



Monitoring Cisco NFVI Performance

The following topics tell you how to display logs to monitor Cisco VIM performance.

- [Logging and Monitoring in Cisco NFVI, on page 1](#)
- [Displaying Cisco VIM Log Files Using the CLI, on page 3](#)
- [Logging Into the Kibana Dashboard, on page 6](#)
- [Rotation of the Cisco VIM Logs, on page 16](#)
- [Snapshot Manager Tool for Elasticsearch, on page 16](#)
- [Remote NFS Backup for Elasticsearch Snapshots, on page 18](#)
- [Network Performance Test with NFVBench, on page 18](#)

Logging and Monitoring in Cisco NFVI

Cisco VIM uses a combination of open source tools to collect and monitor the Cisco OpenStack services including Elasticsearch, Fluentd, and the Kibana dashboard (EFK).

In VIM, we have moved our platform to use Fluentd, instead of logstash. However, to maintain backwards compatibility, the code, and documentation refers to ELK, instead of EFK at various places. In VIM, these two acronyms are interchangeable, however it refers to the presence of EFK in the offering. OpenStack services that followed by EFK include:


- MariaDB—A relational database management system which is based on MySQL. All the OpenStack components store their data in MariaDB.
- HAProxy—HAProxy is a free open source software that provides a high-availability load balancer, and proxy server for TCP and HTTP-based applications that spreads requests across multiple servers.
- Keystone—Keystone is an OpenStack project that provides identity, token, catalog, and policy services for use specifically by projects in the OpenStack.
- Glance—An OpenStack project that allows you to upload and discover data assets that are meant for use with other services.
- Neutron—An OpenStack project that provides the network connectivity between interface devices, such as vNICs, managed by other OpenStack services, such as Nova.
- Nova—An OpenStack project that is designed to provide massively scalable, on demand, self-service access to compute resources.

- HTTP—The Apache HTTP server Project, an effort to develop and maintain an open-source HTTP server.
- Cinder—An OpenStack block storage service that is designed to present storage resources to the users that are consumed by the OpenStack compute project (Nova).
- Memcached—A general purpose distributed memory caching system.
- CloudPulse—Is an OpenStack tool that checks the health of the cloud. CloudPulse includes operator and end-point tests.
- Heat—The main OpenStack Orchestration program. Heat implements an orchestration engine to launch multiple composite cloud applications that is based on text file templates.
- Other OpenStack services—RabbitMQ, Ceph, Open vSwitch, Linux bridge, Neutron VTS (optional), and others.
- VMTP—Integrated control and data plane log for testing the cloud.
- NFVBench—Network performance benchmarking tool.

A Fluentd container resides on each control, compute, and storage nodes. They forward log to the Fluentd-aggr server residing on the management node.

The following figure shows a high-level schematic of the Fluent service assurance architecture.

Figure 1: EFK Service Assurance Architecture



The EFK flow includes:

- Fluentd extracts the relevant data from the logs and tags them so that Kibana can use it later to display useful information about those logs.
- Fluentd sends the logs from all the compute, controller, and storage nodes to the Fluentd-aggr server on the management node.
- Fluentd-aggr in the management node sends the structured logs into the Elasticsearch database.
- Elasticsearch stores the data, indexes it, and supports fast queries against a large amount of log data.
- Kibana visualizes the data that is stored in Elasticsearch using a custom dashboard. You can also add filters to the data to visualize interesting fragments of the log data.

Displaying Cisco VIM Log Files Using the CLI

Cisco VIM log file location depends on the node and log type. Installer logs are found in the management node under the `/var/log/mercury/<install_uid>/` directory. The last 20 log directories are tarred and kept in this directory. These files contain logs related to bootstrap, build orchestration, baremetal, common setup, and OpenStack orchestration.

