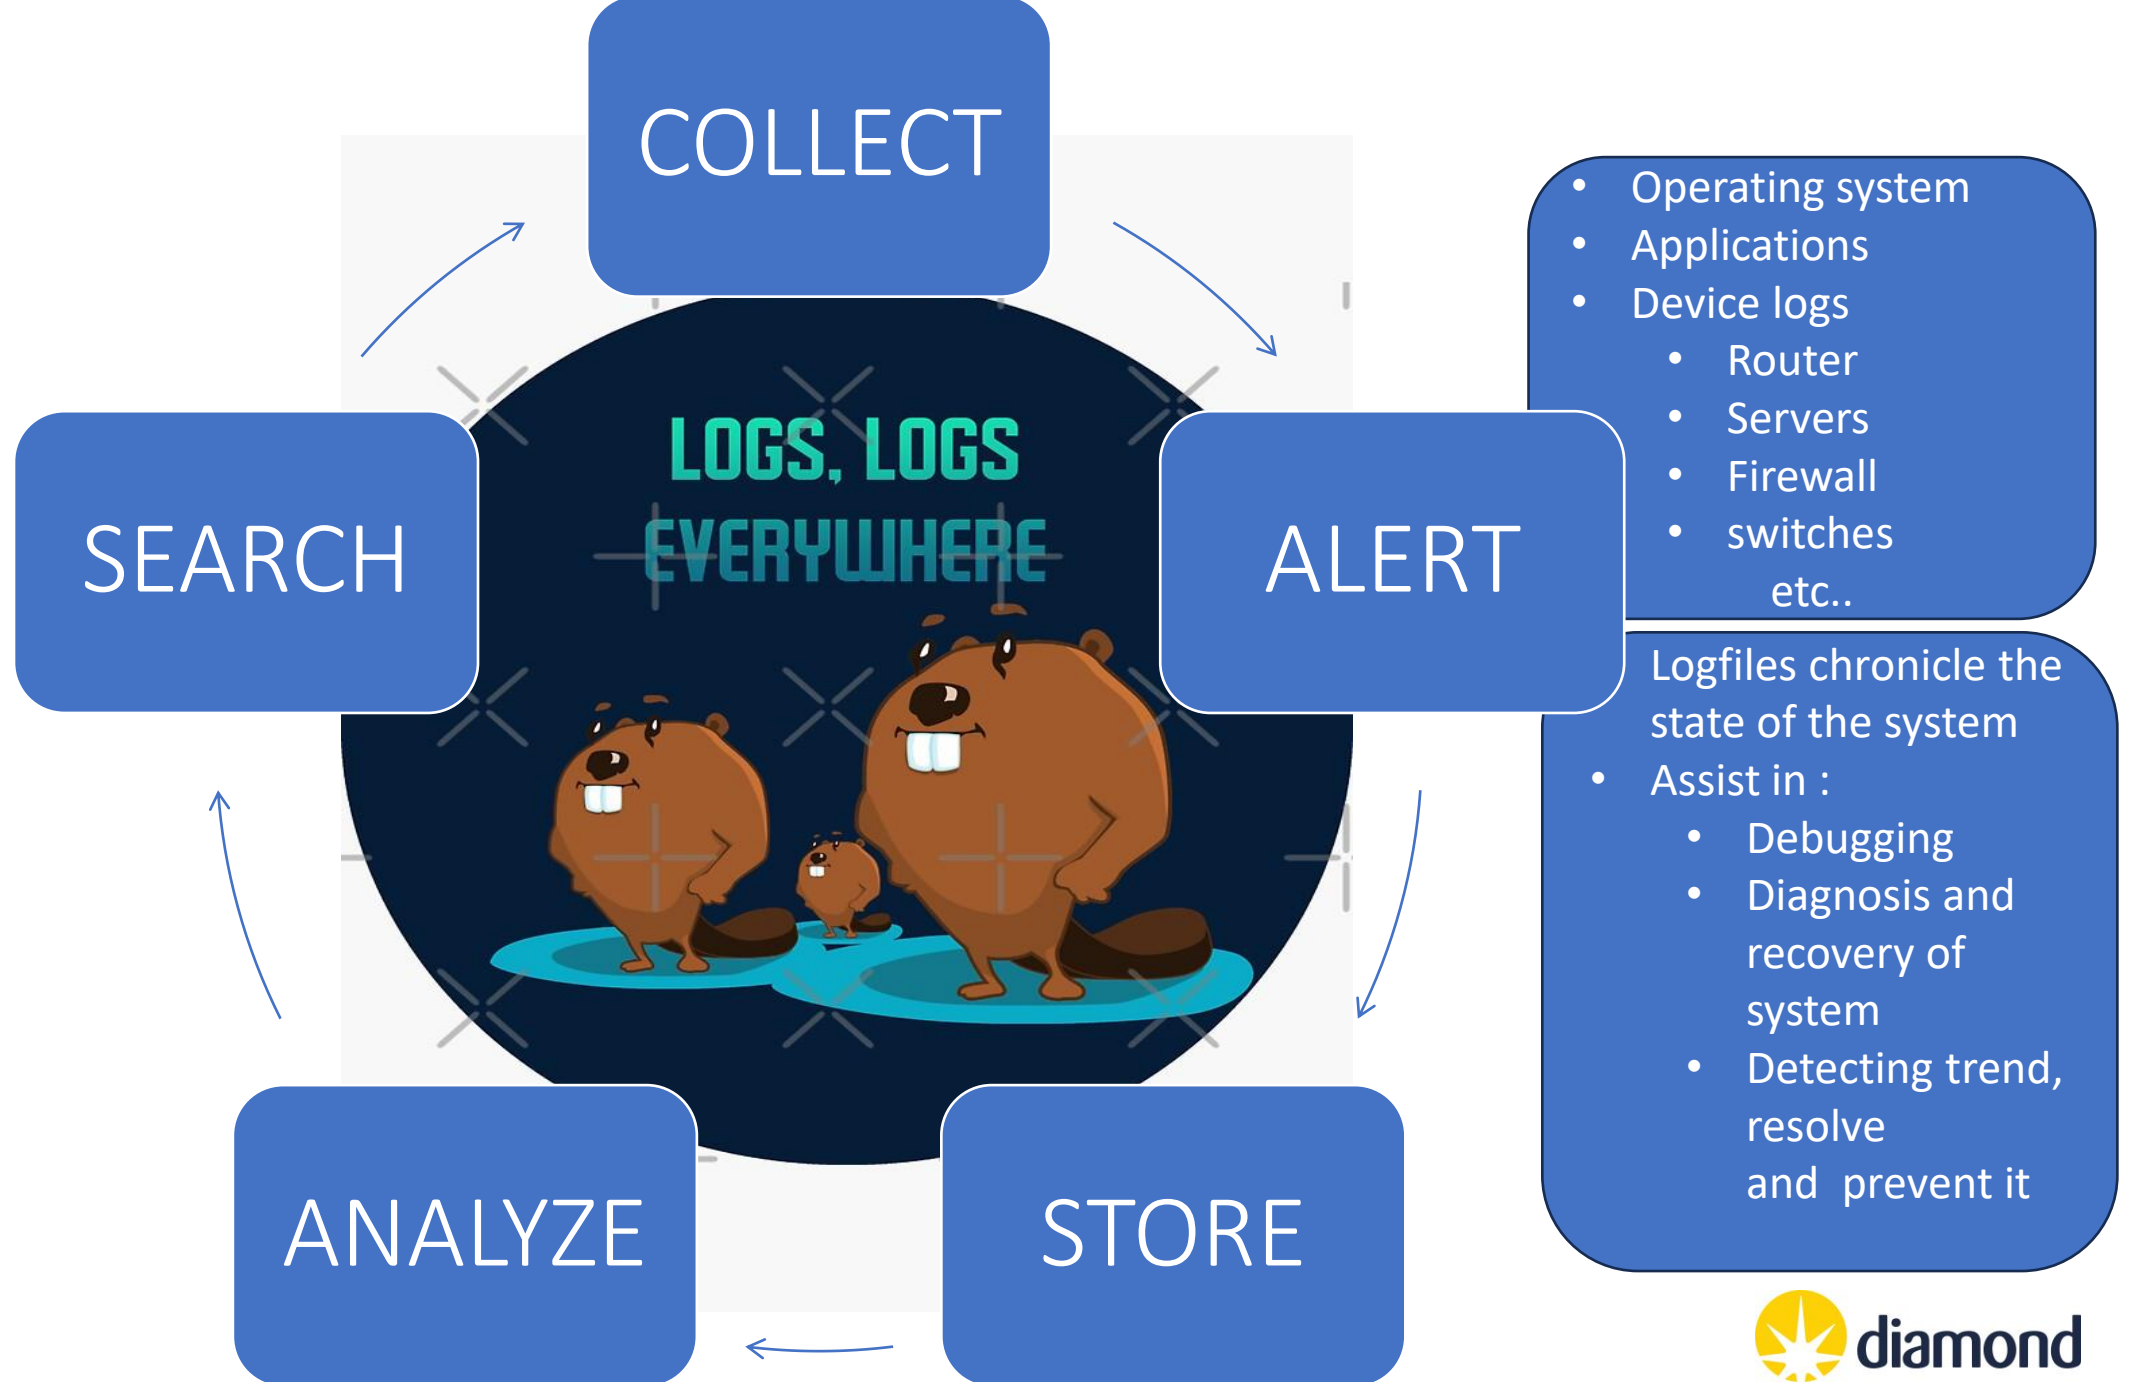


Graylog-as-a-Service: using Kubernetes to rescue a service from ancient hardware

Sonia Taneja, James Thorne

sonia.taneja@diamond.ac.uk, james.thorne@diamond.ac.uk



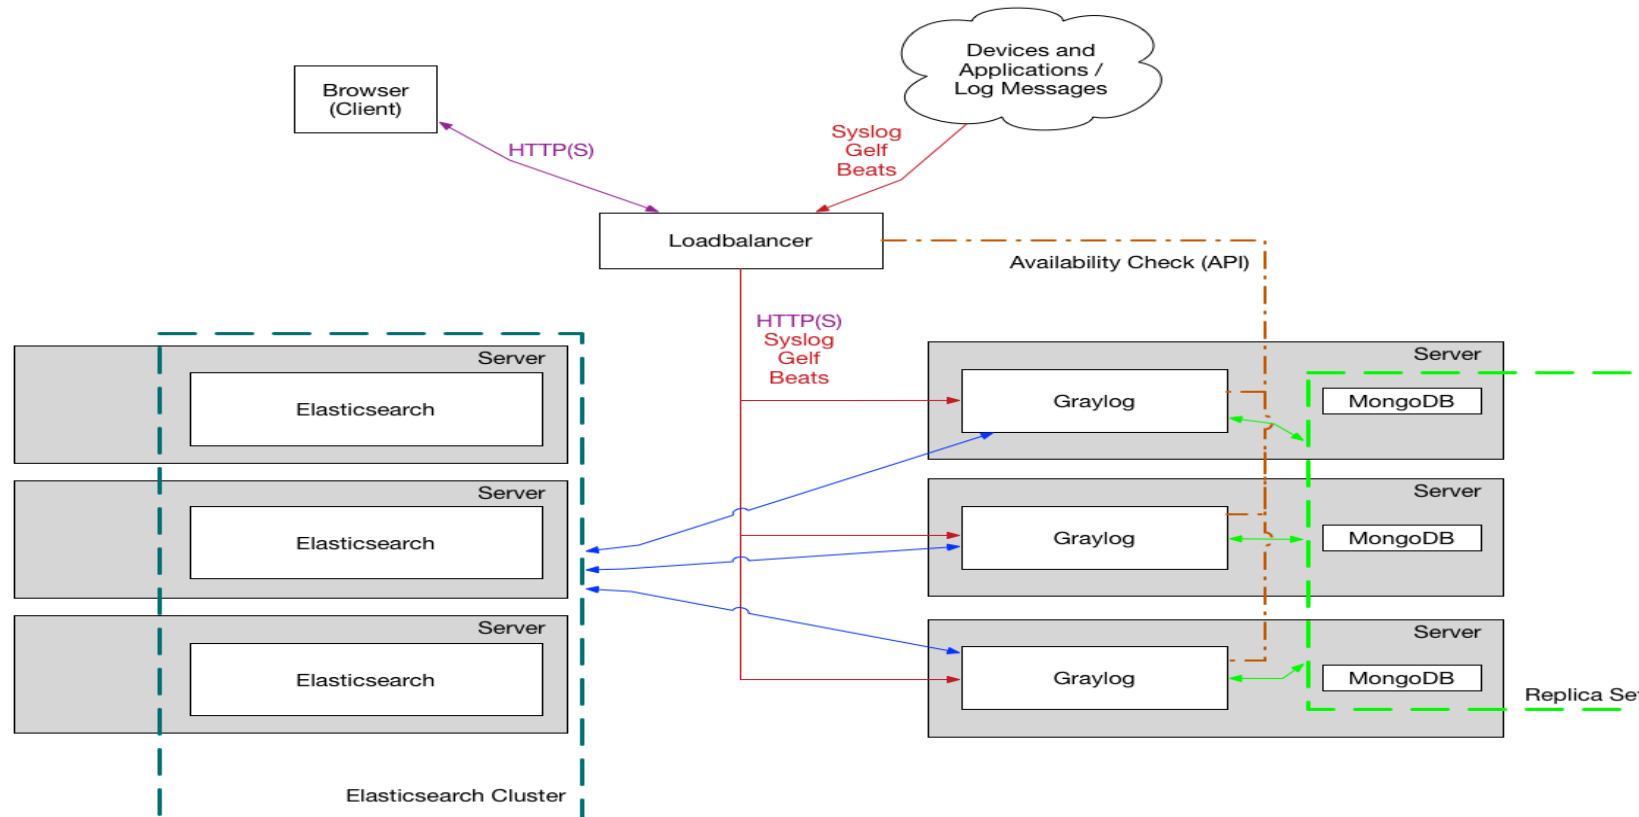
Centralized log management system



- Centralized logging consolidates log files from all systems across on-premises, cloud, and hybrid environments by enabling the following:
 - Collection and aggregation
 - Parsing and normalization
 - Correlation and analysis
 - aggregate, organize, and interpret data so that you can derive meaningful insights while maintaining data integrity
- Open-source solutions around
 - ELK, Graylog, Fluentd, SigNoz, Syslog-ng, Logwatch, Apache-flume

Graylog

- Graylog architecture basically comprises of Graylog, MongoDB and Elasticsearch
 - **Graylog Server** component sits in the middle and works around the data node i.e., elastic nodes
 - **MongoDB** to store meta information and configuration data
 - **Elasticsearch** or Opensearch are the index and search server stores all the log data



What Graylog does - Input

- Receives messages from multiple input protocols like GELF via HTTP/TCP/UDP, syslog, Beats, CEF, Netflow, Apache, kafka etc

graylog Search Views Streams Alerts Dashboards Sources Enterprise System / Inputs

Inputs

Graylog nodes accept data via inputs. Launch or terminate as many inputs as you want here.

Select input Launch new input Find more inputs

- Random HTTP message generator
- Raw/Plaintext AMQP
- Raw/Plaintext Kafka
- Raw/Plaintext TCP
- Raw/Plaintext UDP
- Syslog AMQP
- Syslog Kafka
- Syslog TCP
- Syslog UDP

What Graylog does - Input

- Receives messages from multiple input protocols like GELF via HTTP/TCP/UDP, syslog, Beats, CEF, Netflow, Apache, kafka etc

Launch new *Syslog UDP* input

Global
Should this input start on all nodes

Title
ws355
Select a name of your new input that describes it.

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

Port
5140
Port to listen on.

Receive Buffer Size (optional)
262144
The size in bytes of the recvBufferSize for network connections to this input.

No. of worker threads (optional)
10
Number of worker threads processing network connections for this input.

No. of worker threads (optional)
10
Number of worker threads processing network connections for this input.

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.

Encoding (optional)
UTF-8
Default encoding is UTF-8. Set this to a standard charset name if you want to override the default.

Force rDNS?
Force rDNS resolution of hostname? Use if hostname cannot be parsed. (Be careful if you are sending DNS logs into this input because it can cause a feedback loop.)

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

Store full message?
Store the full original syslog message as full_message?

