[0-5]|2[0-4][0-9]|0[0-1]?[0-9]{1,2}) [.] (?:25[0-5]|2[0-4][0-9]|0[0-1]?[0-9]{1,2}) [.] (?:25[0-5]|2[0-4][0-9]|0[0-1]?[0-9]{1,2})) (?![0-9])
IP (?:%{IPV6}|%{IPV4})
```

Then, in the extractor configuration, we can use these patterns to extract the relevant fields from the line:

```
len=%{NUMBER:length} src=%{IP:srcip} sport=%{NUMBER:srcport} dst=%{IP:dstip} dport=%{NUMBER:dstport}
```

This will add the relevant extracted fields to our log message, allowing Graylog to search on those individual fields, which can lead to more effective search queries by allowing to specifically look for packets that came from a specific source IP instead of also matching destination IPs if one would only search for the IP across all fields.

There are many resources on the web with useful patterns, and one very helpful tool is the [Grok Debugger](#), which allows you to test your patterns while you develop them.

Graylog uses [Java Grok](#) to parse and run Grok patterns.

Normalization

Many log formats are similar to each other, but not quite the same. In particular they often only differ in the names attached to pieces of information.

For example, consider different hardware firewall vendors, whose models log the destination IP in different fields of the message, some use `dstip`, some `dst` and yet others use `destination-address`:

```
2004-10-13 10:37:17 PDT Packet Length=50824, Source address=172.17.22.108, Source_
↳port=829, Destination address=192.168.70.66, Destination port=513
2004-10-13 10:37:17 PDT len=50824 src=172.17.22.108 sport=829 dst=192.168.70.66_
↳dport=513
2004-10-13 10:37:17 PDT length="50824" srcip="172.17.22.108" srcport="829" dstip="192.
↳168.70.66" dstport="513"
```

You can use one or more non-capturing groups to specify the alternatives of the field names, but still be able to extract the a parentheses group in the regular expression. Remember that Graylog will extract data from the first matched group of the regular expression. An example of a regular expression matching the destination IP field of all those log messages from above is:

```
(?:dst|dstip|[dD]estination\saddress)="?(\d{1,3}\.\d{1,3}\.\d{1,3}\.\d{1,3})"?
```

This will only extract the IP address without caring about which of the three naming schemes was used in the original log message. This way you don't have to set up three different extractors.

The standard date converter

Date parser converters for extractors allow you to convert extracted data into timestamps - Usually used to set the timestamp of a message based on some date it contains. Let's assume we have this message from a network device:

```
<131>: foo-bar-dc3-org-de01: Mar 12 00:45:38: %LINK-3-UPDOWN: Interface_
↳GigabitEthernet0/31, changed state to down
```

Extracting most of the data is not a problem and can be done easily. Using the date in the message (*Mar 12 00:45:38*) as Graylog message timestamp however needs to be done with a date parser converter.

Use a standard extractor rule to select the timestamp and apply the *Date* converter with a format string:

```
MMM dd HH:mm:ss
```

(format string table at the end of this page)

Store as field:

Choose a field name. The extracted value will be stored in it. Call it *http_response_code* for example if you are extracting a HTTP response code.

Copy Cut


Do you want to copy or cut from source?

Extractor title:


A descriptive name of this extractor.


Add converter:

Convert to date type

Format string: 

Please note that you cannot use the cutting feature on standard fields like *message* and *source*.

✉ 4765e370-aa42-11e3-a7dd-4c8d79f2b596 

Received by  Cisco System Messages on [fb66b27e / 10.226.163.44](#)

Timestamp: 2014-03-12 00:45:38.000

Index: *graylog2_356*

Actions ▾

facility

local0

level

Error [3]

local_facility

link

local_level

3

message

Interface GigabitEthernet0/31, changed state to down

source

foo-bar-dc3-org-de01

type

updown

Standard date converter format string table

Symbol	Meaning	Presentation	Examples
G	era	text	AD
C	century of era (≥ 0)	number	20
Y	year of era (≥ 0)	year	1996
x	weekyear	year	1996
w	week of weekyear	number	27
e	day of week	number	2
E	day of week	text	Tuesday; Tue
y	year	year	1996
D	day of year	number	189
M	month of year	month	July; Jul; 07
d	day of month	number	10
a	halfday of day	text	PM
K	hour of halfday (0~11)	number	0
h	clockhour of halfday (1~12)	number	12
H	hour of day (0~23)	number	0
k	clockhour of day (1~24)	number	24
m	minute of hour	number	30
s	second of minute	number	55
S	fraction of second	millis	978
z	time zone	text	Pacific Standard Time; PST
Z	time zone offset/id	zone	-0800; -08:00; America/Los_Angeles
'	escape for text	delimiter	
'	single quote	literal	'

The flexible date converter

Now imagine you had one of those devices that send messages that are not so easy to parse because they do not follow a strict timestamp format. Some network devices for example like to send days of the month without adding a padding 0 for the first 9 days. You'll have dates like `Mar 9` and `Mar 10` and end up having problems defining a parser string for that. Or maybe you have something else that is really exotic like just `last wednesday` as timestamp. The flexible date converter is accepting any text data and tries to build a date from that as good as it can.

Examples:

- **Mar 12**, converted at 12:27:00 UTC in the year 2014: 2014-03-12T12:27:00.000
- **2014-3-12 12:27**: 2014-03-12T12:27:00.000
- **Mar 12 2pm**: 2014-03-12T14:00:00.000

Note that the flexible date converter always uses UTC as timezone unless you have timezone information in the parsed text.

Message rewriting with Drools

Graylog can optionally use [Drools Expert](#) to evaluate all incoming messages against a user defined rules file. Each message will be evaluated prior to being written to the outputs.

The rule file location is defined in the Graylog configuration file:

```
# Drools Rule File (Use to rewrite incoming log messages)
rules_file = /etc/graylog.d/rules/graylog.drl
```

The rules file is located on the file system with a `.drl` file extension. The rules file can contain multiple rules, queries and functions, as well as some resource declarations like imports, globals, and attributes that are assigned and used by your rules and queries.

For more information on the DRL rules syntax please read the [Drools User Guide](#).

Getting Started

1. Uncomment the `rules_file` line in the Graylog configuration file.
2. Copy the [sample rules file](#) to the location specified in your Graylog configuration file.
3. Modify the rules file to parse/rewrite/filter messages as needed.

Example rules file

This is an example rules file:

```
import org.graylog2.plugin.Message

rule "Rewrite localhost host"
  when
    m : Message( source == "localhost" )
  then
```

```

        m.addField("source", "localhost.example.com" );
        System.out.println( "[Overwrite localhost rule fired] : " + m.toString() );
end

rule "Drop UDP and ICMP Traffic from firewall"
  when
    m : Message( getField("full_message") matches "(?i).*(ICMP|UDP) Packet (.
↪|\n|\r)*" && source == "firewall" )
  then
    m.setFilterOut(true);
    System.out.println("[Drop all syslog ICMP and UDP traffic] : " + m.toString()
↪);
end

```

Parsing Message and adding fields

In the following script we turn the PID and the src IP into additional fields:

```

import org.graylog2.plugin.Message
import java.util.regex.Matcher
import java.util.regex.Pattern

// Raw Syslog Apr 18 15:34:58 server01 smtp-glass[3371]: NEW (1/0) on=1.1.1.1:9100,
↪src=2.2.2.2:38776, ident=, dst=3.3.3.3:25, id=1303151698.3371
rule "SMTP Glass Logging to GELF"
  when
    m : Message( message matches "^smtp-glass.*" )
  then
    Matcher matcher = Pattern.compile("smtp-glass\\[(\\d+)\\].* src (\\d+\\.\\d+\\.
↪\\d+\\.\\d+)").matcher(m.getMessage());
    if (matcher.find()) {
      m.addField("_pid", Long.valueOf(matcher.group(1)));
      m.addField("_src", matcher.group(2));
    }
  }
end

```

Another example: Adding additional fields and changing the message itself

We send Squid access logs to Graylog using Syslog. The problem is that the *host* field of the message was set to the IP address of the Squid proxy, which not very useful. This rule overwrites the source and adds other fields:

```

import org.graylog2.plugin.Message
import java.util.regex.Matcher
import java.util.regex.Pattern
import java.net.InetAddress;

/*
Raw Syslog: squid[2099]: 1339551529.881 55647 1.2.3.4 TCP_MISS/200 22 GET http://www.
↪google.com/

squid\[\\d+\\]: (\\d+\\.\\d+\\.\\d+) *(\\d+) *(\\d+\\.\\d+\\.\\d+\\.\\d+) *(\\w+\\/\\w+) (\\d+) (\\w+) (.*
matched: 13:1339551529.881
matched: 29:55647
matched: 35:1.2.3.4

```

```
matched: 47:TCP_MISS/200
matched: 60:22
matched: 64:GET
matched: 68:http://www.google.com/
*/

rule "Squid Logging to GELF"
  when
    m : Message( getField("facility") == "local5" )
  then
    Matcher matcher = Pattern.compile("squid\\[[\\d+\\]: (\\d+\\.\\d+) *(\\d+)␣
↪*(\\d+\\.\\d+\\.\\d+\\.\\d+) *(\\w+\\//\\w+) (\\d+) (\\w+) (.*)").matcher(m.getMessage());

    if (matcher.find()) {
      m.addField("facility", "squid");
      InetAddress addr = InetAddress.getByName(matcher.group(3));
      String host = addr.getHostName();
      m.addField("source", host);
      m.addField("message", matcher.group(6) + " " + matcher.group(7));
      m.addField("_status", matcher.group(4));
      m.addField("_size", matcher.group(5));
    }
  end
end
```

Load balancer integration

When running multiple Graylog servers a common deployment scenario is to route the message traffic through an IP load balancer. By doing this we can achieve both a highly available setup, as well as increasing message processing throughput, by simply adding more servers that operate in parallel.

Load balancer state

However, load balancers usually need some way of determining whether a backend service is reachable and healthy or not. For this purpose Graylog exposes a load balancer state that is reachable via its REST API.

There are two ways the load balancer state can change:

- due to a lifecycle change (e.g. the server is starting to accept messages, or shutting down)
- due to manual intervention via the REST API

To query the current load balancer status of a Graylog instance, all you need to do is to issue a HTTP call to its REST API:

```
GET /system/lbstatus
```

The status knows two different states, `ALIVE` and `DEAD`, which is also the `text/plain` response of the resource. Additionally, the same information is reflected in the HTTP status codes: If the state is `ALIVE` the return code will be `200 OK`, for `DEAD` it will be `503 Service unavailable`. This is done to make it easier to configure a wide range of load balancer types and vendors to be able to react to the status.

The resource is accessible without authentication to make it easier for load balancers to access it.

To programmatically change the load balancer status, an additional endpoint is exposed:

```
PUT /system/lbstatus/override/alive
PUT /system/lbstatus/override/dead
```

Only authenticated and authorized users are able to change the status, in the currently released Graylog version this means only admin users can change it.

Graceful shutdown

Often, when running a service behind a load balancer, the goal is to be able to perform zero-downtime upgrades, by taking one of the servers offline, upgrading it, and then bringing it back online. During that time the remaining servers can take the load seamlessly.

By using the load balancer status API described above one can already perform such a task. However, it would still be guesswork when the Graylog server is done processing all the messages it already accepted.

For this purpose Graylog supports a graceful shutdown command, also accessible via the web interface and API. It will set the load balancer status to `DEAD`, stop all inputs, turn on messages processing (should it have been disabled manually previously), and flush all messages in memory to Elasticsearch. After all buffers and caches are processed, it will shut itself down safely.

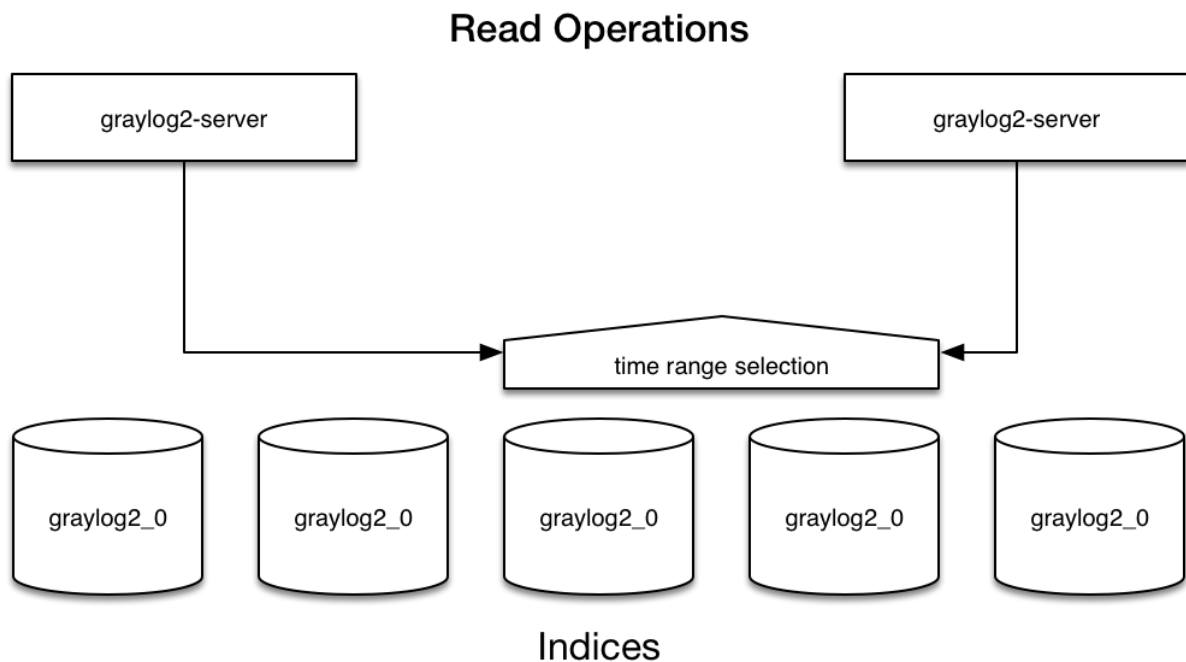
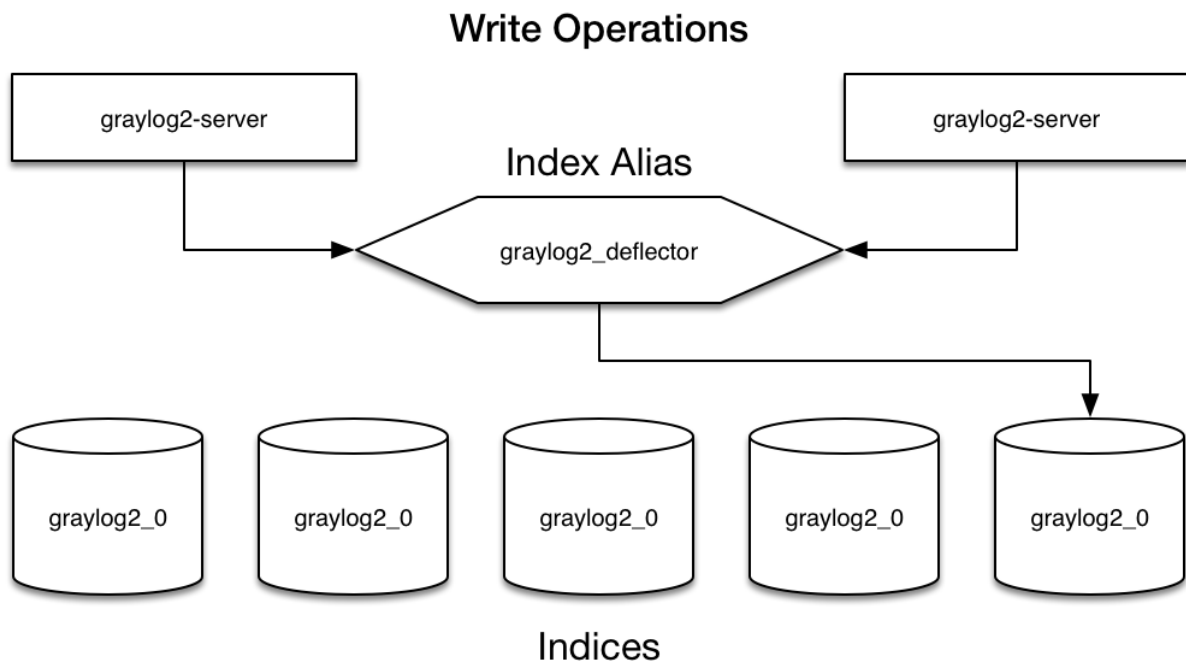
The Graylog index model explained

Overview

Graylog is transparently managing a set of indices to optimise search and analysis operations for speed and low resource utilisation. The system is maintaining an index alias called *graylog_deflector* that is always pointing to the current write-active index. We always have exactly one index to which new messages are appended until the configured maximum size (`elasticsearch_max_docs_per_index` in your `graylog.conf`) is reached.

A background task is running every minute and checks if the maximum size is reached. A new index is created and prepared when that happens. Once the index is considered to be ready to be written to, the `graylog_deflector` is atomically switched to the it. That means that all writing nodes can always write to the deflector alias without even knowing what the currently active write-active index is.

Note that there are also time based retention settings since v1.0 of Graylog. This allows you to instruct Graylog to keep messages based on their age and not the total amount. You can find the corresponding configuration settings in your `graylog.conf`.



Almost every read operation is performed with a given time range. Because Graylog is only writing sequentially it can keep a cached collection of information about which index starts at what point in time. It selects a lists of indices to query when having a time range provided. If no time range is provided it will search in all indices it knows.

Eviction of indices and messages

You have configured the maximum number of indices in your `graylog.conf` (`elasticsearch_max_number_of_indices`). When that number is reached the oldest indices will automatically be deleted. The deleting is performed by the *graylog-server* master node in a background process that is continuously comparing the actual number of indices with the configured maximum:

```
elasticsearch_max_docs_per_index * elasticsearch_max_number_of_indices
= maximum number of messages stored
```

Keeping the metadata in synchronisation

Graylog will notify you when the stored metadata about index time ranges has run out of sync. This can for example happen when you delete indices by hand. The system will offer you to just re-generate all time range information. This may take a few seconds but is an easy task for Graylog.

You can easily re-build the information yourself after manually deleting indices or doing other changes that might cause synchronisation problems:

```
$ curl -XPOST http://127.0.0.1:12900/system/indices/ranges/rebuild
```

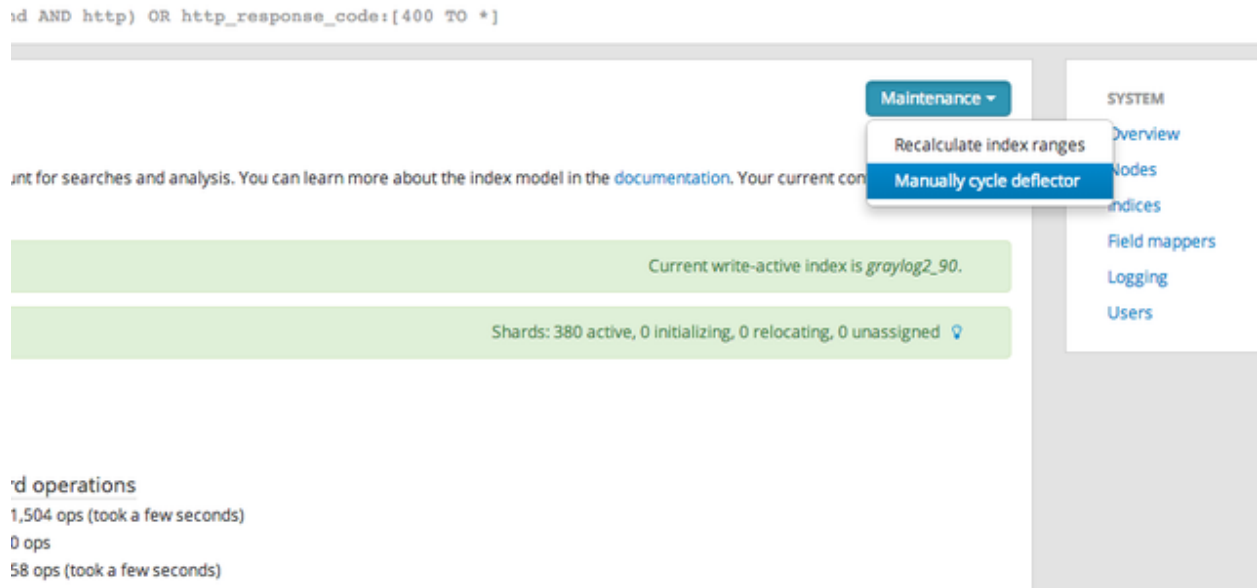
This will trigger a systemjob:

```
INFO : org.graylog2.system.jobs.SystemJobManager - Submitted SystemJob <ef7057c0-5ae3-
↳11e3-b935-4c8d79f2b596> [org.graylog2.indexer.ranges.RebuildIndexRangesJob]
INFO : org.graylog2.indexer.ranges.RebuildIndexRangesJob - Re-calculating index_
↳ranges.
INFO : org.graylog2.indexer.ranges.RebuildIndexRangesJob - Calculated range of_
↳[graylog2_56] in [640ms].
INFO : org.graylog2.indexer.ranges.RebuildIndexRangesJob - Calculated range of_
↳[graylog2_18] in [66ms].
...
INFO : org.graylog2.indexer.ranges.RebuildIndexRangesJob - Done calculating index_
↳ranges for 88 indices. Took 4744ms.
INFO : org.graylog2.system.jobs.SystemJobManager - SystemJob <ef7057c0-5ae3-11e3-b935-
↳4c8d79f2b596> [org.graylog2.indexer.ranges.RebuildIndexRangesJob] finished in_
↳4758ms.
```

Manually cycling the deflector

Sometimes you might want to cycle the deflector manually and not wait until the configured maximum number of messages in the newest index is reached. You can do this either via a REST call against the *graylog-server* master node or via the web interface:

```
$ curl -XPOST http://127.0.0.1:12900/system/deflector/cycle
```



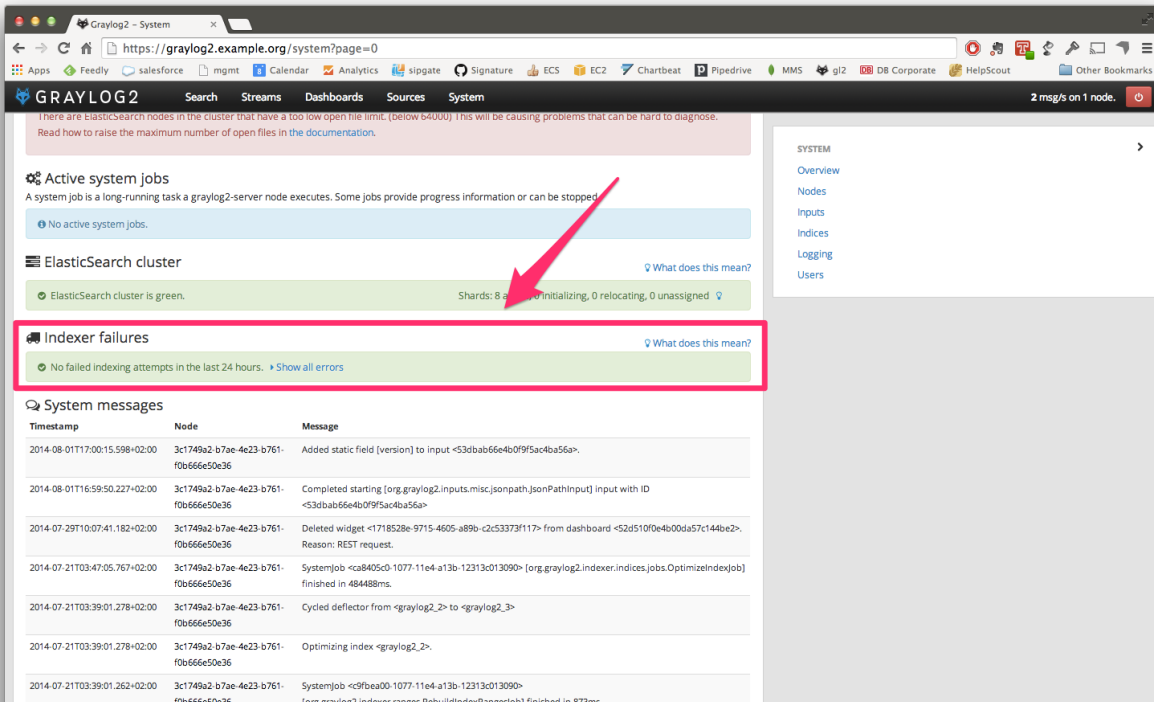
This triggers the following log output:

```
INFO : org.graylog2.rest.resources.system.DeflectorResource - Cycling deflector.
↳Reason: REST request.
INFO : org.graylog2.indexer.Deflector - Cycling deflector to next index now.
INFO : org.graylog2.indexer.Deflector - Cycling from <graylog2_90> to <graylog2_91>
INFO : org.graylog2.indexer.Deflector - Creating index target <graylog2_91>...
INFO : org.graylog2.indexer.Deflector - Done!
INFO : org.graylog2.indexer.Deflector - Pointing deflector to new target index....
INFO : org.graylog2.indexer.Deflector - Flushing old index <graylog2_90>.
INFO : org.graylog2.indexer.Deflector - Setting old index <graylog2_90> to read-only.
INFO : org.graylog2.system.jobs.SystemJobManager - Submitted SystemJob <a05e0d60-5c34-
↳11e3-8df7-4c8d79f2b596> [org.graylog2.indexer.indices.jobs.OptimizeIndexJob]
INFO : org.graylog2.indexer.Deflector - Done!
INFO : org.graylog2.indexer.indices.jobs.OptimizeIndexJob - Optimizing index
↳<graylog2_90>.
INFO : org.graylog2.system.jobs.SystemJobManager - SystemJob <a05e0d60-5c34-11e3-8df7-
↳4c8d79f2b596> [org.graylog2.indexer.indices.jobs.OptimizeIndexJob] finished in
↳334ms.
```

Indexer failures and dead letters

Indexer failures

Every `graylog-server` instance is constantly keeping track about every indexing operation it performs. This is important for making sure that you are not silently losing any messages. The web interface can show you a number of write operations that failed and also a list of failed operations. Like any other information in the web interface this is also available via the REST APIs so you can hook it into your own monitoring systems.



The screenshot shows the Graylog2 System page. The URL is `https://graylog2.example.org/system?page=0`. The page has a navigation bar with links for Search, Streams, Dashboards, Sources, and System. A status bar at the top right indicates "2 msg/s on 1 node".

The main content area is divided into several sections:

- Active system jobs:** A blue box with a gear icon and the text "No active system jobs."
- ElasticSearch cluster:** A green box with a cluster icon and the text "ElasticSearch cluster is green. Shards: 8 available, 0 initializing, 0 relocating, 0 unassigned". A red arrow points to the "Shards" information.
- Indexer failures:** A green box with a warning icon and the text "No failed indexing attempts in the last 24 hours. Show all errors". This section is highlighted with a red box.
- System messages:** A table with columns for Timestamp, Node, and Message.

The System messages table contains the following data:

Timestamp	Node	Message
2014-08-01T17:00:15.598+02:00	3c1749a2-b7ae-4623-b761-f0b666e50e36	Added static field [version] to input <53dbab66e4b09f5ac4ba56a>.
2014-08-01T16:59:50.227+02:00	3c1749a2-b7ae-4623-b761-f0b666e50e36	Completed starting [org.graylog2.inputs.misc.jsonpath.jsonPathInput] input with ID <53dbab66e4b09f5ac4ba56a>
2014-07-29T10:07:41.182+02:00	3c1749a2-b7ae-4623-b761-f0b666e50e36	Deleted widget <1718528e-9715-4605-a89b-c2c53373f117> from dashboard <52d510f0e4b00da57c144be2>. Reason: REST request.
2014-07-21T03:47:05.767+02:00	3c1749a2-b7ae-4623-b761-f0b666e50e36	Systemjob <ca8405cd-1077-11e4-a13b-12313c013090> [org.graylog2.indexer.indices.jobs.OptimizeIndexJob] finished in 484488ms.
2014-07-21T03:39:01.278+02:00	3c1749a2-b7ae-4623-b761-f0b666e50e36	Cycled deflector from <graylog2_2> to <graylog2_3>
2014-07-21T03:39:01.278+02:00	3c1749a2-b7ae-4623-b761-f0b666e50e36	Optimizing index <graylog2_2>.
2014-07-21T03:39:01.262+02:00	3c1749a2-b7ae-4623-b761-f0b666e50e36	Systemjob <c9fbae00-1077-11e4-a13b-12313c013090> [org.graylog2.indexer.ranges.RebuildIndexRangesJob] finished in 873ms.

Information about the indexing failure is stored in a capped MongoDB collection that is limited in size. A lot (many tens of thousands) of failure messages should fit in there but it should not be considered a complete collection of all errors ever thrown.

Dead letters

This is an experimental feature. You can enable the dead letters feature in your `graylog-server.conf` like this:

```
dead_letters_enabled = true
```

Graylog will write every message that could not be written to Elasticsearch into the MongoDB `dead_letters` collection. The messages will be waiting there for you to be processed in some other way. You could write a script that reads every message from there and transforms it in a way that will allow Graylog to accept it.

A dead letter in MongoDB has exactly the structure (in the `message` field) like the message that would have been written to the indices:

```
$ mongo
MongoDB shell version: 2.4.1
connecting to: test
> use graylog2
switched to db graylog2
> db.dead_letters.find().limit(1).pretty()
{
  "_id" : ObjectId("530a951b3004ada55961ee22"),
  "message" : {
    "timestamp" : "2014-02-24 00:40:59.121",
    "message" : "failing",
    "failure" : "haha",
    "level" : NumberLong(6),
    "_id" : "544575a0-9cec-11e3-b502-4c8d79f2b596",
    "facility" : "gelf-rb",
    "source" : "sundaysister",
    "gl2_source_input" : "52ef64d03004faafd4bb0fc2",
    "gl2_source_node" : "fb66b27e-993c-4595-940f-dd521dcdaa93",
    "file" : "(irb)",
    "line" : NumberLong(37),
    "streams" : [ ],
    "version" : "1.0"
  },
  "timestamp" : ISODate("2014-02-24T00:40:59.137Z"),
  "letter_id" : "54466000-9cec-11e3-b502-4c8d79f2b596"
}
```

The `timestamp` is the moment in time when the message could not be written to the indices and the `letter_id` references to the failed indexing attempt and its error message.

Every failed indexing attempt comes with a field called `written` that indicates if a dead letter was created or not:

```
> db.index_failures.find().limit(1).pretty()
{
  "_id" : ObjectId("530a951b3004ada55961ee23"),
  "timestamp" : ISODate("2014-02-24T00:40:59.136Z"),
  "message" : "MapperParsingException[failed to parse [failure]]; nested: ↵
↵NumberFormatException[For input string: \"haha\"]; ",
  "index" : "graylog2_324",
```

```
"written" : true,  
"letter_id" : "54466000-9cec-11e3-b502-4c8d79f2b596",  
"type" : "message"  
}
```

Common indexer failure reasons

There are some common failures that can occur under certain circumstances. Those are explained here:

MapperParsingException

An error message would look like this:

```
MapperParsingException[failed to parse [failure]]; nested: NumberFormatException[For  
↪input string: "some string value"];
```

You tried to write a `string` into a numeric field of the index. The indexer tried to convert it to a number, but failed because the `string` did contain characters that could not be converted.

This can be triggered by for example sending GELF messages with different field types or extractors trying to write strings without converting them to numeric values first. **The recommended solution is to actively decide on field types.** If you sent in a field like `http_response_code` with a numeric value then you should never change that type in the future.

The same can happen with all other field types like for example booleans.

Note that index cycling is something to keep in mind here. The first type written to a field per index wins. If the Graylog index cycles then the field types are starting from scratch for that index. If the first message written to that index has the `http_response_code` set as `string` then it will be a `string` until the index cycles the next time. Take a look at *The Graylog index model explained* for more information.

General information

Graylog comes with a stable plugin API for the following plugin types since Graylog 1.0:

- **Inputs:** Accept/write any messages into Graylog
- **Outputs:** Forward messages to other endpoints in real-time
- **Services:** Run at startup and able to implement any functionality
- **Alarm Callbacks:** Called when a stream alert condition has been triggered
- **Filters:** Transform/drop incoming messages during processing
- **REST API Resources:** A REST resource to expose as part of the `graylog-server` REST API
- **Periodical:** Called at periodical intervals during server runtime

The first for writing a plugin is creating a skeleton that is the same for each type of plugin. The next chapter is explaining how to do this and will then go over to chapters explaining plugin types in detail.

Creating a plugin skeleton

The easiest way to get started is to use our [maven archetype](#) that will create a complete plugin project infrastructure with all required classes, build definitions, and configurations using an interactive wizard.

Maven is a Java widely used build tool that comes pre-installed on many operating systems or can be installed using most package managers. Make sure that it is installed with at least version 3 before you go on.

Use it like this:

```
$ mvn archetype:generate -DarchetypeGroupId=org.graylog -DarchetypeArtifactId=graylog-  
→plugin-archetype
```

It will ask you a few questions about the plugin you are planning to build. Let's say you work for a company called ACMECorp and want to build an alarm callback plugin that creates a JIRA ticket for each alarm that is triggered:

```
groupId: com.acmecorp
artifactId: jira-alarmcallback
version: 1.0.0-SNAPSHOT
package: com.acmecorp
pluginClassName: JiraAlarmCallback
```

Note that you do not have to tell the archetype wizard what kind of plugin you want to build because it is creating the generic plugin skeleton for you but nothing that is related to the actual implementation. More on this in the example plugin chapters later.

You now have a new folder called `jira-alarmcallback` that includes a complete plugin skeleton including Maven build files. Every Java IDE out there can now import the project automatically without any required further configuration.

In [IntelliJ IDEA](#) for example you can just use the *File -> Open* dialog to open the skeleton as a fully configured Java project.

Change some default values

Open the `JiraAlarmCallbackMetaData.java` file and customize the default values like the plugin description, the website URI, and so on. Especially the author name etc. should be changed.

Now go on with implementing the actual login in one of the example plugin chapters below.

Example Alarm Callback plugin

Let's assume you still want to build the mentioned JIRA AlarmCallback plugin. First open the `JiraAlarmCallback.java` file and let it implement the `AlarmCallback` interface:

```
public class JiraAlarmCallback implements AlarmCallback
```

Your IDE should offer you to create the methods you need to implement:

public void initialize(Configuration configuration) throws AlarmCallbackConfigurationException

This is called once at the very beginning of the lifecycle of this plugin. It is common practice to store the `Configuration` as a private member for later access.

public void call(Stream stream, AlertCondition.CheckResult checkResult) throws AlarmCallbackException

This is the actual alarm callback being triggered. Implement your login that creates a JIRA ticket here.

public ConfigurationRequest getRequestedConfiguration()

Plugins can request configurations. The UI in the Graylog web interface is generated from this information and the filled out configuration values are passed back to the plugin in `initialize(Configuration configuration)`.

This is an example configuration request:

```
final ConfigurationRequest configurationRequest = new ConfigurationRequest();
configurationRequest.addField(new TextField(
    "service_key", "Service key", "", "JIRA API token. You can find this token in_
↪your account settings.",
    ConfigurationField.Optional.NOT_OPTIONAL)); // required, must be filled out
```

```
configurationRequest.addField(new BooleanField(
    "use_https", "HTTPs", true,
    "Use HTTP for API communication?"));
```

public String getName()

Return a human readable name of this plugin.

public Map<String, Object> getAttributes()

Return attributes that might be interesting to be shown under the alarm callback in the Graylog web interface. It is common practice to at least return the used configuration here.

public void checkConfiguration() throws ConfigurationException

Throw a `ConfigurationException` if the user should have entered missing or invalid configuration parameters.

Registering the plugin

You now have to register your plugin in the `JiraAlarmCallbackModule.java` file to make `graylog-server` load the alarm callback when launching. The reason for the manual registering is that a plugin could consist of multiple plugin types. Think of the generated plugin file as a bundle of multiple plugins.

Register your new plugin using the `configure()` method:

```
@Override
protected void configure() {
    addAlarmCallback(JiraAlarmCallback.class);
}
```

Building plugins

Building the plugin is easy because the archetype has created all necessary files and settings for you. Just run `mvn package` from the plugin directory:

```
$ mvn package
```

This will generate a `.jar` file in `target/` that is the complete plugin file:

```
$ ls target/jira-alarmcallback-1.0.0-SNAPSHOT.jar
target/jira-alarmcallback-1.0.0-SNAPSHOT.jar
```

Installing and loading plugins

The only thing you need to do to run the plugin in Graylog is to copy the `.jar` file to your plugins folder that is configured in your `graylog.conf`. The default is just `plugins/` relative from your `graylog-server` directory.

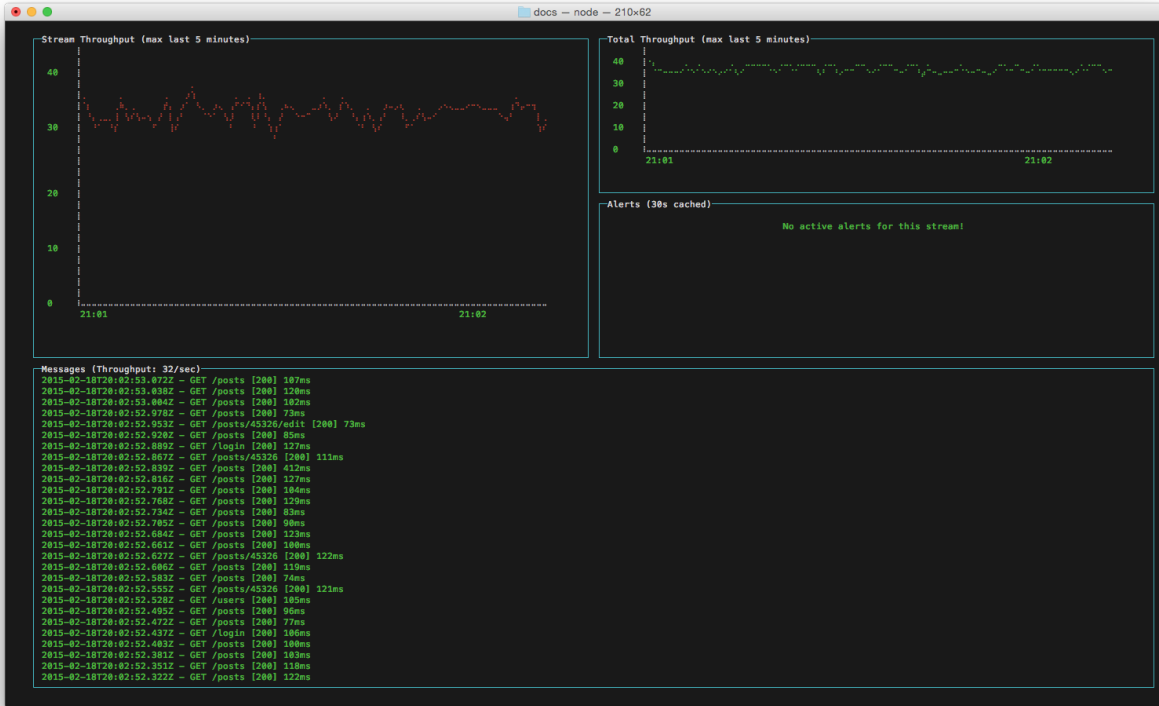
Restart `graylog-server` and the plugin should be available to use from the web interface immediately.

External dashboards

There are other frontends that are connecting to the Graylog REST API and display data or information in a special way.

CLI stream dashboard

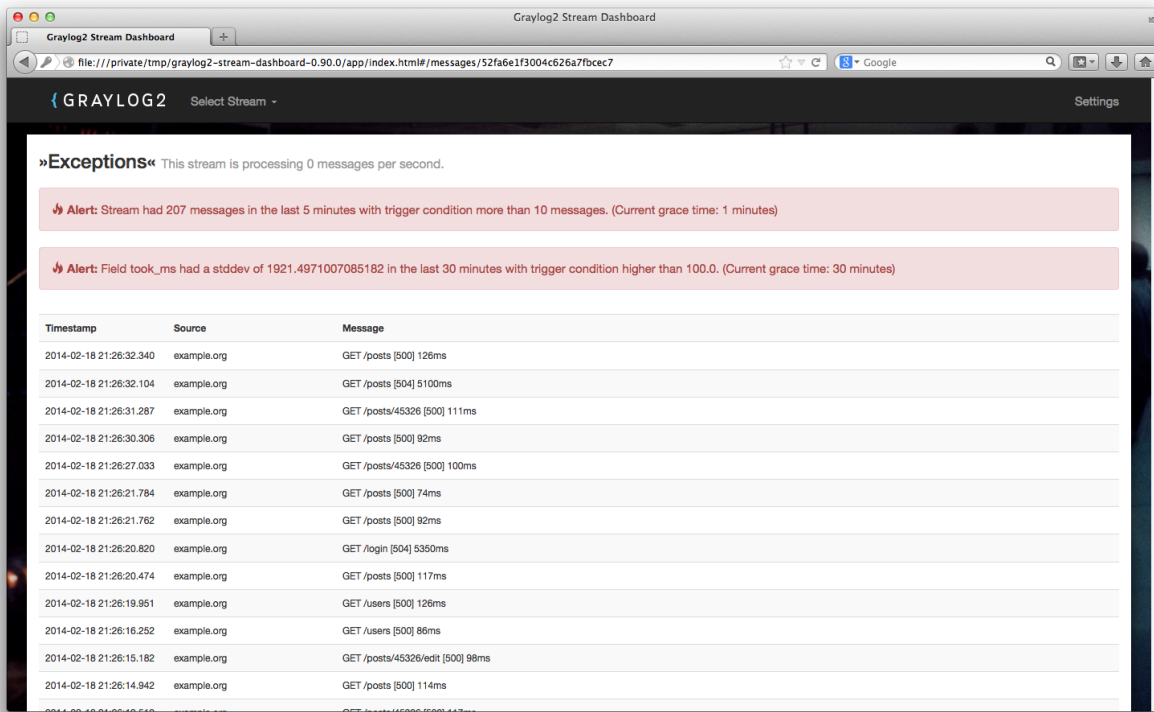
This official Graylog dashboard which is developed by us is showing live information of a specific stream in your terminal. For example it is the perfect companion during a deployment of your platform: Run it next to the deployment output and show information of a stream that is catching all errors or exceptions on your systems.



The CLI stream dashboard documentation is available on [GitHub](#).

Browser stream dashboard

This official Graylog dashboard is showing live information of a specific stream in a web browser. It will display and automatically reload the most recent messages and alerts of a stream and is perfect to display on large screens in your office.



The browser stream dashboard documentation is available on [GitHub](#).

The thinking behind the Graylog architecture and why it matters to you

A short history of Graylog

The Graylog project was started by Lennart Koopmann some time around 2009. Back then the most prominent log management software vendor issued a quote for a one year license of their product that was so expensive that he decided to write a log management system himself. Now you might call this a bit over optimistic (*I'll build this in two weeks*, end of quote) but the situation was hopeless: There was basically no other product on the market and especially no open source alternatives.

The log management market today

Things have changed a bit since 2009. Now there are viable open source projects with serious products and a growing list of SaaS offerings for log management.

Architectural considerations

Graylog has been successful in providing log management software **because it was built for log management from the beginning**. Software that stores and analyzes log data must have a very specific architecture to do it efficiently. It is more than just a database or a full text search engine because it has to deal with both text data and metrics data on a time axis. Searches are always bound to a time frame (relative or absolute) and only going back into the past because future log data has not been written yet. **A general purpose database or full text search engine that could also store and index the private messages of your online platform for search will never be able to effectively manage your log data**. Adding a specialized frontend on top of it makes it look like it could do the job in a good way but is basically just putting lipstick on the wrong stack.

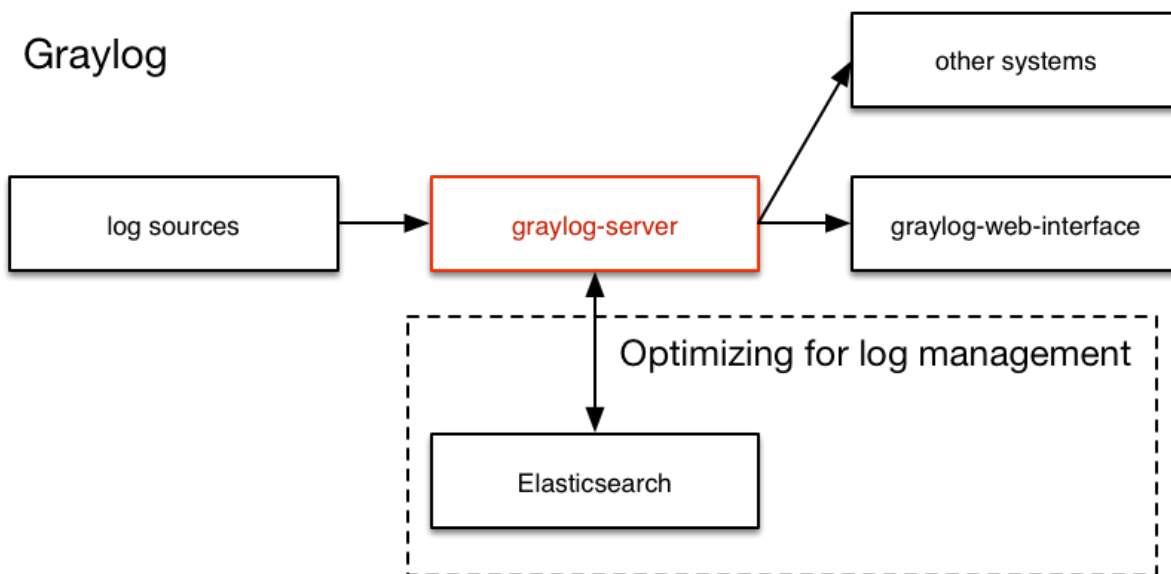
A log management system has to be constructed of several services that take care of processing, indexing, and data access. The most important reason is that you need to scale parts of it horizontally with your changing use cases and usually the different parts of the system have different hardware requirements. All services must be tightly integrated to allow efficient management and configuration of the system as a whole. A data ingestion or forwarder tool is hard to tedious to manage if the configuration **has** to be stored on the client machines and is not possible via for example

REST APIs controlled by a simple interface. A system administrator needs to be able to log into the web interface of a log management product and select log files of a remote host (that has a forwarder running) for ingestion into the tool.

You also want to be able to see the health and configuration of all forwarders, data processors and indexers in a central place because the whole log management stack can easily involve thousands of machines if you include the log emitting clients into this calculation. You need to be able to see which clients are forwarding log data and which are not to make sure that you are not missing any important data.

Graylog is coming the closest to the Splunk architecture:

- **Graylog was solely built as a log management system from the first line of code.** This makes it very efficient and easy to use.
- The `graylog-server` component sits in the middle and works around shortcomings of Elasticsearch (a full text search engine, not a log management system) for log management. It also builds an abstraction layer on top of it to make data access as easy as possible without having to select indices and write tedious time range selection filters, etc. - Just submit the search query and Graylog will take care of the rest for you.
- All parts of the system are tightly integrated and many parts speak to each other to make your job easier.
- Like Wordpress makes MySQL a good solution for blogging, Graylog makes Elasticsearch a good solution for logging. You should never have a system or frontend query Elasticsearch directly for log management so we are putting `graylog-server` in front of it.



Unlimited data collection

Volume based license models are making your job unnecessary hard. Price is a big factor here but it is even worse that volume based license models make you (or your manager makes you) try to save volume. This means that you will be finding yourself thinking about which data really needs to be ingested. The big problem is that you do not know what you might need the data for in the moment you are sending (or not sending) it. We have seen operations teams during a downtime wishing that they had collected the data of a certain log file that was now not searchable. **This is counter-productive and dangerous. You can be limited by disk space or other resources but never by the license that somebody bought.**

It is also a law of the market that you have to build your volume pricing model on the amount of data that is usually collected **today**. The amount of generated data has increased dramatically and vendors are nailed to their pricing model from 2008. This is why you get quotes that fill you with sadness in today's world.

Blackboxes

Closed source systems tend to become black boxes that you cannot extend or adapt to fit the needs of your use case. This is an important thing to consider especially for log management software. The use cases can range from simple syslog centralization to ultra flexible data bus requirements. A closed source system will always make you depending on the vendor because there is no way to adapt. As your setup reaches a certain point of flexibility you might hit a wall earlier than expected.

Consider spending a part of the money you would spend for the wrong license model for developing your own plugins or integrations.

The future

Graylog is the only open source log management system that will be able to deliver functionality and scaling in a way that Splunk does. It will be possible to replace Elasticsearch with something that is really suited for log data analysis without even changing the public facing APIs.

Graylog 1.0.2

Released: 2015-04-28

<https://www.graylog.org/graylog-v1-0-2-has-been-released/>

- Regular expression and Grok test failed when example message is a JSON document or contains special characters (Graylog2/graylog2-web-interface#1190, Graylog2/graylog2-web-interface#1195)
- “Show message terms” was broken (Graylog2/graylog2-web-interface#1168)
- Showing message indices was broken (Graylog2/graylog2-web-interface#1211)
- Fixed typo in SetIndexReadOnlyJob (Graylog2/graylog2-web-interface#1206)
- Consistent error messages when trying to create graphs from non-numeric values (Graylog2/graylog2-web-interface#1210)
- Fix message about too few file descriptors for Elasticsearch when number of file descriptors is unlimited (Graylog2/graylog2-web-interface#1220)
- Deleting output globally which was assigned to multiple streams left stale references (Graylog2/graylog2-server#1113)
- Fixed problem with sending alert emails (Graylog2/graylog2-server#1086)
- TokenizerConverter can now handle mixed quoted and un-quoted k/v pairs (Graylog2/graylog2-server#1083)

Graylog 1.0.1

Released: 2015-03-16