If the installer fails, look at the last tar.gz file for logs, for example:

```
[root@mgmtnode mercury]# ls -lrt
total 20
drwxr-xr-x. 2 root root 80 Jul 19 23:42 573f2b7f-4463-4bfa-b57f-98a4a769aced
drwxr-xr-x. 2 root root 4096 Jul 20 03:29 installer
drwxr-xr-x. 2 root root 79 Jul 20 03:29 e9117bc5-544c-4bda-98d5-65bffa56a18f
drwxr-xr-x. 2 root root 79 Jul 20 04:54 36cdf8b5-7a35-4e7e-bb79-0cfb1987f550
drwxr-xr-x. 2 root root 79 Jul 20 04:55 bd739014-fdf1-494e-adc0-98b1fba510bc
drwxr-xr-x. 2 root root 79 Jul 20 04:55 e91c4a6c-ae92-4fef-8f7c-cafa9f5dc1a3
drwxr-xr-x. 2 root root 79 Jul 20 04:58 1962b2ba-ff15-47a6-b292-25b7fb84cd28
drwxr-xr-x. 2 root root 79 Jul 20 04:59 d881d453-f6a0-448e-8873-a7c51d8cc442
drwxr-xr-x. 2 root root 78 Jul 20 05:04 187a15b6-d425-46a8-a4a2-e78b65e008b6
drwxr-xr-x. 2 root root 4096 Jul 20 06:47 d0346cdd-5af6-4058-be86-1330f7ae09d1
drwxr-xr-x. 2 root root 79 Jul 20 17:09 f85c8c6c-32c9-44a8-b649-b63fdb11a79a
drwxr-xr-x. 2 root root 67 Jul 20 18:09 179ed182-17e4-4f1f-a44d-a3b6c16cf323
drwxr-xr-x. 2 root root 68 Jul 20 18:13 426cb05f-b1ee-43ce-862d-5bb4049cc957
drwxr-xr-x. 2 root root 68 Jul 20 18:13 1d2eec9d-f4d8-4325-9eb1-7d96d23e30fc
drwxr-xr-x. 2 root root 68 Jul 20 18:13 02f62a2f-3f59-46a7-9f5f-1656b8721512
drwxr-xr-x. 2 root root 68 Jul 20 18:14 c7417be9-473e-49da-b6d0-d1ab8fb4b1fc
drwxr-xr-x. 2 root root 68 Jul 20 18:17 b4d2077b-c7a9-46e7-9d39-d1281fba9baf
drwxr-xr-x. 2 root root 68 Jul 20 18:35 21972890-3d45-4642-b41d-c5fadfeba21a
drwxr-xr-x. 2 root root 80 Jul 20 19:17 d8b1b54c-7fc1-4ea6-83a5-0e56ff3b67a8
drwxr-xr-x. 2 root root 80 Jul 20 19:17 23a3cc35-4392-40bf-91e6-65c62d973753
drwxr-xr-x. 2 root root 80 Jul 20 19:17 7e831ef9-c932-4b89-8c81-33a45ad82b89
drwxr-xr-x. 2 root root 80 Jul 20 19:18 49ea0917-f9f4-4f5d-82d9-b86570a02dad
drwxr-xr-x. 2 root root 80 Jul 20 19:18 21589a61-5893-4e30-a70e-55ad0dc2e93f
drwxr-xr-x. 2 root root 80 Jul 20 19:22 6ae6d136-7f87-4fc8-92b8-64cd542495bf
drwxr-xr-x. 2 root root 4096 Jul 20 19:46 1c6f4547-c57d-4dcc-a405-ec509306ee25
drwxr-xr-x. 2 root root 68 Jul 20 21:20 c6dcc98d-b45b-4904-a217-d25001275c85
drwxr-xr-x. 2 root root 68 Jul 20 21:40 ee58d5d6-8b61-4431-9f7f-8cab2c331637
drwxr-xr-x. 2 root root 4096 Jul 20 22:06 243cb0f8-5169-430d-a5d8-48008a00d5c7
drwxr-xr-x. 2 root root 4096 Jul 20 22:16 188d53da-f129-46d9-87b7-c876b1aea70c
```

Cisco VIM autobackup logs are found in the following location:

```
# CVIM autobackup logs (auto-backup enabled by default)
/var/log/mercury/autobackup_2.2.x_2018-03-19_15-11-10.log

# cobbler apache log (may be needed for PXE troubleshooting)
/var/log/cobblerhttpd/access_log
/var/log/cobblerhttpd/error_log

# VMTP logs
/var/log/vmtp/vmtp.log
```

