

Graylog Hitchhiker's Guide to the Galaxy

Resources/Sources	3
Logging basics	3
What to search	3
Graylog components	3
Resource/sources	3
Graylog stream	4
Graylog inputs	4
Graylog alerts	5
Graylog Extractors/grok patterns	6
Install/Setup Graylog	7
Install/Setup Graylog	7
Install/Setup MongoDB	7
Install/Setup Elasticsearch	8
Install/Setup Graylog	8
Install/Setup Nginx and OpenSSL	9
Install/Setup FirewallD	9
Graylog Client Setup	10
Install/Setup Filebeat on CentOS 7 64-bit	10
Install/Setup Packetbeat on Centos 7 64-bit	11
Install/Setup Winlogbeat on Windows 7 64-bit	11
Graylog searching How-to	11
Resources/sources	11
Search by timeframe	12
String based searches	12
Search by key:value pair	14
Multiple key:value pairs	14
Graylog Create an alert	15
Resources/sources	15

Create Graylog stream to filter logs
Creating alert for unauthorized SSH login

15
16

Resources/Sources

- <http://docs.graylog.org/en/2.2/>
- <http://edbaker.weebly.com/blog/windows-and-logstash-quick-n-dirty>
- <https://github.com/elastic/beats/blob/master/winlogbeat/docs/getting-started.asciidoc>
- http://docs.graylog.org/en/2.2/pages/getting_started/stream_alerts.html

Logging basics

- Log - A saved event of an observable occurrence in an information system that actually happened at some point in time.
- Elements of a log
 - Who, When, or What performed the activity
 - Type of action
 - Ex: Authorize, create, read, update, delete, and accept such as a network connection
 - Identifiers - Files accessed, query parameters, and etc
 - Before and after values of action performed
-

What to search

- Step one: know what to look for.
 - It seems simple in theory but in practice/real life it's not. So break things down.
 - If we know they comprised a Windows system then only look at Windows logs.
 - We may know the account compromised so we can lookup authorized user logins.
 - We may know the data they exfiltrated from the network then we can look up who has access to that file.
- When logging is done correctly we should have records of all events and the before and after value/result of each event.