<https://www.graylog.org/graylog-v1-0-1-has-been-released/>

- Properly log stack traces (Graylog2/graylog2-server#970)

- Update REST API browser to new Graylog logo
- Avoid spamming the logs if the original input of a message in the disk journal can't be loaded (Graylog2/graylog2-server#1005)
- Allows reader users to see the journal status (Graylog2/graylog2-server#1009)
- Compatibility with MongoDB 3.0 and Wired Tiger storage engine (Graylog2/graylog2-server#1024)
- Respect `rest_transport_uri` when generating entity URLs in REST API (Graylog2/graylog2-server#1020)
- Properly map `NodeNotFoundException` (Graylog2/graylog2-web-interface#1137)
- Allow replacing all existing Grok patterns on bulk import (Graylog2/graylog2-web-interface#1150)
- Configuration option for discarding messages on error in AMQP inputs (Graylog2/graylog2-server#1018)
- Configuration option of maximum HTTP chunk size for HTTP-based inputs (Graylog2/graylog2-server#1011)
- Clone alarm callbacks when cloning a stream (Graylog2/graylog2-server#990)
- Add `hasField()` and `getField()` methods to `MessageSummary` class (Graylog2/graylog2-server#923)
- Add per input parse time metrics (Graylog2/graylog2-web-interface#1106)
- Allow the use of <https://logging.apache.org/log4j/extras/> `log4j-extras` classes in `log4j` configuration (Graylog2/graylog2-server#1042)
- Fix updating of input statistics for Radio nodes (Graylog2/graylog2-web-interface#1022)
- Emit proper error message when a regular expression in an Extractor doesn't match example message (Graylog2/graylog2-web-interface#1157)
- Add additional information to system jobs (Graylog2/graylog2-server#920)
- Fix false positive message on LDAP login test (Graylog2/graylog2-web-interface#1138)
- Calculate saved search resolution dynamically (Graylog2/graylog2-web-interface#943)
- Only enable LDAP test buttons when data is present (Graylog2/graylog2-web-interface#1097)
- Load more than 1 message on Extractor form (Graylog2/graylog2-web-interface#1105)
- Fix NPE when listing alarm callback using non-existent plugin (Graylog2/graylog2-web-interface#1152)
- Redirect to nodes overview when node is not found (Graylog2/graylog2-web-interface#1137)
- Fix documentation links to integrations and data sources (Graylog2/graylog2-web-interface#1136)
- Prevent accidental indexing of web interface by web crawlers (Graylog2/graylog2-web-interface#1151)
- Validate grok pattern name on the client to avoid duplicate names (Graylog2/graylog2-server#937)
- Add message journal usage to nodes overview page (Graylog2/graylog2-web-interface#1083)
- Properly format numbers according to locale (Graylog2/graylog2-web-interface#1128, Graylog2/graylog2-web-interface#1129)

Graylog 1.0.0

Released: 2015-02-19

<https://www.graylog.org/announcing-graylog-v1-0-ga/>

- No changes since Graylog 1.0.0-rc.4

Graylog 1.0.0-rc.4

Released: 2015-02-13

<https://www.graylog.org/graylog-v1-0-rc-4-has-been-released/>

- Default configuration file locations have changed. [Graylog2/graylog2-server#950](#)
- Improved error handling on search errors. [Graylog2/graylog2-server#954](#)
- Dynamically update dashboard widgets with keyword range. [Graylog2/graylog2-server#956](#), [Graylog2/graylog2-web-interface#958](#)
- Prevent duplicate loading of plugins. [Graylog2/graylog2-server#948](#)
- Fixed password handling when editing inputs. [Graylog2/graylog2-web-interface#1103](#)
- Fixed issues getting Elasticsearch cluster health. [Graylog2/graylog2-server#953](#)
- Better error handling for extractor imports. [Graylog2/graylog2-server#942](#)
- Fixed structured syslog parsing of keys containing special characters. [Graylog2/graylog2-server#845](#)
- Improved layout on Grok patterns page. [Graylog2/graylog2-web-interface#1109](#)
- Improved formatting large numbers. [Graylog2/graylog2-web-interface#1111](#)
- New Graylog logo.

Graylog 1.0.0-rc.3

Released: 2015-02-05

<https://www.graylog.org/graylog-v1-0-rc-3-has-been-released/>

- Fixed compatibility with MongoDB version 2.2. [Graylog2/graylog2-server#941](#)
- Fixed performance regression in process buffer handling. [Graylog2/graylog2-server#944](#)
- Fixed data type for the `max_size_per_index` config option value. [Graylog2/graylog2-web-interface#1100](#)
- Fixed problem with indexer error page. [Graylog2/graylog2-web-interface#1102](#)

Graylog 1.0.0-rc.2

Released: 2015-02-04

<https://www.graylog.org/graylog-v1-0-rc-2-has-been-released/>

- Better Windows compatibility. [Graylog2/graylog2-server#930](#)
- Added helper methods for the plugin API to simplify plugin development.
- Fixed problem with input removal on radio nodes. [Graylog2/graylog2-server#932](#)
- Improved buffer information for input, process and output buffers. [Graylog2/graylog2-web-interface#1096](#)
- Fixed API return value incompatibility regarding node objects. [Graylog2/graylog2-server#933](#)
- Fixed reloading of LDAP settings. [Graylog2/graylog2-server#934](#)
- Fixed ordering of message input state labels. [Graylog2/graylog2-web-interface#1094](#)

- Improved error messages for journal related errors. [Graylog2/graylog2-server#931](#)
- Fixed browser compatibility for stream rules form. [Graylog2/graylog2-web-interface#1095](#)
- Improved grok pattern management. [Graylog2/graylog2-web-interface#1099](#), [Graylog2/graylog2-web-interface#1098](#)

Graylog 1.0.0-rc.1

Released: 2015-01-28

<https://www.graylog.org/graylog-v1-0-rc-1-has-been-released/>

- Cleaned up internal metrics when input is terminating. [Graylog2/graylog2-server#915](#)
- Added Telemetry plugin options to example `graylog.conf`. [Graylog2/graylog2-server#914](#)
- Fixed problems with user permissions on streams. [Graylog2/graylog2-web-interface#1058](#)
- Added information about different rotation strategies to REST API. [Graylog2/graylog2-server#913](#)
- Added better error messages for failing inputs. [Graylog2/graylog2-web-interface#1056](#)
- Fixed problem with JVM options in `bin/radioctrl` script. [Graylog2/graylog2-server#918](#)
- Fixed issue with updating input configuration. [Graylog2/graylog2-server#919](#)
- Fixed password updating for reader users by the admin. [Graylog2/graylog2-web-interface#1075](#)
- Enabled the `message_journal_enabled` config option by default. [Graylog2/graylog2-server#924](#)
- Add REST API endpoint to list reopened indices. [Graylog2/graylog2-web-interface#1072](#)
- Fixed problem with GELF stream output. [Graylog2/graylog2-server#921](#)
- Show an error message on the indices page if the Elasticsearch cluster is not available. [Graylog2/graylog2-web-interface#1070](#)
- Fixed a problem with stopping inputs. [Graylog2/graylog2-server#926](#)
- Changed output configuration display to mask passwords. [Graylog2/graylog2-web-interface#1066](#)
- Disabled message journal on radio nodes. [Graylog2/graylog2-server#927](#)
- Create new message representation format for search results in alarm callback messages. [Graylog2/graylog2-server#923](#)
- Fixed stream router to update the stream engine if a stream has been changed. [Graylog2/graylog2-server#922](#)
- Fixed focus problem in stream rule modal windows. [Graylog2/graylog2-web-interface#1063](#)
- Do not show new dashboard link for reader users. [Graylog2/graylog2-web-interface#1057](#)
- Do not show stream output menu for reader users. [Graylog2/graylog2-web-interface#1059](#)
- Do not show user forms of other users for reader users. [Graylog2/graylog2-web-interface#1064](#)
- Do not show permission settings in the user profile for reader users. [Graylog2/graylog2-web-interface#1055](#)
- Fixed extractor edit form with no messages available. [Graylog2/graylog2-web-interface#1061](#)
- Fixed problem with node details page and JVM locale settings. [Graylog2/graylog2-web-interface#1062](#)
- Improved page layout for Grok patterns.
- Improved layout for the message journal information. [Graylog2/graylog2-web-interface#1084](#), [Graylog2/graylog2-web-interface#1085](#)

- Fixed wording on radio inputs page. [Graylog2/graylog2-web-interface#1077](#)
- Fixed formatting on indices page. [Graylog2/graylog2-web-interface#1086](#)
- Improved error handling in stream rule form. [Graylog2/graylog2-web-interface#1076](#)
- Fixed time range selection problem for the sources page. [Graylog2/graylog2-web-interface#1080](#)
- Several improvements regarding permission checks for user creation. [Graylog2/graylog2-web-interface#1088](#)
- Do not show stream alert test button for reader users. [Graylog2/graylog2-web-interface#1089](#)