Cisco VIM RestAPI log location

```
# CVIM RestAPI logs
/var/log/mercury_restapi/restapi.log

# CIM RestAPI apache logs (TCP port 8445)
/var/log/httpd/mercury_access.log
/var/log/httpd/mercury_error.log

# CIM RestAPI log-directory logs (TCP port 8008)
/var/log/httpd/access_log
/var/log/httpd/error_log
```

EFK log location

```
# Elasticsearch-fluentd-Kibana
/var/log/elasticsearch/
/var/log/fluentd-aggr/
/var/log/kibana/
/var/log/curator/
```

```
# HAProxy TLS certificate expiration check
/var/log/curator/certchecker.log
```

Viewing Cisco VIM Logs

```
# list logs sorted reverse on time
ls -lrt /var/log/mercury/
# untar logs
tar xvzf /var/log/mercury/<UUID>/mercury_install_2018-3-20_10-2.tar.gz -C /tmp/
```

Cisco VIM Configuration Files

```
# example configuration files
/root/openstack-configs/setup_data.yaml.B_Series_EXAMPLE
/root/openstack-configs/setup_data.yaml.C_Series_EXAMPLE

# system maintained setup files - do not modify directly
# always supply user copy of setup_data.yaml
# when using ciscovim client
/root/openstack-configs/setup_data.yaml

# system inventory in pretty format
/root/openstack-configs/mercury_servers_info

# passwords store
/root/openstack-configs/secrets.yaml

# openstack configuration file
/root/openstack-configs/openstack_config.yaml

# RestAPI password
/opt/cisco/ui_config.json

# Insight password
/opt/cisco/insight/secrets.yaml
```

Enabling debug logs for certain OpenStack Services

```
# openstack config file
/root/openstack-configs/openstack_config.yaml

# help
ciscovim help

# list openstack keys
ciscovim list-openstack-configs

# help on reconfigure sub-command
ciscovim help reconfigure

# how to execute subcommand, example below
# important note: reconfigure requires a maintenance window
ciscovim reconfigure --setopenstackconfig KEYSTONE_DEBUG_LOGGING,CINDER_DEBUG_LOGGING
```

On controller and compute nodes, all services are run within their respective Docker™ containers.

To list the Docker containers in the node, execute the following:

```
[root@control-server-2 ~]# docker ps -a
CONTAINER ID          IMAGE                                     PORTS                NAMES                COMMAND
258b2cald46a         172.31.228.164:5000/mercury-rhel7-osp8/nova-scheduler:4780
"/usr/bin/my_init /no" 25 minutes ago     Up 25 minutes       novascheduler_4780
ffe70809bbe0         172.31.228.164:5000/mercury-rhel7-osp8/nova-novncproxy:4780
"/usr/bin/my_init /st" 25 minutes ago     Up 25 minutes       novanovncproxy_4780
```

```
12b92bcb9dc0          172.31.228.164:5000/mercury-rhel7-osp8/nova-consoleauth:4780
"/usr/bin/my_init /st" 26 minutes ago Up 26 minutes
```

.....

```
novaconsoleauth_4780
7295596f5167          172.31.228.164:5000/mercury-rhel7-osp8/nova-api:4780
"/usr/bin/my_init /no" 27 minutes ago Up 27 minutes          novaapi_4780
```

To view the Docker logs of any container, execute the following on the corresponding host:

```
ls -l /var/log/<service_name>/<log_filename>
e.g. ls -l /var/log/keystone/keystone.log
```

To get into a specific container, execute the following commands:

```
[root@control-server-2 ~]# alias | grep container
root@control-server-2 ~]# source /root/.bashrc
#execute the alias:
[root@control-server-2 ~]# novaapi
novaapi_4761 [nova@control-server-2 /]$
novaapi_4761 [nova@control-server-2 /]$ exit
exit
```