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

Cancel Launch Input



What Graylog does - Input

- Receives messages from multiple input protocols like GELF via HTTP/TCP/UDP, syslog, Beats, CEF, Netflow, Apache, kafka etc

Filter by title Filter Reset

Global inputs 1 configured

ws355 Syslog UDP 2 RUNNING

Show received messages Manage extractors Stop input More actions ▾

```
allow_override_date: true
bind_address: 0.0.0.0
charset_name: UTF-8
expand_structured_data: false
force_rdns: false
number_worker_threads: 10
override_source: <empty>
port: 5140
recv_buffer_size: 262144
store_full_message: false
```

Throughput / Metrics

1 minute average rate: 0 msg/s

Network IO: ▼0B ▲0B (total: ▼6.5KiB ▲0B)

Empty messages discarded: 0

[Show details](#)



What Graylog does – Index sets

- Stores messages in Elasticsearch index sets for graphing

Create Index Set

Create a new index set that will let you configure the retention, sharding, and replication of messages coming from one or more streams.

Title	<input type="text" value="ws355"/>
Description	<input type="text" value="Sonia's workstation syslog"/>
Index prefix	<input type="text" value="graylog-test-ws355"/>
Analyzer	<input type="text" value="standard"/>
Index shards	<input type="text" value="4"/>
Index replicas	<input type="text" value="0"/>
Max. number of segments	<input type="text" value="1"/>

Descriptive name of the index set.

Add a description of this index set.

A **unique** prefix used in Elasticsearch indices belonging to this index set. The prefix must start with a letter or number, and can only contain letters, numbers, '_', '-' and '+'.
Elasticsearch analyzer for this index set.

Number of Elasticsearch shards used per index in this index set.

Number of Elasticsearch replicas used per index in this index set.

Maximum number of segments per Elasticsearch index after optimization (force merge).

Index optimization after rotation

Disable index optimization after rotation

Disable Elasticsearch index optimization (force merge) after rotation.

Field type refresh interval

seconds

How often the field type information for the active write index will be updated.

Index Rotation Configuration

Graylog uses multiple indices to store documents in. You can configure the strategy it uses to determine when to rotate the currently active write index.

Select rotation strategy

Max size per index (in bytes)

1.0GiB

Maximum size of an index before it gets rotated

Index Retention Configuration

Graylog uses a retention strategy to clean up old indices.

Select retention strategy

Max number of indices

Maximum number of indices to keep before **deleting** the oldest ones



What Graylog does - Streams

- Assigns messages to stream

Create stream

Title
ws355

Description
syslog

Index Set
ws355

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

Remove matches from 'Default Stream'
Don't assign messages that match this stream to the 'Default Stream'.

Cancel Create stream

New Stream Rule

Type
match input

Value
ws355 (Syslog UDP)

Inverted

Description (optional)

Result: *gl_source_input* must match input ws355 (Syslog UDP: 642c0736a6a5bf53650ac73a)

The server will try to convert to strings or numbers based on the matcher type as well as it can.

[Take a look at the matcher code on GitHub](#)

Regular expressions use Java syntax.

Cancel Create Rule

What Graylog does - Streams

- Assigns messages to stream

? Search Reset Show 10

All events index set Graylog Events
Stream containing all events created by Graylog
No configured rules.

Pause Stream Manage Rules Share More Actions

All system events index set Graylog System Events
Stream containing all system events created by Graylog
No configured rules.

Pause Stream Manage Rules Share More Actions

Default Stream index set Default index set Default
Contains messages that are not explicitly routed to other streams
The default stream cannot have rules.

Pause Stream Manage Rules Share More Actions

ws355 index set ws355 stopped
syslog
No configured rules.

Start Stream Manage Rules Share More Actions

What Graylog does - Streams

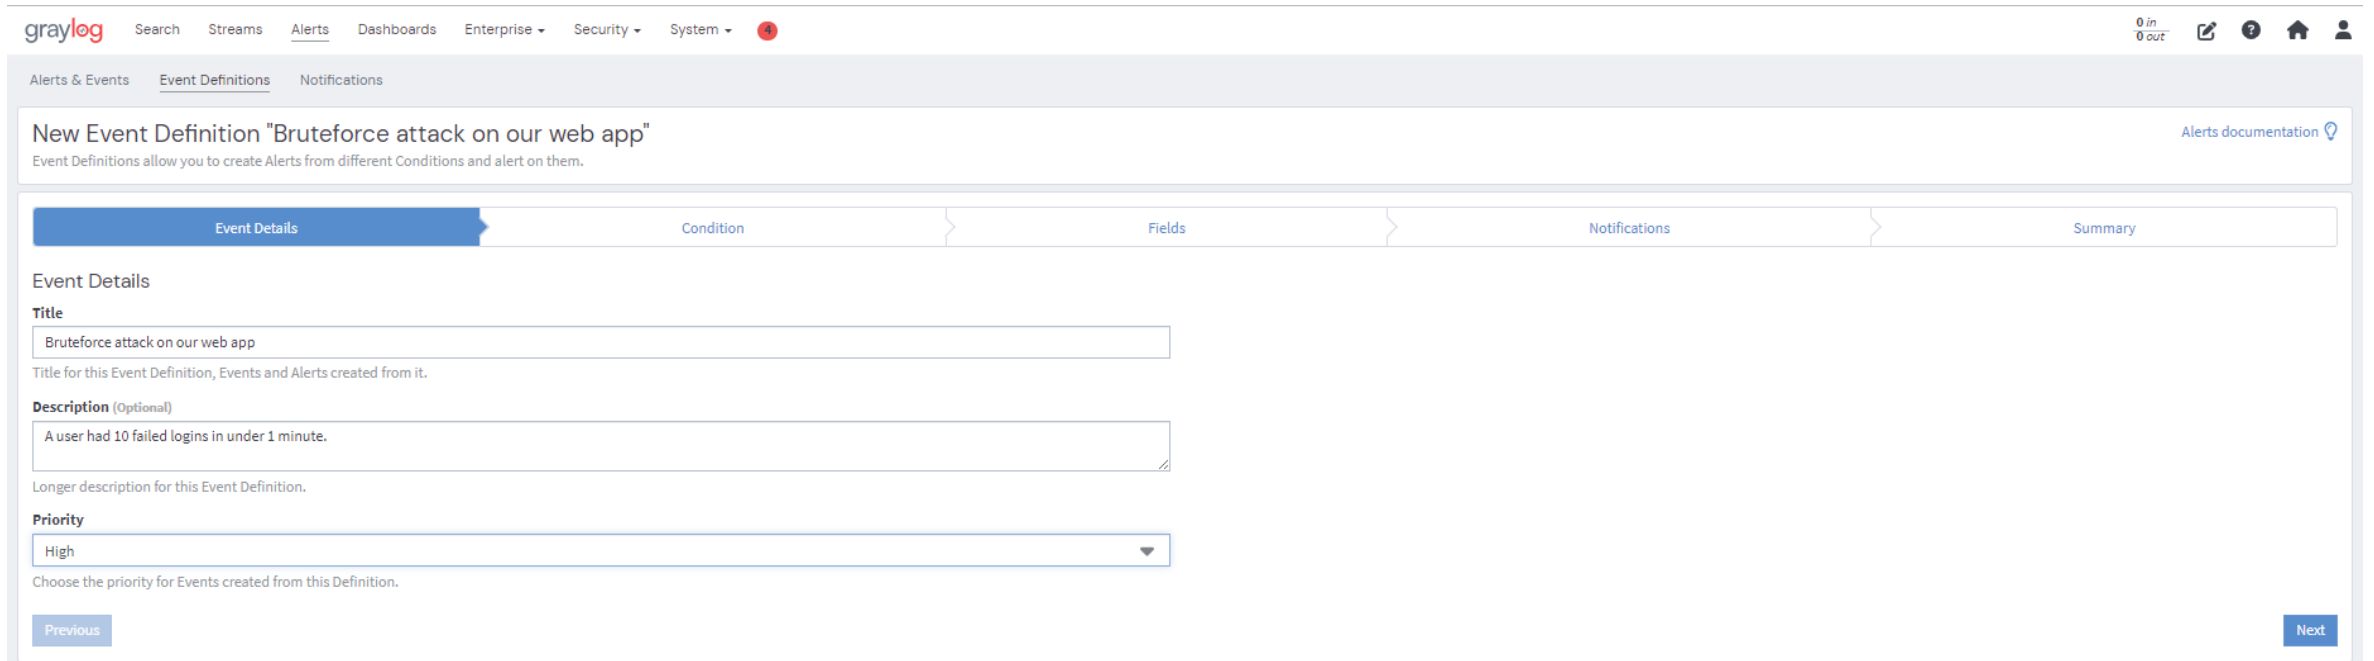
- Assigns messages to stream

The screenshot displays the Graylog search interface. At the top, there's a search bar with the query 'ws355' and a time range from '5 minutes ago' to 'Now'. Below the search bar is a 'Message Count' bar chart showing message volume over time. The chart has a y-axis from 0 to 10 and an x-axis from 12:10:00 to 12:14:30 on Apr 4, 2023. Three bars are visible: one at 12:13:30 with a count of 6, one at 12:14:00 with a count of 10, and one at 12:14:30 with a count of 8. Below the chart is the 'All Messages' section, showing a list of messages. The selected message is from 'ws355' at '2023-04-04 12:14:23.792' with the message text 'Started Hostname Service.'. The message details are shown in a table format.