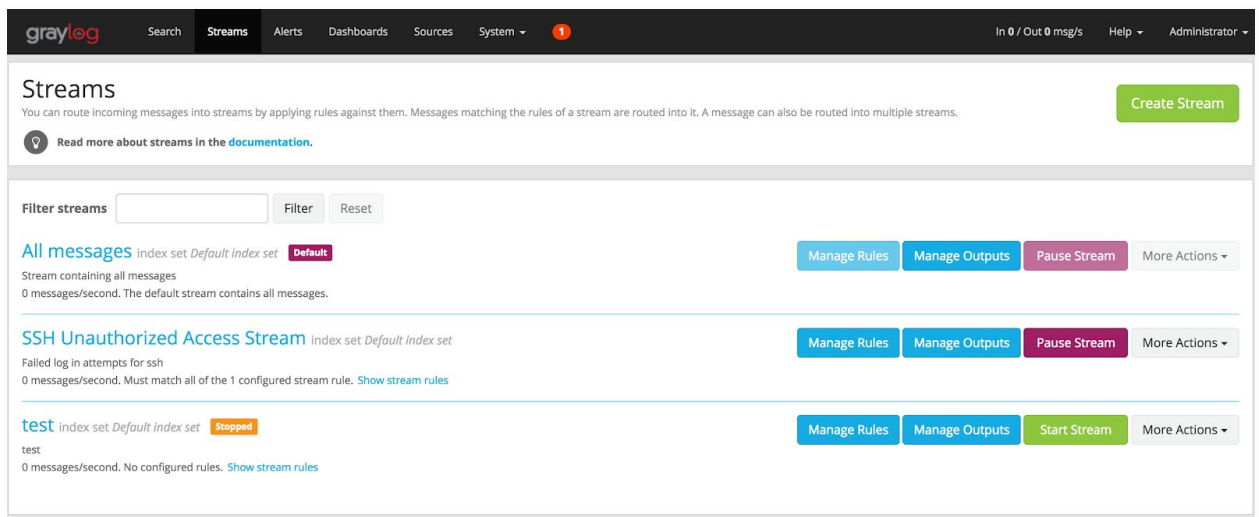
Graylog components

Resource/sources

- <http://docs.graylog.org/en/2.2/pages/streams.html>

Graylog stream

- Graylog > Streams
- [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



Graylog inputs

- Graylog > System > inputs
- This will tell Graylog to accept the log [messages](#).
- Input types
 - Beats (Filebeat, Packetbeat, Winlogbeat)
 - Syslog
 - Json via HTTP
 - GELF

Launch new input: Syslog UDP ✕

Node(s) to spawn input on:
Select the node you want to spawn this input on.

ac472930 / graylog

or:

Global input (started on all nodes)

Title
Select a name of your new input that describes it.

Syslog UDP

Bind address
Address to listen on. For example 0.0.0.0 or 127.0.0.1

127.0.0.1

Port
Port to listen on.

5140

Receive Buffer Size (optional)
The size in bytes of the recvBufferSize for network connections to this input.

262144

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.

Force rDNS? (optional)
Force rDNS resolution of hostname? Use if hostname cannot be parsed.

Allow overriding date? (optional)
Allow to override with current date if date could not be parsed?

Store full message? (optional)
Store the full original syslog message as full_message?

Expand structured data? (optional)
Expand structured data elements by prefixing attributes with their SD-ID?

Close Launch

Graylog alerts

- Graylog > Alerts tab
- Instruct Graylog to create [alerts](#) or send e-mail notifications when predefined events meet condition(s).
- Alert conditions are based off Graylog streams

graylog Search Streams Alerts Dashboards Sources System 11/07/2017 08:00 msgs Help Administrator

Alerts overview

Alerts are triggered when conditions you define are satisfied. Graylog will automatically mark alerts as resolved once the status of your conditions change.

[Manage conditions](#) [Manage notifications](#)

[Read more about alerting in the documentation.](#)

Alerts

Check your alerts status from here. Currently displaying all alerts.

[Refresh](#) [Show unresolved alerts](#)

SSH Unauthorized Access Alert on stream *SSH Unauthorized Access Stream* Resolved

Triggered at 2017-03-24 14:12:10, resolved at 2017-03-24 14:17:10.

Reason: Stream had 6 messages in the last 5 minutes with trigger condition more than 3 messages. (Current grace time: 15 minutes)

Type: Message Count Alert Condition

« < 1 > »

Graylog Extractors/grok patterns

- Graylog > System> Grok patterns
- [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.
 - 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.
- [Grok debugger/creator](#)

graylog Search Streams Alerts Dashboards Sources System / Inputs

New extractor for input Syslog

Extractors are applied on every message that is received by an input. Use them to extract and transform any text data into fields that allow you easy filtering and analysis later on.

Find more information about extractors in the [documentation](#).

Example message

```
mgmt-mongo sshd[16144]: Invalid user oracle from 203.0.113.42
```

Wrong example? You can [load another message](#).

Extractor configuration

Extractor type Regular expression

Source field message

Regular expression [Try](#)

The regular expression used for extraction. First matcher group is used. Learn more in the [documentation](#).

Extractor preview

```
oracle
```

Condition

- Always try to extract
- Only attempt extraction if field contains string
- Only attempt extraction if field matches regular expression

Extracting only from messages that match a certain condition helps you avoiding wrong or unnecessary extractions and can also save CPU resources.

Store as field

Choose a field name to store the extracted value. It can only contain alphanumeric characters and underscores. Example: `http_response_code`.

Extraction strategy Copy Cut

Do you want to copy or cut from source? You cannot use the cutting feature on standard fields like `message` and `source`.

Extractor title

A descriptive name for this extractor.

Add converter [Add](#)

Add converters to transform the extracted value.