If the Docker status indicates a container is down (based on output of “docker ps -a”), collect the Docker service logs as well:

```
cd /etc/systemd/system/multi-user.target.wants/
ls docker* # get the corresponding service name from the output
systemctl status <service_name> -n 1000 > /root/filename # redirects the output to the file
```

For storage nodes running Ceph, execute the following to check the cluster status:

```
ceph -v # on monitor nodes (controller), show's ceph version
ceph -s # on monitor nodes (controller), show cluster status
ceph osd lspools #on monitor nodes (controller),list pools
ceph mon stat # summarize monitor status
ceph-disk list # on OSD / storage nodes; List disks, partitions, and Ceph OSDs
rbd list images # on monitor nodes (controller); dump list of image snapshots
rbd list volumes # on monitor nodes (controller); dump list of volumes
```

Logging Into the Kibana Dashboard

Kibana is an open source data visualization platform that is used to explore Cisco VIM logs.

To log into the Kibana dashboard:

Step 1 Using a terminal client, use SSH to log into your management node and enter the password to login.

The following command shows the management node has an IP address of 17.0.0.2:

```
# ssh root@17.0.0.2
root@17.0.0.2's password
```

Step 2 In the SSH terminal session, locate the line containing KIBANA_PASSWORD in /root/installer-{tag-id}/openstack-configs/secrets.yaml. Note the value of the KIBANA_PASSWORD. It is used in Step 4.

```
cat /root/installer-{tag-id}/openstack-configs/secrets.yaml
...
KIBANA_PASSWORD: <note this value>
...
```

Step 3 Navigate to the `http://<management_node_ip_address>:5601`.

Note Kibana uses the HTTPS + TLS to provide a secure connection between the browser and the Kibana service.

By default Kibana uses the certificate located at `/var/www/mercury/mercury.<cert|key>` or you can provide your own certificates in `/root/openstack-configs/` directory (using the same `mercury.<cert|key>` file names).

Note If you are accessing Kibana for the first time, by default it shows self-signed certificate. Some browsers display the warning message *Your connection is not private*. Click **Proceed** to access the Kibana link. A window dialog box appears.

Step 4 Enter the Username and Password:

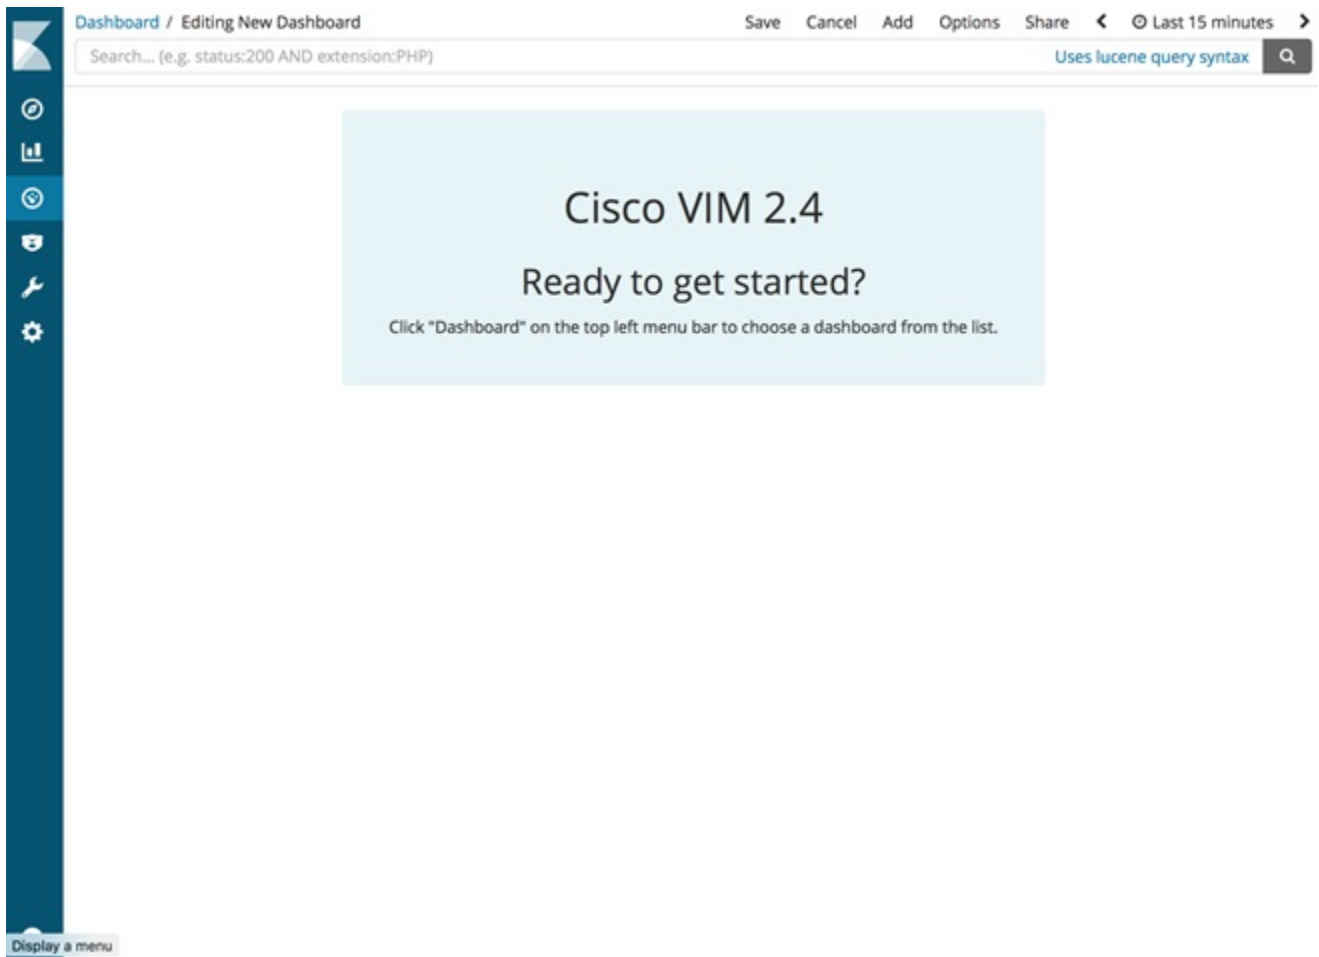
Authentication Required

Username

Password

User Name: admin


Password: <value of ELK_PASSWORD from Step 2>. The Kibana dashboard appears which displays the Cisco VIM service and installer logs.



Step 5 Click the Dashboard to choose the desired dashboard.

Note We recommend you not to use visualize/Timelion/DevTools or Management options on the left side.

Figure 2: Lists of Dashboards




The following are the list of dashboards:

- **Hostlogs Dashboard:** Provides log information of the system for the cloud nodes. This displays entries from the host logs-* index in Elasticsearch. It contains the log from /var/log/messages file on each server.
- **Installer Dashboard:** Provides information about the management node and the installation process. It can only read uncompressed files. Hence, it reads the files prior to the cloud installation. This displays entries from the installer-* index in Elasticsearch.
- **OpenStack Dashboard:** (openstack-* index) Provides log information about all the OpenStack processes. This displays entries from the openstack-* index in Elasticsearch.
- **VMTP Dashboard:** Provides log information about the VMTP runs performed against the cloud. It displays entries from the vmtp-* index in Elasticsearch.

For Example: if you click **OpenStack Dashboard** link the following screen appears.

Figure 3: OpenStack Dashboard



You can switch on from one dashboard to another by selecting the appropriate dashboard from the right top bar menu.

All dashboards have generic and specific fields.

The generic ones are:

- **Title:** Title is seen at the top left of the page. Title shows which dashboard is being displayed. For Example: OpenStack Dashboard.
- **Time:** Time is seen at the top right of the page. Time indicates the time schedule for the log information. You can modify the time to indicate absolute, relative time in the past or specify automatically refresh rates.
- **Search bar:** Search bar is an input field where you can enter a query in the Lucene syntax format to filter the logs by specific fields (which depend on the fields for the index being selected)
- **Add a filter tab:** Use this tab to introduce filters graphically.