- Fixed node processing status not updating on the nodes page. [Graylog2/graylog2-web-interface#1090](#)
- Fixed filename handling on Windows. [Graylog2/graylog2-server#928](#), [Graylog2/graylog2-server#732](#)

Graylog 1.0.0-beta.2

Released: 2015-01-21

<https://www.graylog.org/graylog-v1-0-beta-3-has-been-released/>

- Fixed stream alert creation. [Graylog2/graylog2-server#891](#)
- Suppress warning message when PID file doesn't exist. [Graylog2/graylog2-server#889](#)
- Fixed an error on outputs page with missing output plugin. [Graylog2/graylog2-server#894](#)
- Change default heap and garbage collector settings in scripts.
- Add extractor information to log message about failing extractor.
- Fixed problem in SplitAndIndexExtractor. [Graylog2/graylog2-server#896](#)
- Improved rendering time for indices page. [Graylog2/graylog2-web-interface#1060](#)
- Allow user to edit its own preferences. [Graylog2/graylog2-web-interface#1049](#)
- Fixed updating stream attributes. [Graylog2/graylog2-server#902](#)
- Stream throughput now shows combined value over all nodes. [Graylog2/graylog2-web-interface#1047](#)
- Fixed resource leak in JVM PermGen memory. [Graylog2/graylog2-server#907](#)
- Update to gelfclient-1.1.0 to fix DNS resolving issue. [Graylog2/graylog2-server#882](#)
- Allow arbitrary characters in user names (in fact in any resource url). [Graylog2/graylog2-web-interface#1005](#), [Graylog2/graylog2-web-interface#1006](#)
- Fixed search result CSV export. [Graylog2/graylog2-server#901](#)
- Skip GC collection notifications for parallel collector. [Graylog2/graylog2-server#899](#)
- Shorter reconnect timeout for Radio AMQP connections. [Graylog2/graylog2-server#900](#)
- Fixed random startup error in Radio. [Graylog2/graylog2-server#911](#)
- Fixed updating an alert condition. [Graylog2/graylog2-server#912](#)
- Add system notifications for journal related warnings. [Graylog2/graylog2-server#897](#)
- Add system notifications for failing outputs. [Graylog2/graylog2-server#741](#)
- Improve search result pagination. [Graylog2/graylog2-web-interface#834](#)
- Improved regex error handling in extractor testing. [Graylog2/graylog2-web-interface#1044](#)
- Wrap long names for node metrics. [Graylog2/graylog2-web-interface#1028](#)

- Fixed node information progress bars. [Graylog2/graylog2-web-interface#1046](#)
- Improve node buffer utilization readability. [Graylog2/graylog2-web-interface#1046](#)
- Fixed username alert receiver form field. [Graylog2/graylog2-web-interface#1050](#)
- Wrap long messages without break characters. [Graylog2/graylog2-web-interface#1052](#)
- Hide list of node plugins if there aren't any plugins installed.
- Warn user before leaving page with unpinned graphs. [Graylog2/graylog2-web-interface#808](#)

Graylog 1.0.0-beta.2

Released: 2015-01-16

<https://www.graylog.org/graylog-v1-0-0-beta2/>

- SIGAR native libraries are now found correctly (for getting system information)
- plugins can now state if they want to run in server or radio
- Fixed LDAP settings testing. [Graylog2/graylog2-web-interface#1026](#)
- Improved RFC5425 syslog message parsing. [Graylog2/graylog2-server#845](#)
- JVM arguments are now being logged on start. [Graylog2/graylog2-server#875](#)
- Improvements to log messages when Elasticsearch connection fails during start.
- Fixed an issue with AMQP transport shutdown. [Graylog2/graylog2-server#874](#)
- After index cycling the System overview page could be broken. [Graylog2/graylog2-server#880](#)
- Extractors can now be edited. [Graylog2/graylog2-web-interface#549](#)
- Fixed saving user preferences. [Graylog2/graylog2-web-interface#1027](#)
- Scripts now return proper exit codes. [Graylog2/graylog2-server#886](#)
- Grok patterns can now be uploaded in bulk. [Graylog2/graylog2-server#377](#)
- During extractor creation the test display could be offset. [Graylog2/graylog2-server#804](#)
- Performance fix for the System/Indices page. [Graylog2/graylog2-web-interface#1035](#)
- A create dashboard link was shown to reader users, leading to an error when followed. [Graylog2/graylog2-web-interface#1032](#)
- Content pack section was shown to reader users, leading to an error when followed. [Graylog2/graylog2-web-interface#1033](#)
- Failing stream outputs were being restarted constantly. [Graylog2/graylog2-server#741](#)

Graylog2 0.92.4

Released: 2015-01-14

<https://www.graylog.org/graylog2-v0-92-4/>

- [SERVER] Ensure that Radio inputs can only be started on server nodes ([Graylog2/graylog2-server#843](#))
- [SERVER] Avoid division by zero when finding rotation anchor in the time-based rotation strategy ([Graylog2/graylog2-server#836](#))

- [SERVER] Use username as fallback if display name in LDAP is empty (Graylog2/graylog2-server#837)

Graylog 1.0.0-beta.1

Released: 2015-01-12

<https://www.graylog.org/graylog-v1-0-0-beta1/>

- Message Journaling
- New Widgets
- Grok Extractor Support
- Overall stability and resource efficiency improvements
- Single binary for `graylog2-server` and `graylog2-radio`
- Inputs are now editable
- Order of field charts rendered inside the search results page is now maintained.
- Improvements in focus and keyboard behaviour on modal windows and forms.
- You can now define whether to disable expensive, frequent real-time updates of the UI in the settings of each user. (For example the updating of total messages in the system)
- Experimental search query auto-completion that can be enabled in the user preferences.
- The API browser now documents server response payloads in a better way so you know what to expect as an answer to your call.
- Now using the standard Java ServiceLoader for plugins.

Graylog2 0.92.3

Released: 2014-12-23

<https://www.graylog.org/graylog2-v0-92-3/>

- [SERVER] Removed unnecessary instrumentation in certain places to reduce GC pressure caused by many short living objects (Graylog2/graylog2-server#800)
- [SERVER] Limit Netty worker thread pool to 16 threads by default (see `rest_worker_threads_max_pool_size` in `graylog2.conf`)
- [WEB] Fixed upload of content packs when a URI path prefix (`application.context` in `graylog2-web-interface.conf`) is being used (Graylog2/graylog2-web-interface#1009)
- [WEB] Fixed display of metrics of type Counter (Graylog2/graylog2-server#795)

Graylog2 0.92.1

Released: 2014-12-11

<https://www.graylog.org/graylog2-v0-92-1/>

- [SERVER] Fixed name resolution and overriding sources for network inputs.

- [SERVER] Fixed wrong delimiter in GELF TCP input.
- [SERVER] Disabled the output cache by default. The output cache is the source of all sorts of interesting problems. If you want to keep using it, please read the upgrade notes.
- [SERVER] Fixed message timestamps in GELF output.
- [SERVER] Fixed connection counter for network inputs.
- [SERVER] Added warning message if the receive buffer size (SO_RECV) couldn't be set for network inputs.
- [WEB] Improved keyboard shortcuts with most modal dialogs (e. g. hitting Enter submits the form instead of just closing the dialogs).
- [WEB] Upgraded to play2-graylog2 1.2.1 (compatible with Play 2.3.x and Java 7).

Graylog2 0.92.0

Released: 2014-12-01

<https://www.graylog.org/graylog2-v0-92/>

- [SERVER] **IMPORTANT SECURITY FIX:** It was possible to perform LDAP logins with crafted wildcards. (A big thank you to Jose Tozo who discovered this issue and disclosed it very responsibly.)
- [SERVER] Generate a system notification if garbage collection takes longer than a configurable threshold.
- [SERVER] Added several JVM-related metrics.
- [SERVER] Added support for Elasticsearch 1.4.x which brings a lot of stability and resilience features to Elasticsearch clusters.
- [SERVER] Made version check of Elasticsearch version optional. Disabling this check is not recommended.
- [SERVER] Added an option to disable optimizing Elasticsearch indices on index cycling.
- [SERVER] Added an option to disable time-range calculation for indices on index cycling.
- [SERVER] Lots of other performance enhancements for large setups (i.e. involving several Radio nodes and multiple Graylog2 Servers).
- [SERVER] Support for Syslog Octet Counting, as used by syslog-ng for syslog via TCP (#743)
- [SERVER] Improved support for structured syslog messages (#744)
- [SERVER] Bug fixes regarding IPv6 literals in `mongodb_replica_set` and `elasticsearch_discovery_zen_ping_unicast_hosts`
- [WEB] Added additional details to system notification about Elasticsearch max. open file descriptors.
- [WEB] Fixed several bugs and inconsistencies regarding time zones.