[Create extractor](#)

Install/Setup Graylog

Install/Setup Graylog

1. yum update -y && yum install upgrade -y
2. yum install epel-release -y && yum update -y
3. yum install -y vim net-tools pwgen
4. yum install java-1.8.0-openjdk-headless.x86_64

Install/Setup MongoDB

5. cat > /etc/yum.repos.d/mongodb-org-3.2.repo << EOF

```
[mongodb-org-3.2]
```

```
name=MongoDB Repository
```

```
baseurl=https://repo.mongodb.org/yum/redhat/$releasever/mongodb-org/3.2/x86_64/
```

```
gpgcheck=1
```

```
enabled=1
```

```
gpgkey=https://www.mongodb.org/static/pgp/server-3.2.asc
```

```
EOF
```

6. yum install -y mongodb-org
7. chkconfig --add mongod
8. systemctl daemon-reload
9. systemctl enable mongod.service
10. systemctl start mongod.service

Install/Setup Elasticsearch

1. rpm --import <https://packages.elastic.co/GPG-KEY-elasticsearch>
2. cat > /etc/yum.repos.d/elasticsearch.repo << EOF

```
[elasticsearch-2.x]
```

```
name=Elasticsearch repository for 2.x packages
```

```
baseurl=https://packages.elastic.co/elasticsearch/2.x/centos
```

```
gpgcheck=1
```

```
gpgkey=https://packages.elastic.co/GPG-KEY-elasticsearch
```

```
enabled=1
```

```
EOF
```

3. yum install -y elasticsearch
4. yum update --exclude=elasticsearch-2.x
 - a. GRAYLOG NEEDS ELASTICSEARCH 2
5. sed -i 's/# cluster.name: my-application/cluster.name: graylog/g'
/etc/elasticsearch/elasticsearch.yml
6. chkconfig --add elasticsearch
7. systemctl daemon-reload
8. systemctl enable elasticsearch.service
9. systemctl restart elasticsearch.service

Install/Setup Graylog

1. rpm -Uvh https://packages.graylog2.org/repo/packages/graylog-2.1-repository_latest.rpm
2. yum install graylog-server
3. echo -n yourpassword | sha256sum
 - a. Copy output text
4. sed -i 's/root_password_sha2 =/root_password_sha2 = <hash from above>/g'
/etc/graylog/server/server.conf

5. pwgen -N 1 -s 96
 - a. Copy output text
6. sed -i 's/password_secret =/password_secret = <pwgen output> /g' /etc/graylog/server/server.conf
7. chkconfig --add graylog-server
8. systemctl daemon-reload
9. systemctl enable graylog-server.service
10. systemctl start graylog-server.service

Install/Setup Nginx and OpenSSL

1. yum install nginx -y
2. mkdir /etc/nginx/ssl
3. openssl req -x509 -nodes -days 365 -newkey rsa:2048 -keyout /etc/nginx/ssl/nginx.key -out /etc/nginx/ssl/nginx.crt
4. openssl dhparam -out /etc/ssl/certs/dhparam.pem 2048
5. sed -i -e '38,87d' /etc/nginx/nginx.conf
6. cat > /etc/nginx/conf.d/graylog.conf << EOF
 - a. [See resource above for examples](#)
7. sudo setsebool -P httpd_can_network_connect 1
8. systemctl enable nginx
9. systemctl start nginx

Install/Setup FirewallD

1. yum install firewalld -y
2. systemctl enable firewalld
3. systemctl start firewalld
4. firewall-cmd --permanent --add-service=ssh
5. firewall-cmd --permanent --add-service=http
6. firewall-cmd --permanent --add-service=https
7. firewall-cmd --permanent --add-port=5044/tcp
8. firewall-cmd --reload

Graylog Client Setup

Install/Setup Filebeat on CentOS 7 64-bit

- Filebeat helps you keep the simple things simple by offering a lightweight way to forward and centralize logs and files.