For more information on using Kibana, see the *Kibana documentation* (Version 5.5.1).


Cisco VIM stores the OpenStack logs in Elasticsearch. The Elasticsearch snapshots all the indices (where the data is stored) which are rotated on a periodic basis. You may not see the older data in Kibana if the data is rotated out and/or deleted.

Logs keep being visualized in Kibana as they are being updated in Elasticsearch on the Discover tab. To debug something on kibana, you can program the Kibana dashboard to auto-refresh at specific intervals (by default is off). To enable auto-refresh, click the date at the top right corner of the dashboard and click Auto-refresh to configure the desired value.

Figure 4: Auto-Refresh



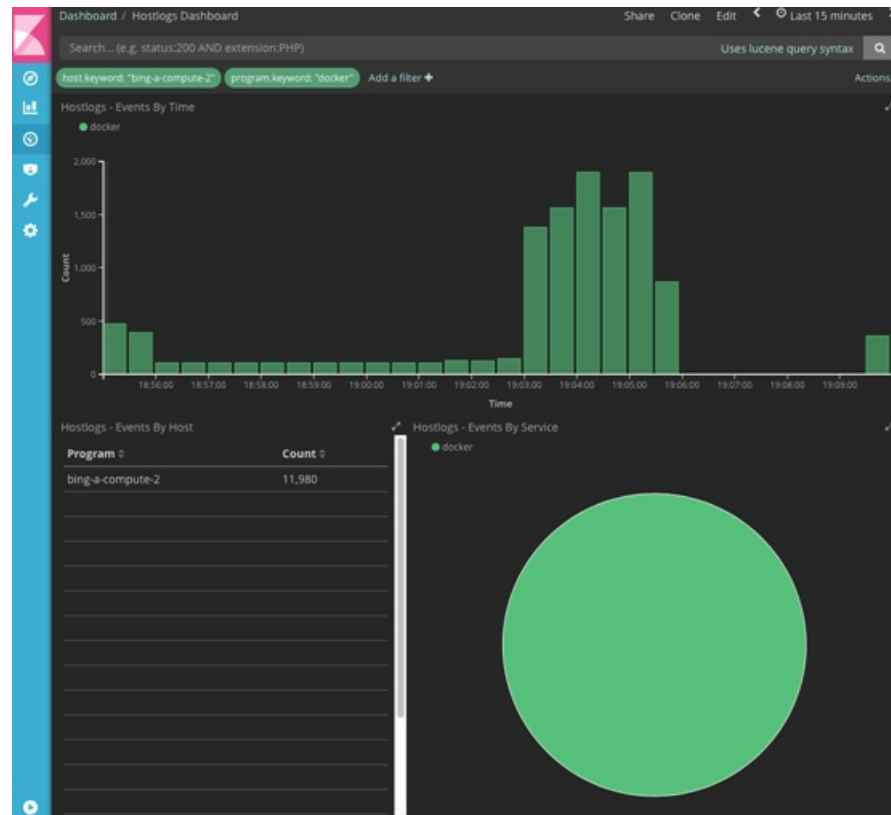
you can click **play/pause** button on the top navigator bar to continue/pause the refreshing of logs events:



a) Few examples on usage of filters in Openstack dashboard to gather useful information

- On the Hostlogs Dashboard, in the Events by Host panel, choose a hostname and click the + or - symbol that appears close to the hostname to include or exclude that server from the filter. Then, click the desired slice on the Events By Service panel to add the docker service to the section.
- Under the **Search** field, you see included sections in green and excluded sections in red.