Timestamp	application_name
2023-04-04 12:14:23.792	systemd
Received by	facility
ws355 on 3674befb / graylog-test-0.graylog-test.graylog-test.svc.cluster.local	system daemon
Stored in index	facility_num
graylog-test-ws355_0	3
Routed into streams	level
<ul style="list-style-type: none">ws355	6
message	Started Hostname Service.
process_id	1
source	ws355
timestamp	2023-04-04 12:14:23.792

What Graylog does - Alerts

- Triggers user-defined alerts per stream



The screenshot displays the Graylog web interface for creating a new event definition. The top navigation bar includes 'graylog', 'Search', 'Streams', 'Alerts', 'Dashboards', 'Enterprise', 'Security', and 'System'. A notification in the top right corner shows '0 in 0 out' and icons for help, home, and user profile.

The main content area is titled 'New Event Definition "Bruteforce attack on our web app"'. Below the title is a breadcrumb trail: 'Alerts & Events > Event Definitions > Notifications'. A sub-header explains: 'Event Definitions allow you to create Alerts from different Conditions and alert on them.' A link for 'Alerts documentation' is visible in the top right of this section.

The configuration is presented in a multi-step wizard with five tabs: 'Event Details', 'Condition', 'Fields', 'Notifications', and 'Summary'. The 'Event Details' tab is currently active and contains the following fields:

- Title:** A text input field containing 'Bruteforce attack on our web app'. Below it, a note states: 'Title for this Event Definition, Events and Alerts created from it.'
- Description (Optional):** A text area containing 'A user had 10 failed logins in under 1 minute.' Below it, a note states: 'Longer description for this Event Definition.'
- Priority:** A dropdown menu currently set to 'High'. Below it, a note states: 'Choose the priority for Events created from this Definition.'

Navigation buttons 'Previous' and 'Next' are located at the bottom left and right of the configuration area, respectively.

What Graylog does - Alerts

- Triggers user-defined alerts per stream

Event Details | **Filter & Aggregation** | Fields | Notifications | Summary

Event Condition

Configure how Graylog should create Events of this kind. You can later use those Events as input on other Conditions, making it possible to build powerful Conditions based on others.

Condition Type

Filter & Aggregation

Choose the type of Condition for this Event.

Filter

Add information to filter the log messages that are relevant for this Event Definition.

Search Query

Login failed for user

Search query that Messages should match. You can use the same syntax as in the Search page, including declaring Query Parameters from Lookup Tables by using the `$newParameter$` syntax.

Streams (Optional)

All messages x

Select streams the search should include. Searches in all streams if empty.

Search within the last

1 minutes

Execute search every

10 seconds

Enable

Should this event definition be executed automatically?

Available Conditions

Filter & Aggregation
Create Events from log messages by filtering them and (optionally) aggregating their results to match a given condition. These Events can be used as input for a Correlation Rule.

Event Correlation
Correlate previously defined Events to identify meaningful incidents. This will create new Events that you can later use.

How many Events will Filter & Aggregation create?

Filter Preview

Timestamp	Message
2022-11-28T18:20:26.464Z	2022-11-28T18:20:26.464Z GET /login [200] 59ms
2022-11-28T18:20:26.329Z	2022-11-28T18:20:26.329Z GET /login [200] 41ms
2022-11-28T18:20:26.301Z	2022-11-28T18:20:26.301Z GET /login [200] 45ms
2022-11-28T18:20:25.977Z	2022-11-28T18:20:25.977Z GET /login [200] 65ms
2022-11-28T18:20:25.843Z	2022-11-28T18:20:25.843Z GET /login [200] 37ms
2022-11-28T18:20:25.316Z	2022-11-28T18:20:25.316Z GET /login [200] 53ms
2022-11-28T18:20:25.261Z	2022-11-28T18:20:25.261Z GET /login [200] 39ms
2022-11-28T18:20:25.166Z	2022-11-28T18:20:25.166Z GET /login [200] 64ms
2022-11-28T18:20:25.118Z	2022-11-28T18:20:25.118Z GET /login [200] 60ms
2022-11-28T18:20:25.009Z	2022-11-28T18:20:25.009Z POST /login [201] 131ms



What Graylog does - Alerts

- Triggers user-defined alerts per stream

Event Details | Filter & Aggregation | Fields | **Notifications** | Summary

Notifications (optional) [Manage Notifications](#)

Is this Event important enough that requires your attention? Make it an Alert by adding Notifications to it.

Notification	Type	Actions
Email Notification Toronto EST	Email Notification	Remove from Event

[Add Notification](#)

[Previous](#) [Next](#)

Notification Settings

Grace Period
 5 minutes

Graylog sends Notifications for Alerts every time they occur. Set a Grace Period to control how long Graylog should wait before sending Notifications again. Note that Events with keys will have a Grace Period for each different key value.

Message Backlog
 50

Number of messages to be included in Notifications.

What Graylog does - Alerts

- Triggers user-defined alerts per stream

Event Details | Filter & Aggregation | Fields | Notifications | **Summary**

Event Summary

Details

Title
Bruteforce attack on our webb app

Description
A user had 10 failed logins in under 1 minute.

Priority
High

Filter & Aggregation

Type
Aggregation

Search Query
Login failed for user

Streams
[All messages](#)

Search within
1 minutes

Execute search every
10 seconds

Enable scheduling
no

Group by Field(s)
user

Create Events if
`count() >= 10`

Notifications

Settings
Grace Period is set to 5 minutes
Notifications will not include any messages.

Email Notification Toronto EST
Email Notification
[More detail](#)