- Operating System: Linux and Windows
1. yum update -y && yum upgrade -y
 2. sudo rpm --import <https://packages.elastic.co/GPG-KEY-elasticsearch>
 3. cat > /etc/yum.repos.d/elastic.repo << EOF


```
[elastic-5.x]
name=Elastic repository for 5.x packages
baseurl=https://artifacts.elastic.co/packages/5.x/yum
gpgcheck=1
gpgkey=https://artifacts.elastic.co/GPG-KEY-elasticsearch
enabled=1
autorefresh=1
type=rpm-md
EOF
```
 4. yum install filebeat -y
 5. mkdir /etc/filebeat/conf.d/
 6. cp /etc/filebeat/filebeat.yml /etc/filebeat/filebeat.yml.bak
 7. cat > /etc/filebeat/filebeat.yml << EOF


```
filebeat:
  registry_file: /var/lib/filebeat/registry
  config_dir: /etc/filebeat/conf.d

output.logstash:
  hosts: ["<hostname/IP Addr>:5044"]
EOF
```
 8. cat > /etc/filebeat/conf.d/logging.yml << EOF


```
filebeat.prospectors:
- paths:
  - /var/log/*
input_type: log
EOF
```
 9. systemctl enable filebeat
 10. systemctl start filebeat

Install/Setup Packetbeat on Centos 7 64-bit

- Packetbeat is a lightweight network packet analyzer that sends data to Logstash or Elasticsearch.
- Operating system: Linux and Windows

Install/Setup Winlogbeat on Windows 7 64-bit

- Winlogbeat
- Operating system: Windows
- 1. Download the Winlogbeat zip file from the [downloads page](#).
- 2. Extract the contents into C:\Program Files.
- 3. Rename the winlogbeat-<version> directory to Winlogbeat.
- 4. Open a PowerShell prompt as an Administrator (right-click on the PowerShell icon and select Run As Administrator).
 - a. If you are running Windows XP, you may need to download and install PowerShell.
- 5. cd 'C:\Program Files\Winlogbeat'
- 6. .\install-service-winlogbeat.ps1
 - a. Run ExecutionPolicy UnRestricted
- 7. Edit winlogbeat.yml
 - a. Comment out “#output.elasticsearch:
#hosts:
- localhost:9200”
 - b. Uncomment “output.elasticsearch:
hosts:
- <IP Addr of graylog>:9200”
 - c. Save, exit
- 8. Start-Service winlogbeat

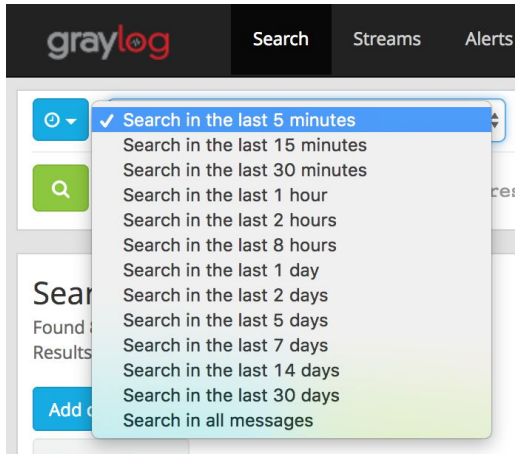
Graylog searching How-to

Resources/sources

- <https://www.youtube.com/watch?v=vxmAIDZe1j0>

Search by timeframe

1. Select the “Search” tab
2. Select the drop menu for time



String based searches

- Let's search for domain names
 1. Enter "google.com" into the search.
 2. Select the search icon



Search in the last 1 day



google.com

3. This search will populate results within the specified time frame

The screenshot shows a search interface with a search bar at the top left containing "Search in the last 1 day". Below the search bar, there's a "Search result" section with a "Histogram" chart. The histogram shows a bar at 06:00 and another at 03:00. Below the histogram is a "Messages" table with columns for "Timestamp" and "source". The table lists several messages with their timestamps and sources.

