

Installing Graylog2 on Arch Linux

07/17/2014 07:58 AM - Daniel Curtis

Status:	Suspended	Start date:	07/17/2014
Priority:	Normal	Due date:	
Assignee:	Daniel Curtis	% Done:	70%
Category:	Logging Server	Estimated time:	4.00 hours
Target version:		Spent time:	4.00 hours

Description

Prerequisites

You will need the following environment:

- Arch Linux
- 2GB RAM (I ran into issues using less)
- 80GB HD

You will need to have the following services installed on either the host you are running graylog2-server on or on dedicated machines:

1. ElasticSearch v0.90.10
2. MongoDB (as recent stable version as possible, at least v2.0)

Install Java Runtime Environment 7

- Start by installing Java 7:

```
pacman -S jre7-openjdk
```

Install MongoDB

- Install MongoDB:

```
pacman -S mongodb
```

Start and enable MongoDB to start at boot:

```
systemctl start mongod.service
systemctl enable mongod.service
```

Install ElasticSearch v0.90.10

NOTE: You must use ElasticSearch v0.90.10 to avoid compatibility problems.

```
cd /opt
wget https://download.elasticsearch.org/elasticsearch/elasticsearch/elasticsearch-0.90.10.tar.gz
tar xzf elasticsearch-0.90.10.tar.gz
mv elasticsearch-0.90.10 elasticsearch
cd elasticsearch
bin/elasticsearch -f
```

You can test if it works using:

```
curl -X GET http://localhost:9200/
```

- Create a separate user and group for Elasticsearch to run as:

```
groupadd -g 700 elasticsearch
useradd -u 700 -g elasticsearch -c "Elasticsearch User" -d /var/lib/elasticsearch elasticsearch
```

- Create the Elasticsearch directory structure

```
mkdir -p /var/lib/elasticsearch
chown elasticsearch:elasticsearch /var/lib/elasticsearch

mkdir -p /var/log/elasticsearch
chown elasticsearch:elasticsearch /var/log/elasticsearch

mkdir -p /etc/elasticsearch
ln -s /opt/elasticsearch/config/elasticsearch.yml /etc/elasticsearch/elasticsearch.yml
ln -s /opt/elasticsearch/config/logging.yml /etc/elasticsearch/logging.yml
```

Next, edit the configuration file `/etc/elasticsearch/elasticsearch.yml` to change the default cluster name, to set the config path, and to use the directories you created in the previous step.

The important thing for Elasticsearch is that you configure cluster.name: graylog2.

- By default, Elasticsearch uses the cluster name "elasticsearch".

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

1. Change the following configuration parameters:

```
# Cluster name identifies your cluster for auto-discovery.
# cluster.name: elasticsearch
cluster.name: graylog2

# Path to directory containing configuration (this file and logging.yml):
# path.conf: /path/to/conf
path.conf: /etc/elasticsearch/elasticsearch.yml

# Path to directory where to store index data allocated for this node.
# path.data: /path/to/data
path.data: /var/lib/elasticsearch

# Path to log files:
# path.logs: /path/to/logs
path.logs: /var/log/elasticsearch
```

For best performance, you will want to adjust the processors directive in the `elasticsearch.yml` configuration file.

- Then add it to your configuration file.

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

```
processors: 17
Running Elasticsearch
```

- You can run Elasticsearch directly like this. The `-d` option runs it as a daemon:

```
/opt/elasticsearch/bin/elasticsearch -f
```

- Test that it's running:

```
curl localhost:9200
```

```
{
  "status" : 200,
  "name" : "NFL Superpro",
  "version" : {
    "number" : "1.0.1",
    "build_hash" : "5c03844e1978e5cc924dab2a423dc63ce881c42b",
    "build_timestamp" : "2014-02-25T15:52:53Z",
    "build_snapshot" : false,
    "lucene_version" : "4.6"
  },
  "tagline" : "You Know, for Search"
}
```

Make sure to also read these pages:

- [Graylog2 architecture high level overview](#)
- [The Graylog2 index model explained](#)
- [Configuring and tuning Elasticsearch for Graylog2](#)

Install the Graylog2 server and web interface

- Download the package from the Graylog site:

```
wget https://github.com/Graylog2/graylog2-server/releases/download/0.20.5/graylog2-server-0.20.5.tgz
wget https://github.com/Graylog2/graylog2-web-interface/releases/download/0.20.5/graylog2-web-interface-0.20.5.tgz
wget https://github.com/Graylog2/graylog2-server/releases/download/0.20.5/graylog2-radio-0.20.5.tgz
```

- Extract the archive:

```
tar xfz graylog2-server-0.20.5.tgz && tar xfz graylog2-web-interface-0.20.5.tgz && tar xzf graylog2-radio-0.20.5.tgz
cd graylog2-server-0.20.5
```

Configuration

- Now copy the example configuration file:

```
cp graylog2.conf.example /etc/graylog2.conf
```

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

Configure at least these variables in `/etc/graylog2.conf`:

```
is_master = true
```

- Set only one graylog2-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 -s 96`. If you run multiple Graylog2 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 Graylog2 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.