Figure 5: Hostlogs Dashboard



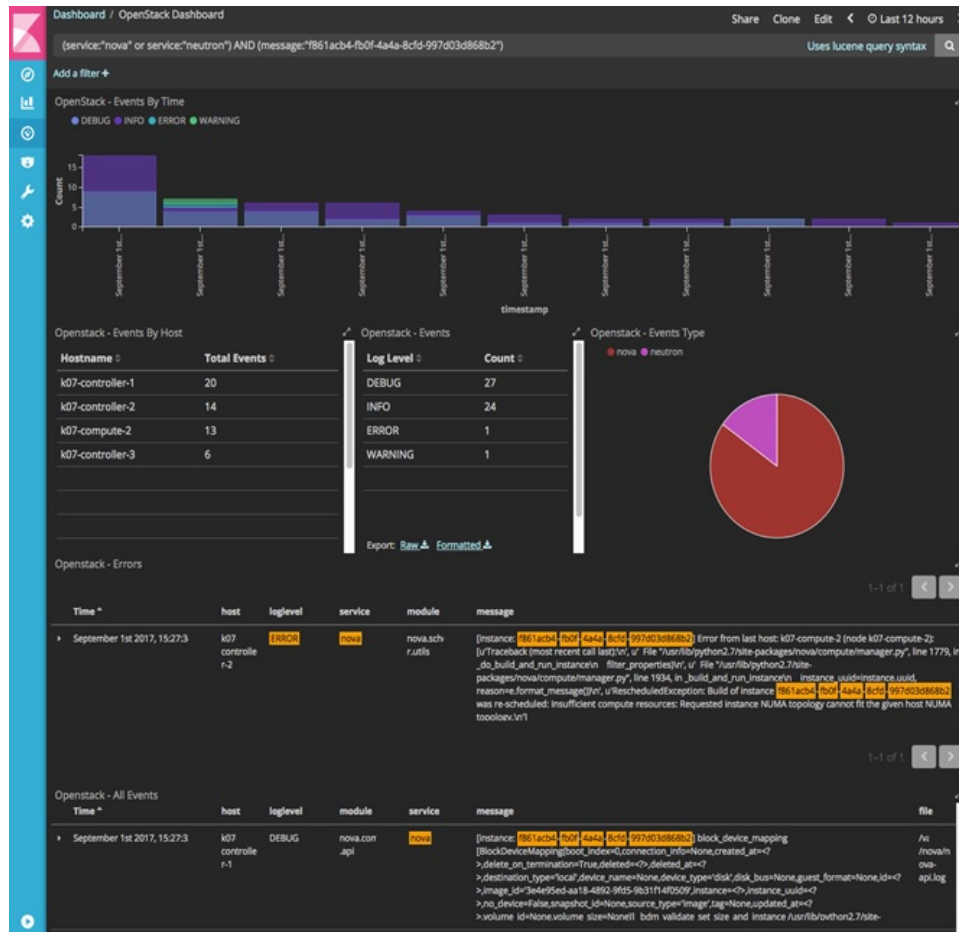
| Time | logdate | host | program | message |
|-----------------------------|---------|----------------|---------|---|
| September 6th 2017, 19:09:4 | - | bing-compute-2 | docker | 2017-09-07 02:09:44.342 8 ERROR oslo.messaging.drivers.impl_rabbit [] Failed to consume message from queue: (0, 0) (403) ACCESS_REFUSED - Login was refused using authentication mechanism AMQPPLAIN. For details see the broker logfile. |
| September 6th 2017, 19:09:4 | - | bing-compute-2 | docker | 2017-09-07 02:09:44.467 8 DEBUG neutron.plugins.ml2.drivers.openvswitch.agent.ovs_neutron_agent [req-809cbc46-3eaa-492f-a352-f81b34f82210 - ...] Agent rpc_loop - iteration:1397 started rpc_loop /usr/lib/python2.7/site-packages/neutron/plugins/ml2/drivers/openvswitch/agent/ovs_neutron_agent.py:1965 |
| September 6th 2017, 19:09:4 | - | bing-compute-2 | docker | 2017-09-07 02:09:44.472 8 DEBUG neutron.agent.linux.utils [req-809cbc46-3eaa-492f-a352-f81b34f82210 - ...] Running command: [ps, '-ppid', '85', '-o', 'pid='] create_process /usr/lib/python2.7/site-packages/neutron/agent/linux/utils.py:89 |
| September 6th 2017, 19:09:4 | - | bing-compute-2 | docker | 2017-09-07 02:09:44.055 8 ERROR oslo.messaging.drivers.impl_rabbit [] [5d14257f-d6f1-4b33-8532-4d80421b66ea] AMQP server on 10.23.222.122:5672 is unreachable: <AMQPError: unknown error>. Trying again in 1 seconds. Client port: None |
| September 6th 2017, 19:09:4 | - | bing-compute-2 | docker | 2017-09-07 02:09:44.096 8 ERROR oslo.messaging.drivers.impl_rabbit [] [9521d890-240f-46a4-9ff6-5037ad52d6b6] AMQP server on 10.23.222.122:5672 is unreachable: <AMQPError: unknown error>. Trying again in 1 seconds. Client port: None |
| September 6th 2017, 19:09:4 | - | bing-compute-2 | docker | 2017-09-07 02:09:44.574 8 DEBUG neutron.agent.linux.utils [req-809cbc46-3eaa-492f-a352-f81b34f82210 - ...] Exit code: 0 execute /usr/lib/python2.7/site-packages/neutron/agent/linux/utils.py:150 |
| September 6th 2017, 19:09:4 | - | bing-compute-2 | docker | 2017-09-07 02:09:44.470 8 DEBUG neutron.plugins.ml2.drivers.openvswitch.agent.openflow.native.ofswitch [req-809cbc46-3eaa-492f-a352-f81b34f82210 - ...] ofctl request: version=0x4,msg_type=0x12,msg_len=0x38,aid=0x496923cb,OFFFlowStatsRequest(cookie=0,cookie_mask=0,flags=0,match=OFFMatchoam_fields=(),out_group=4294967295,out_port=4294967295,table_id=23,type=1) result TOFFFlowStatsReslvbody= |