Timestamp	source
2017-03-23 05:31:06.869	freeipa.hackinglab.tech
23-Mar-2017 01:31:05.465	client 10.140.100.15#25958 (lh3.google.com): query: lh3.google.com IN A + (10.140.100.253)
2017-03-23 05:30:21.864	freeipa.hackinglab.tech
23-Mar-2017 01:30:19.051	client 10.140.100.15#30653 (clients2.google.com): query: clients2.google.com IN A + (10.140.100.253)
2017-03-23 05:29:28.857	freeipa.hackinglab.tech
23-Mar-2017 01:29:26.722	client 10.140.100.15#41269 (clients4.google.com): query: clients4.google.com IN A + (10.140.100.253)
2017-03-23 05:27:58.849	freeipa.hackinglab.tech
23-Mar-2017 01:27:57.748	client 10.140.100.15#2493 (clients6.google.com): query: clients6.google.com IN A + (10.140.100.253)
2017-03-23 05:24:11.830	freeipa.hackinglab.tech
23-Mar-2017 01:24:09.075	client 10.140.100.15#18024 (clients4.google.com): query: clients4.google.com IN A + (10.140.100.253)

4. Select an entry from above to see all the detail
a. We can see below that all entries are a key:value pair

The screenshot shows a detailed view of a message. It includes a "Received by" section with a link to the message, a "Stored in index" section with the index name, and a "Routed into streams" section with a link to all messages. The main part of the view is a list of fields and their values, including "GREEDYDATA", "facility", "file", "input_type", "message", "name", "offset", "query", "source", "timestamp", and "type".

2017-03-23 05:31:06.869 freeipa.hackinglab.tech
23-Mar-2017 01:31:05.465 client 10.140.100.15#25958 (lh3.google.com): query: lh3.google.com IN A + (10.140.100.253)

✉ eb61c980-0f89-11e7-90e2-000c295c7189

Received by
Beats input on [a508cf4f / logging.hackinglab.tech](#)

Stored in index
graylog_0

Routed into streams
• [All messages](#)

GREEDYDATA
23-Mar-2017 01:31:05.465 client 10.140.100.15#25958 (lh3.google.com):

facility
filebeat

file
/var/named/data/named.run

input_type
log

message
23-Mar-2017 01:31:05.465 client 10.140.100.15#25958 (lh3.google.com): query: lh3.google.com IN A + (10.140.100.253)