- [WEB] Improved graphs and diagrams
- [WEB] Allow to update dashboards when browser window is not on focus (#738)
- [WEB] Bug fixes regarding timezone handling
- Numerous internal bug fixes

Graylog2 0.92.0-rc.1

Released: 2014-11-21

<https://www.graylog.org/graylog2-v0-92-rc-1/>

- [SERVER] Generate a system notification if garbage collection takes longer than a configurable threshold.
- [SERVER] Added several JVM-related metrics.
- [SERVER] Added support for Elasticsearch 1.4.x which brings a lot of stability and resilience features to Elasticsearch clusters.
- [SERVER] Made version check of Elasticsearch version optional. Disabling this check is not recommended.
- [SERVER] Added an option to disable optimizing Elasticsearch indices on index cycling.
- [SERVER] Added an option to disable time-range calculation for indices on index cycling.
- [SERVER] Lots of other performance enhancements for large setups (i. e. involving several Radio nodes and multiple Graylog2 Servers).
- [WEB] Upgraded to Play 2.3.6.
- [WEB] Added additional details to system notification about Elasticsearch max. open file descriptors.
- [WEB] Fixed several bugs and inconsistencies regarding time zones.
- Numerous internal bug fixes

Graylog2 0.91.3

Released: 2014-11-05

<https://www.graylog.org/graylog2-v0-90-3-and-v0-91-3-has-been-released/>

- Fixed date and time issues related to DST changes
- Requires Elasticsearch 1.3.4; Elasticsearch 1.3.2 had a bug that can cause index corruptions.
- The `mongodb_replica_set` configuration variable now supports IPv6
- Messages read from the on-disk caches could be stored with missing fields

Graylog2 0.91.3

Released: 2014-11-05

<https://www.graylog.org/graylog2-v0-90-3-and-v0-91-3-has-been-released/>

- Fixed date and time issues related to DST changes
- The `mongodb_replica_set` configuration variable now supports IPv6
- Messages read from the on-disk caches could be stored with missing fields

Graylog2 0.92.0-beta.1

Released: 2014-11-05

<https://www.graylog.org/graylog2-v0-92-beta-1/>

- Content packs
- [SERVER] SSL/TLS support for Graylog2 REST API
- [SERVER] Support for time based retention cleaning of your messages. The old message count based approach is still the default.
- [SERVER] Support for Syslog Octet Counting, as used by syslog-ng for syslog via TCP (Graylog2/graylog2-server#743)
- [SERVER] Improved support for structured syslog messages (Graylog2/graylog2-server#744)
- [SERVER] Bug fixes regarding IPv6 literals in `mongodb_replica_set` and `elasticsearch_discovery_zen_ping_unicast_hosts`
- [WEB] Revamped “Sources” page in the web interface
- [WEB] Improved graphs and diagrams
- [WEB] Allow to update dashboards when browser window is not on focus (Graylog2/graylog2-web-interface#738)
- [WEB] Bug fixes regarding timezone handling
- Numerous internal bug fixes

Graylog2 0.91.1

Released: 2014-10-17

<https://www.graylog.org/two-new-graylog2-releases/>

- Messages written to the persisted master caches were written to the system with unreadable timestamps, leading to
- errors when trying to open the message.
- Extractors were only being deleted from running inputs but not from all inputs
- Output plugins were not always properly loaded
- You can now configure the `alert_check_interval` in your `graylog2.conf`
- Parsing of configured Elasticsearch unicast discovery addresses could break when including spaces

Graylog2 0.90.1

Released: 2014-10-17

<https://www.graylog.org/two-new-graylog2-releases/>

- Messages written to the persisted master caches were written to the system with unreadable timestamps, leading to errors when trying to open the message.
- Extractors were only being deleted from running inputs but not from all inputs

- Output plugins were not always properly loaded
- You can now configure the `alert_check_interval` in your `graylog2.conf`
- Parsing of configured Elasticsearch unicast discovery addresses could break when including spaces

Graylog2 0.91.0-rc.1

Released: 2014-09-23

<https://www.graylog.org/graylog2-v0-90-has-been-released/>

- Optional Elasticsearch v1.3.2 support

Graylog2 0.90.0

Released: 2014-09-23

<https://www.graylog.org/graylog2-v0-90-has-been-released/>

- Real-time data forwarding to Splunk or other systems
- Alert callbacks for greater flexibility
- New disk-based architecture for buffering in load spike situations
- Improved graphing
- Plugin API
- Huge performance and stability improvements across the whole stack
- Small possibility of losing messages in certain scenarios has been fixed
- Improvements to internal logging from threads to avoid swallowing Graylog2 error messages
- Paused streams are no longer checked for alerts
- Several improvements to timezone handling
- JavaScript performance fixes in the web interface and especially a fixed memory leak of charts on dashboards
- The GELF HTTP input now supports CORS
- Stream matching now has a configurable timeout to avoid stalling message processing in case of too complex rules or erroneous regular expressions
- Stability improvements for Kafka and AMQP inputs
- Inputs can now be paused and resumed
- Dozens of bug fixes and other improvements

Graylog2 0.20.3

Released: 2014-08-09

<https://www.graylog.org/graylog2-v0-20-3-has-been-released/>

- Bugfix: Storing saved searches was not accounting custom application contexts

- Bugfix: Editing stream rules could have a wrong a pre-filled value
- Bugfix: The create dashboard link was shown even if the user has no permission to so. This caused an ugly error page because of the missing permissions.
- Bugfix: graylog2-radio could lose numeric fields when writing to the message broker
- Better default batch size values for the Elasticsearch output
- Improved `rest_transport_uri` default settings to avoid confusion with loopback interfaces
- The deflector index is now also using the configured index prefix

Graylog2 0.20.2

Released: 2014-05-24

<https://www.graylog.org/graylog2-v0-20-2-has-been-released/>

- Search result highlighting
- Reintroduces AMQP support
- Extractor improvements and sharing
- Graceful shutdowns, Lifecycles, Load Balancer integration
- Improved stream alert emails
- Alert annotations
- CSV exports via the REST API now support chunked transfers and avoid heap size problems with huge result sets
- Login now redirects to page you visited before if there was one
- More live updating information in node detail pages
- Empty dashboards no longer show lock/unlock buttons
- Global inputs now also show IO metrics
- You can now easily copy message IDs into native clipboard with one click
- Improved message field selection in the sidebar
- Fixed display of floating point numbers in several places
- Now supporting application contexts in the web interface like `http://example.org/graylog2`
- Several fixes for LDAP configuration form
- Message fields in the search result sidebar now survive pagination
- Only admin users are allowed to change the session timeout for reader users
- New extractor: Copy whole input
- New converters: uppercase/lowercase, flexdate (tries to parse any string as date)
- New stream rule to check for presence or absence of fields
- Message processing now supports trace logging
- Better error message for ES discovery problems
- Fixes to GELF HTTP input and it holding open connections

- Some timezone fixes
- CSV exports now only contain selected fields
- Improvements for bin/graylog* control scripts
- UDP inputs now allow for custom receive buffer sizes
- Numeric extractor converter now supports floating point values
- Bugfix: Several small fixes to system notifications and closing them
- Bugfix: Carriage returns were not escaped properly in CSV exports
- Bugfix: Some AJAX calls redirected to the startpage when they failed
- Bugfix: Wrong sorting in sources table
- Bugfix: Quickvalues widget was broken with very long values
- Bugfix: Quickvalues modal was positioned wrong in some cases
- Bugfix: Indexer failures list could break when you had a lot of failures
- Custom application prefix was not working for field chart analytics
- Bugfix: Memory leaks in the dashboards
- Bugfix: NullPointerException when Elasticsearch discovery failed and unicast discovery was disabled
- Message backlog in alert emails did not always include the correct number of messages
- Improvements for message outputs: No longer only waiting for filled buffers but also flushing them regularly. This avoids problems that make Graylog2 look like it misses messages in cheap benchmark scenarios combined with only little throughput.