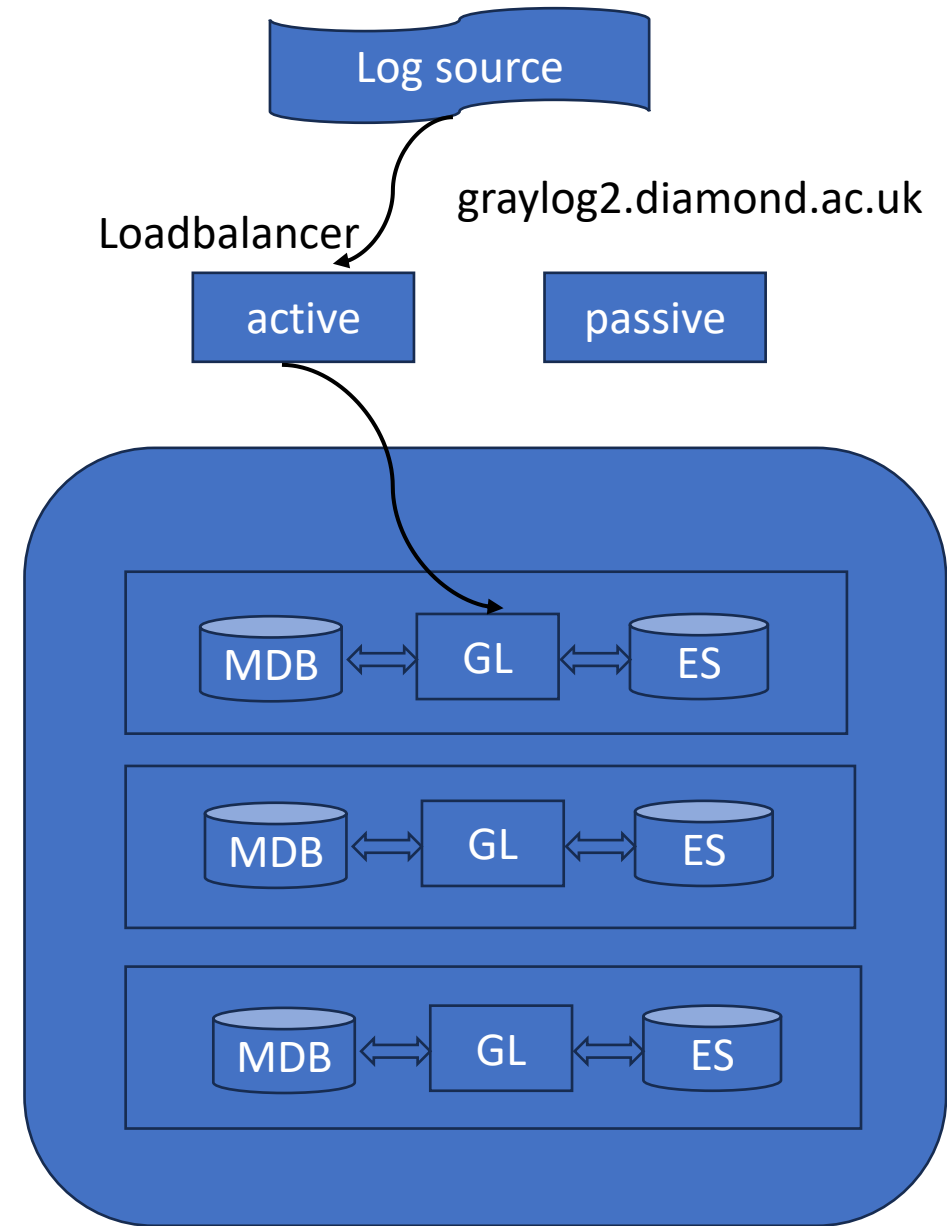
Cancel Create event definition

What Graylog does

- Receives messages from multiple input protocols like GELF via HTTP/TCP/UDP, syslog, Beats, CEF, Netflow, Apache, kafka etc
- Stores messages in Elasticsearch index sets for graphing
- Assigns messages to stream
- Triggers user-defined alerts per stream
- Provides user friendly interface for search, alerting and analysis of data
- Powerful search capabilities - answers to complex queries in milliseconds
- Perform real-time analysis of terabytes of machine data
- The GUI has a range of graphs and widgets to clearly visualize log and event data

Graylog - In use at Diamond

- A single, monolithic, ancient version of Graylog running on an ancient hardware, cluster comprising of 3 servers each with:
 - CPU: 24
 - Memory: 192GB
 - Disk: 890GB
- Caters to the needs of users from various groups
- Upgrade to newer version
- User's requirements
 - Logging from various applications that goes onto file and Graylog as well
 - Powerful search capabilities
 - Possibility to define events and an Alert
 - Good retention period for logs
 - No interference from other groups
 - Upgraded permissions system enabling better access control



Solution - Graylog-as-a-Service on Kubernetes

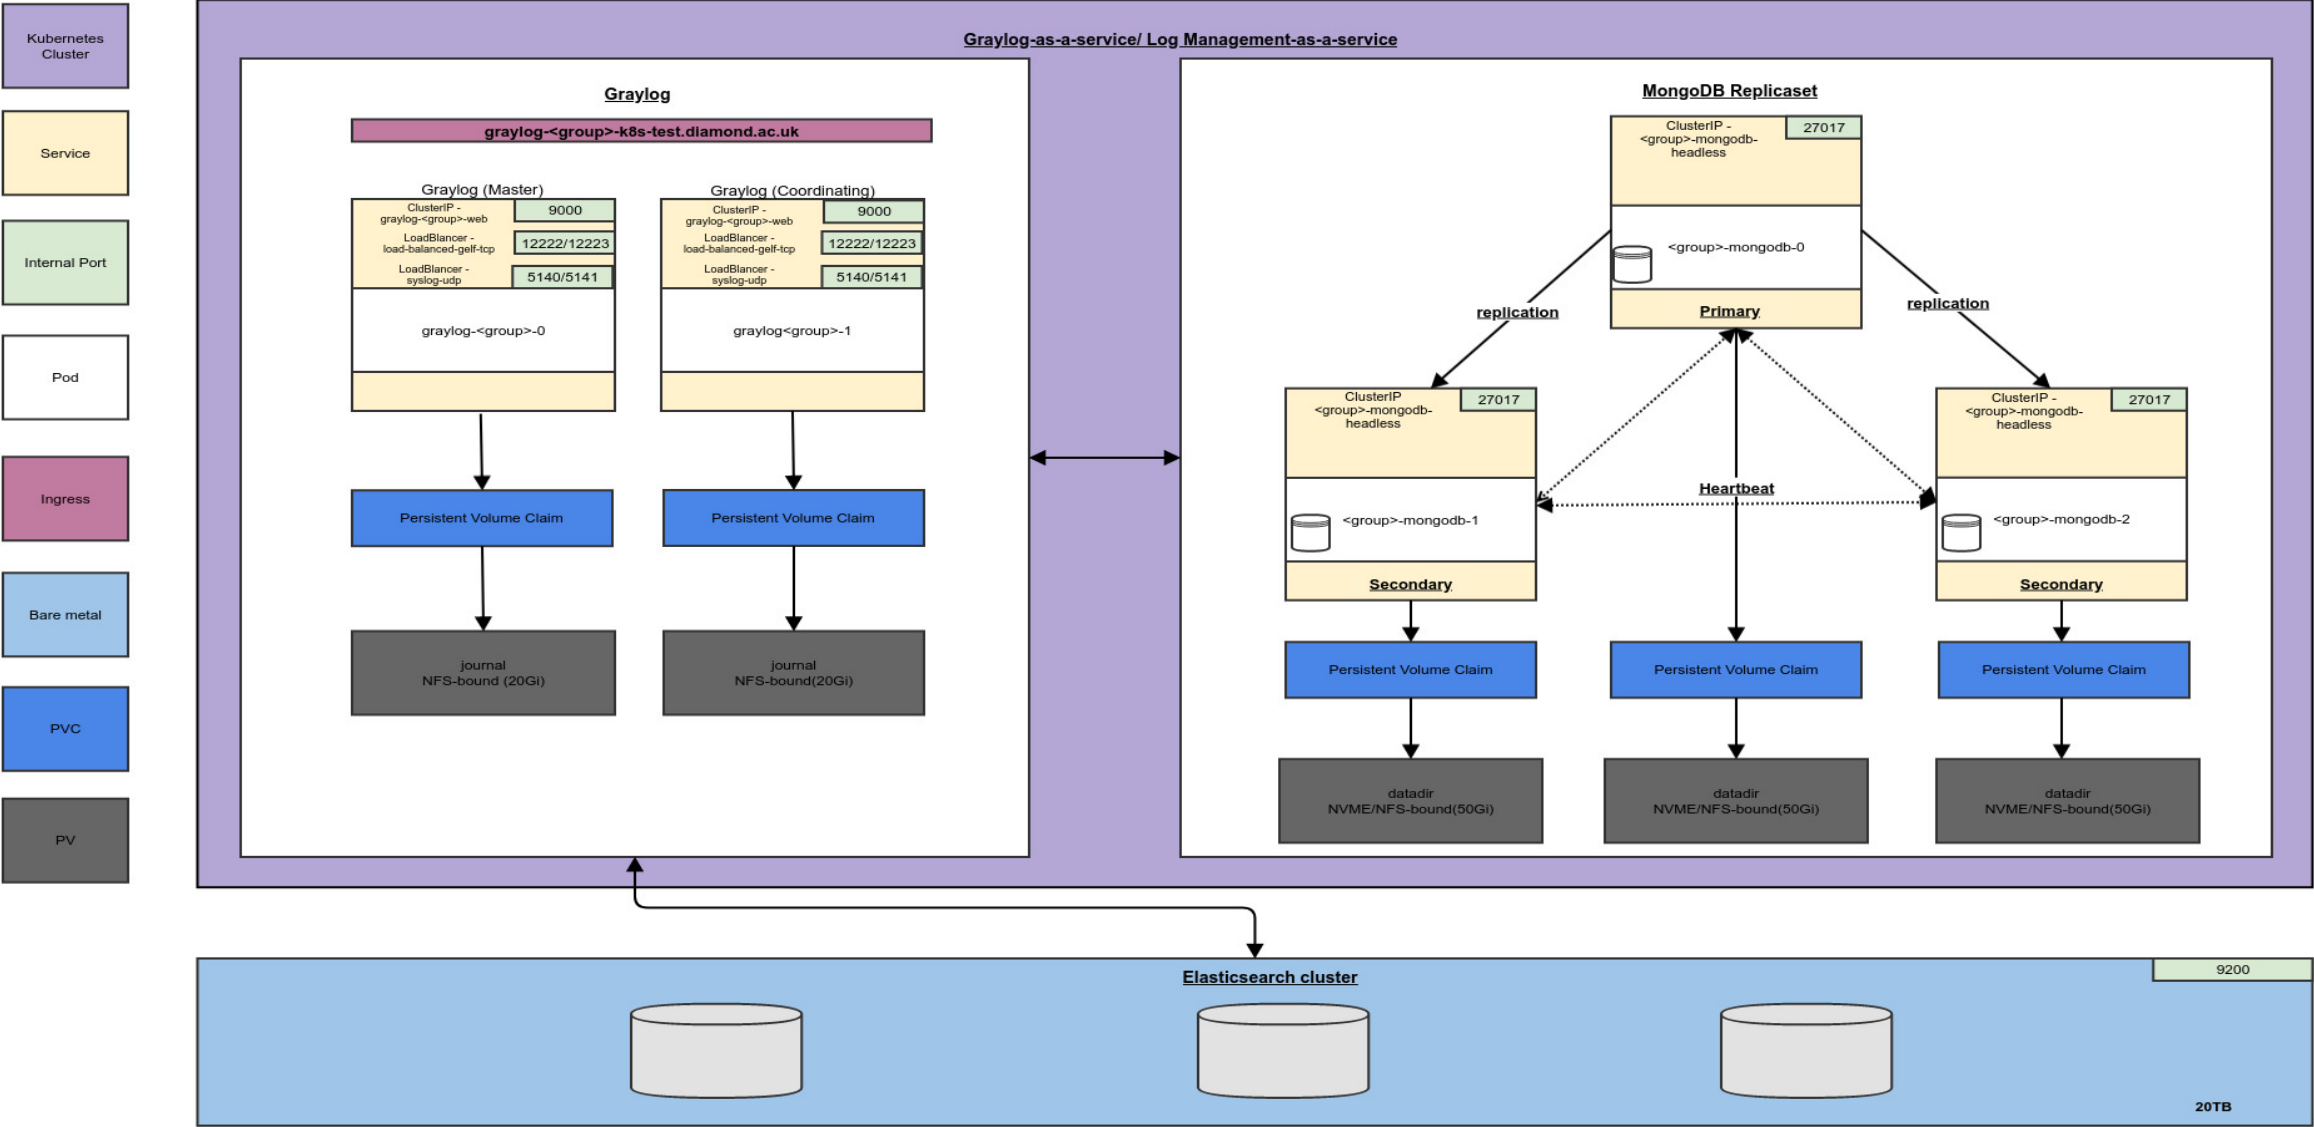
- Implemented Graylog architecture which is distributed across Cloud infrastructure and traditional physical Infrastructure
- Instances per group
- Advantages:
 - Better resource allocation and isolation
 - No interference from other groups
 - Better access control setup
 - Simplified deployment and management
 - Easy to scale the application horizontally
 - Built-in high availability and self-healing capability from Kubernetes
 - Load balanced and better monitoring