name
freeipa.hackinglab.tech

offset
3446492

query
lh3.google.com

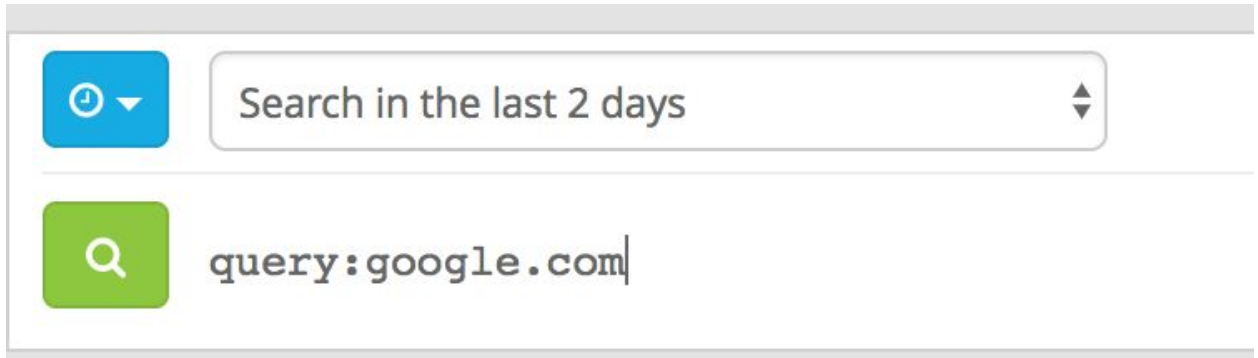
source
freeipa.hackinglab.tech

timestamp
2017-03-23T05:31:06.869Z

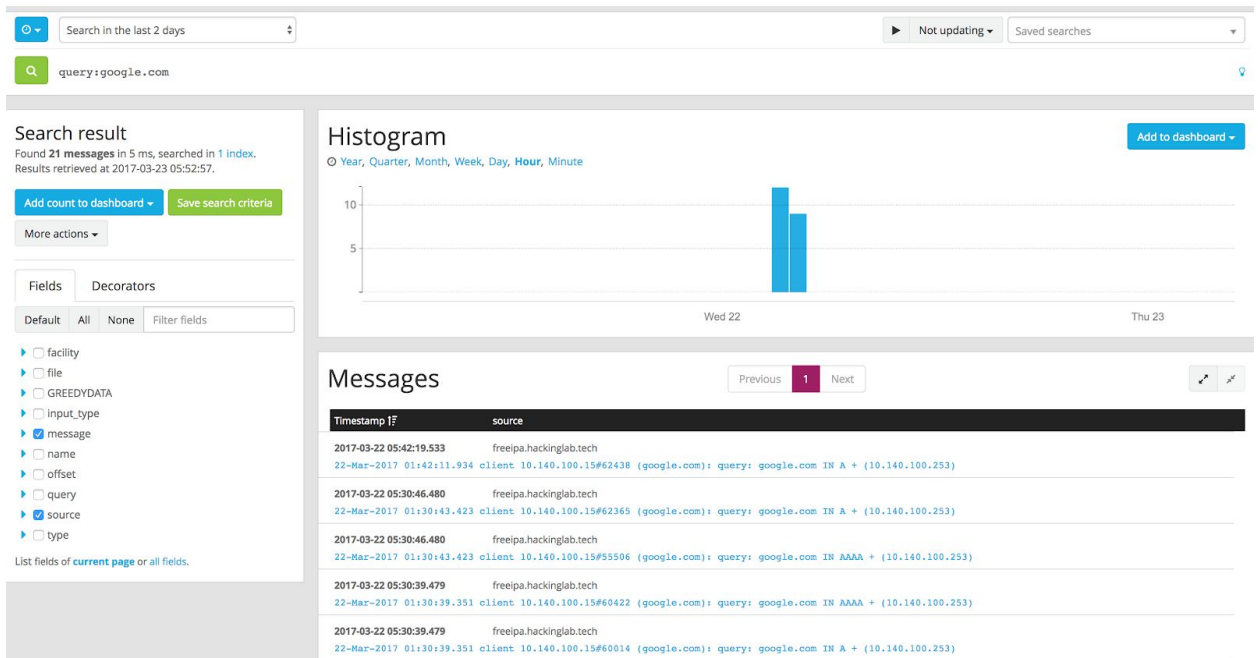
type
syslog

Search by key:value pair

- Let's search for every log entry where the key is query and the value is google.com
 - Enter "query:google.com" into the search
 - Select the search icon



- The search will populate entries

A screenshot of a search results page. The top section shows the search bar with "query:google.com" and a "Search in the last 2 days" dropdown. Below the search bar, there is a "Search result" section with a "Histogram" chart. The histogram shows a single bar for "Wed 22" with a count of approximately 10. Below the histogram is a "Messages" section with a table of log entries. The table has columns for "Timestamp" and "source".

Timestamp	source
2017-03-22 05:42:19.533	freeipa.hackinglab.tech
22-Mar-2017 01:42:11.934 client 10.140.100.15#62438 (google.com): query: google.com IN A + (10.140.100.253)	
2017-03-22 05:30:46.480	freeipa.hackinglab.tech
22-Mar-2017 01:30:43.423 client 10.140.100.15#62365 (google.com): query: google.com IN A + (10.140.100.253)	
2017-03-22 05:30:46.480	freeipa.hackinglab.tech
22-Mar-2017 01:30:43.423 client 10.140.100.15#5506 (google.com): query: google.com IN AAAA + (10.140.100.253)	
2017-03-22 05:30:39.479	freeipa.hackinglab.tech
22-Mar-2017 01:30:39.351 client 10.140.100.15#60422 (google.com): query: google.com IN AAAA + (10.140.100.253)	
2017-03-22 05:30:39.479	freeipa.hackinglab.tech
22-Mar-2017 01:30:39.351 client 10.140.100.15#60014 (google.com): query: google.com IN A + (10.140.100.253)	