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

- The first start should be performed without the `bin/graylog2ctl` script to easily see warnings, errors or problems:

```
java -jar graylog2-server.jar --debug
```

The server will try to write a node_id to graylog2-server-node-id. It won't start if it can't write there because of for example missing permissions.

See the [startup parameters](#) page to learn more about available startup parameters. Note that you might have to be root to bind to port 514 for syslog.

You should see a line like this in the debug output if graylog2-server successfully connected to your ElasticSearch cluster:

```
2013-10-01 12:13:22,382 DEBUG: org.elasticsearch.transport.netty - [graylog2-server] connected to node [[Unuscione, Angelo][thN_gIBkQDm2ab7k-2Zaaw][inet[/10.37.160.227:9300]]]
```

This line indicates that your graylog2-server instance is up and ready to accept messages:

```
2013-10-01 12:13:53,149 INFO : org.graylog2.Core - Graylog2 up and running.
```

- Now exit and start the server using the control script:

```
cd bin/  
./graylog2ctl start
```

This will start your graylog2-server in the background. Find logs in logs/.

IMPORTANT: All graylog2-server instances must have synchronized time. We strongly recommend to use NTP on all machines of your Graylog2 infrastructure.

That's it! Now go on by installing the graylog2-webinterface to finish your installation.

Install Graylog2 web interface

We will download and install the Graylog2 v.0.20.2 web interface in /opt with the following commands:

```
cd /opt;
```

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

```
sudo ln -s graylog2-web-interface-0.20.2 graylog2-web-interface
```

Next, we want to configure the web interface's secret key, the application.secret parameter in graylog2-web-interface.conf. We will generate another key, as we did with the Graylog2 server configuration, and insert it with sed, like so:

```
SECRET=$(pwgen -s 96 1)  
sudo -E sed -i -e 's/application\.secret=""/application\.secret="'$SECRET'"/' /opt/graylog2-web-interface/conf/graylog2-web-interface.conf
```

Now open the web interface configuration file, with this command:

```
sudo vi /opt/graylog2-web-interface/conf/graylog2-web-interface.conf
```

Now we need to update the web interface's configuration to specify the graylog2-server.uris parameter. This is a comma delimited list of the server REST URLs. Since we only have one Graylog2 server node, the value should match that of rest_listen_uri in the Graylog2 server configuration (i.e. "http://127.0.0.1:12900/").

```
graylog2-server.uris="http://127.0.0.1:12900/"
```

The Graylog2 web interface is now configured. Let's start it up to test it out:

```
sudo /opt/graylog2-web-interface/bin/graylog2-web-interface
```

Troubleshooting

Problems with IPv6 vs. IPv4?

If your graylog2-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 graylog2ctl script or from wherever you are calling the graylog2-server.jar:

```
java -Djava.net.preferIPv4Stack=true -jar graylog2-server.jar
```

Resources

*<http://support.torch.sh/help/kb/graylog2-server/installing-graylog2-server-v020x-on-nix-systems>

- <https://wiki.joyent.com/wiki/display/jpc2/Installing+Elasticsearch+From+Source+on+SmartOS>
- <http://www.elasticsearch.org/downloads/0-90-10/>

Related issues:

Copied to GNU/Linux Administration - Support #435: Installing Graylog2 on Deb...

Suspended **07/17/2014**

History

#1 - 07/17/2014 08:45 AM - Daniel Curtis

- Description updated

- % Done changed from 20 to 30

#2 - 07/17/2014 09:00 AM - Daniel Curtis

- Description updated

#3 - 07/17/2014 09:01 AM - Daniel Curtis

- Description updated

#4 - 07/17/2014 09:03 AM - Daniel Curtis

- Description updated

#5 - 07/17/2014 09:17 AM - Daniel Curtis

- Description updated

- % Done changed from 30 to 60

#6 - 07/17/2014 09:37 AM - Daniel Curtis

- Description updated

#7 - 07/21/2014 07:50 AM - Daniel Curtis

- Description updated

- % Done changed from 60 to 70

#8 - 07/21/2014 09:38 AM - Daniel Curtis

- Description updated

#9 - 07/23/2014 08:18 AM - Daniel Curtis

- Description updated

#10 - 08/12/2014 07:28 AM - Daniel Curtis

- Copied to Support #435: Installing Graylog2 on Debian 7 added

#11 - 02/15/2015 08:39 PM - Daniel Curtis

- Project changed from 94 to GNU/Linux Administration

- Category set to Source Code Management

#12 - 06/04/2017 08:23 PM - Daniel Curtis

- Category changed from Source Code Management to Logging Server

- Status changed from In Progress to Suspended