Graylog-as-a-service on Kubernetes

- Group namespaces on Kubernetes production cluster
 - >> `kubectl describe namespace <group-namespace-name>`
- Installed using helm charts
 - Graylog – kongz/graylog
 - >> `helm repo add kongz https://charts.kong-z.com`
 - >> `helm install <releaseName> kongz/graylog -n <group-namespace-name> -f values.yaml`
 - MongoDB – bitnami/mongodb
 - >> `helm repo add bitnami https://charts.bitnami.com/bitnami`
 - >> `helm install <releaseName> bitnami/mongodb -n <group-namespace-name> -f values.yaml`
- Customized manifest file – values.yaml
- ConfigMap for keystore
 - >> `kubectl create configmap graylog-keystore --namespace <group-namespace-name> --from-file=<location_of_cacerts.jks>`
 - >> `kubectl get cm graylog-keystore -o yaml|less`
- Secret/SealedSecret
 - MongoDB – one SealedSecret
 - Graylog – three SealedSecret
 - `mongodb://username:password@POD1_NAME.SERVICE_NAME.NAMESPACE_NAME.svc.cluster.local:27017,POD2_NAME.SERVICE_NAME.NAMESPACE_NAME.svc.cluster.local:27017,POD3_NAME.SERVICE_NAME.NAMESPACE_NAME.svc.cluster.local:27017/DEFAULT_AUTH_DB?replicaSet=REPLICASET_NAME"`
 - `https://<elastic_user>:<password>@node1:9200,https://<elastic_user>:<password>@node2:9200,https://<elastic_user>:<password>@node3:9200`
- Each group namespace has two Statefulsets
 - Graylog – Statefulset with 2 replicas
 - MongoDB – Statefulset with 3 replicas
- NVME-based PersistentVolumeClaim and podAntiAffinityPreset to hard