Multiple key:value pairs

- Let's search for every log entry where the key is and the values google.com

Graylog Create an alert

Resources/sources

- http://docs.graylog.org/en/2.2/pages/getting_started/stream_alerts.html

Create Graylog stream to filter logs

1. Login into graylog
2. Select “Stream” tab
3. Select “Create Stream”
 - a. Enter “SSH Unauthorized Access Stream” for title
 - b. Enter “Failed login attempts for ssh” for description
 - c. Select “Default index” for index

Creating Stream

Title
SSH Unauthorized Access Stream

Description
Failed log in attempts for ssh

Index Set
Default index set

Messages that match this stream will be written to the configured index set.

Remove matches from 'All messages' stream

Remove messages that match this stream from the 'All messages' stream which is assigned to every message by default.

Cancel Save

- d. Select “Save”
4. Select “Manage rules” for “SSH Unauthorized Access Stream”
 5. Select “Add stream rule”
 - a. Enter “message” for field
 - b. Select “contain” for type

- c. Enter “pam_unix(sshd:auth): authentication failure” for value

The screenshot shows a 'New Stream Rule' dialog box with the following fields and values:

- Field:** message
- Type:** contain
- Value:** pam_unix(sshd:auth): authentication failure
- Description (optional):** (empty)
- Result:** Field message must contain pam_unix(sshd:auth): authentication failure

There is a 'Save' button highlighted in purple at the bottom right.

- d. Select “Save”
6. Select “I’m done”

Creating alert for unauthorized SSH login

1. Log into graylog
2. Select the “Alert” tab
3. Select “Manage conditions”
4. Select “Add new condition”
 - a. Select “SSH Unauthorized Access Stream” for Alert of stream
 - b. For this example select “Message Count Alert Condition”

Condition

Define the condition to evaluate when triggering a new alert.

Alert on stream

SSH Unauthorized Access Stream

Select the stream that the condition will use to trigger alerts.

Condition type

Message Count Alert Condition

Select the condition type that will be used.

Add alert condition

- c. Select “Add alert condition”
5. Field content for rule

- Enter "SSH Unauthorized Access Alert" for title
- Enter "5" for time range events occur in
- Select "more than" for threshold type
- Enter "3" for threshold
- Enter "15" for grace period

Create new Message Count Alert Condition ✕

Message Count Alert Condition description

This condition is triggered when the number of messages is higher/lower than a defined threshold in a given time range.

Title

The alert condition title

Time Range

Evaluate the condition for all messages received in the given number of minutes

Threshold Type

Select condition to trigger alert: when there are more or less messages than the threshold

Threshold

Value which triggers an alert if crossed

Grace Period

Number of minutes to wait after an alert is resolved, to trigger another alert

Message Backlog

The number of messages to be included in alert notifications

- Select "Save"

6. Select "Alert" tab

graylog
In 0 / Out 0 msg/s Help Administrator

Search
Streams
Alerts
Dashboards
Sources
System

Alerts overview Manage conditions Manage notifications

Alerts are triggered when conditions you define are satisfied. Graylog will automatically mark alerts as resolved once the status of your conditions change.

[Read more about alerting in the documentation.](#)

Alerts Refresh Show unresolved alerts

Check your alerts status from here. Currently displaying all alerts.

SSH Unauthorized Access Alert on stream *SSH Unauthorized Access Stream* Unresolved

Triggered at 2017-03-24 14:12:10, **still ongoing**.

Reason: Stream had 6 messages in the last 5 minutes with trigger condition more than 3 messages. (Current grace time: 15 minutes)

Type: Message Count Alert Condition

« < 1 > »