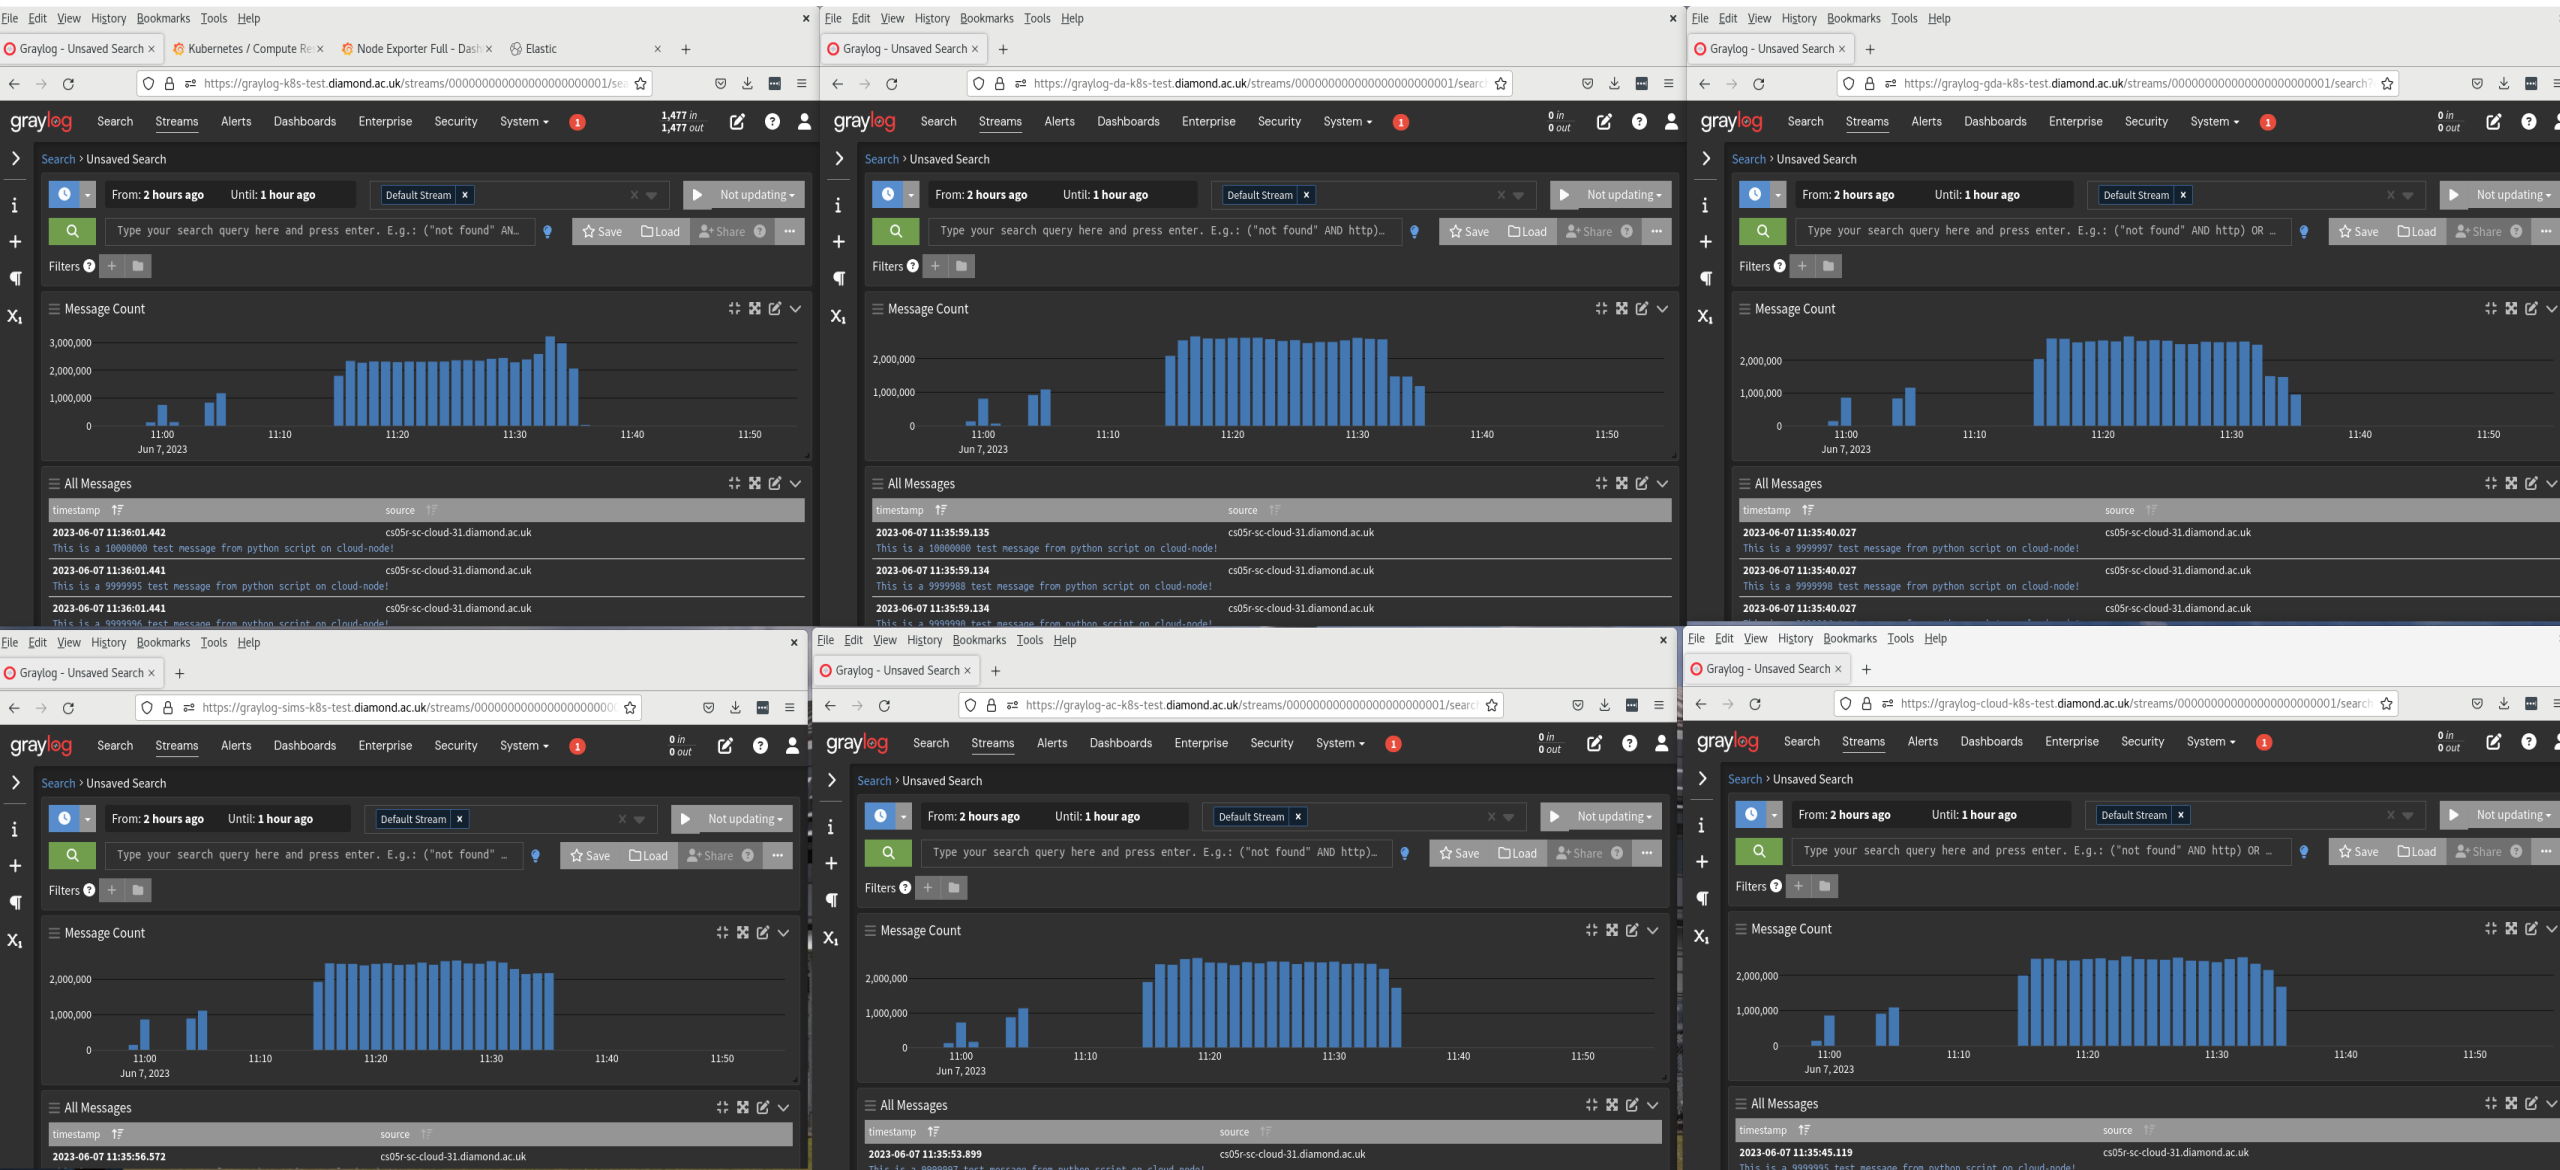
Diagram



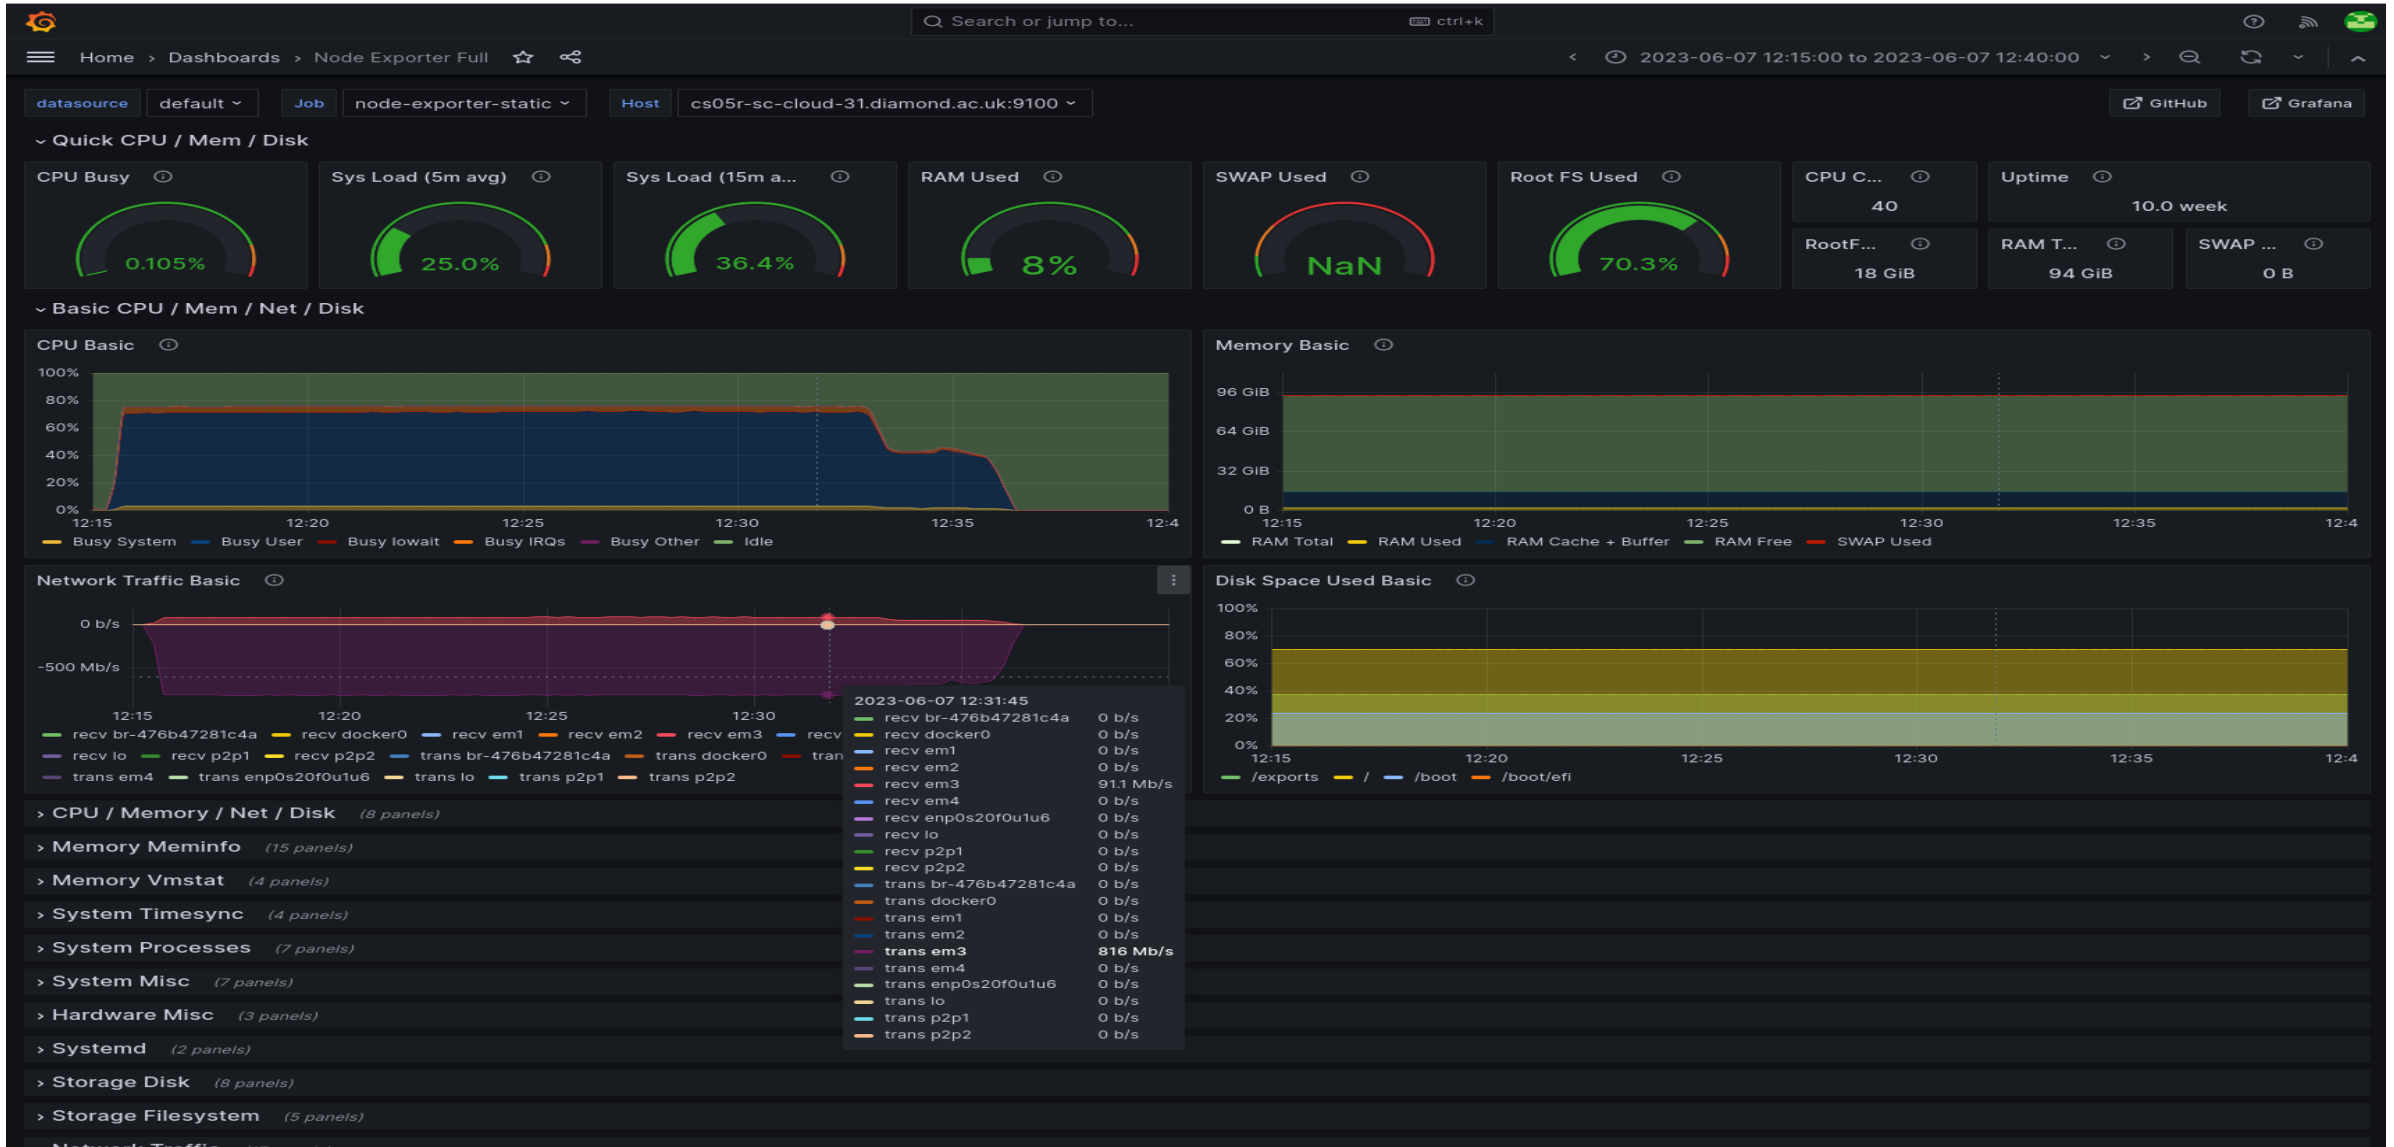
Benchmarking

- Purpose: Benchmarking on new Graylog instances on Kubernetes by sending large number of messages to all the instances concurrently
- How: Bash wrapper script around python script which uses graypy
- 5 concurrent iteration of 10 million messages to each Graylog instance
- Results: this architectural design survived 2.5 million messages/minute on each instance (total 6). We maxed-out the Network transfer speed of 816 Mb/s on node which has link of 10GB

Benchmarking



Benchmarking



Final thoughts

- Log management system in place is a necessity
- Graylog offers solution that can capture, store, and perform real-time analysis of terabytes of machine data
- Migrated from single monolithic system to a mixed model of Graylog-as-a-service on Kubernetes
- From the Graylog user perspective instances owned per group compared to a central instance has its own benefits in terms of better resource allocation and isolation and no interference from other group.
- Reliable, high available, performant and upgraded permission system with an admin domain
- Soon going into production



Thank you!

Questions?