b) To know the log events in the Openstack for a given VM by writing the filter directly on the Search field:

Note The uuid of the VM is identified by executing `openstack nova list` or looking at the horizon website.

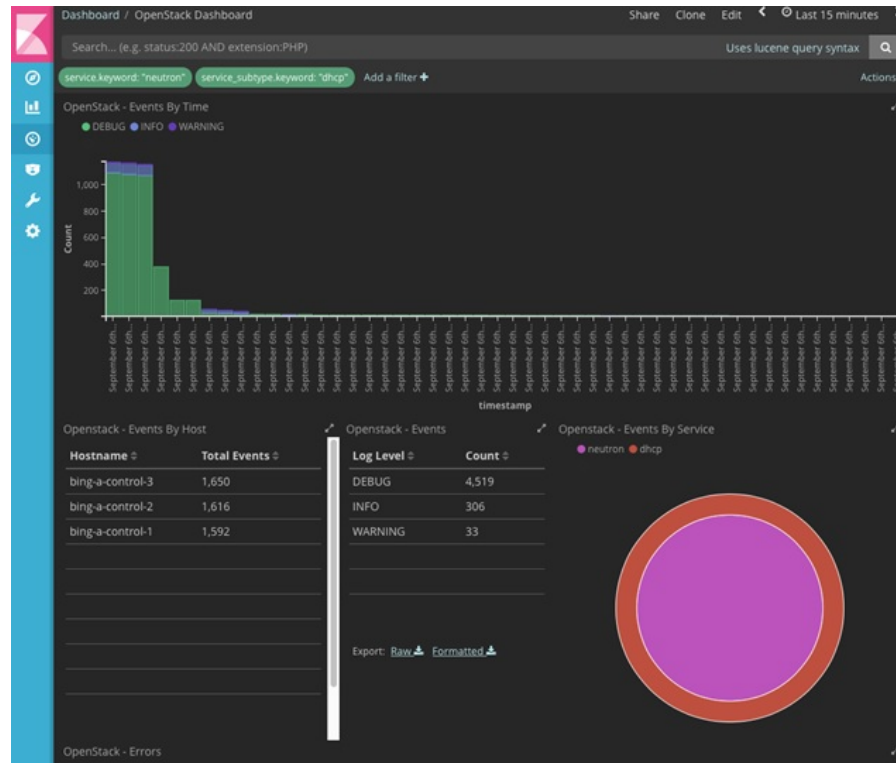
- Write the Lucene query (`service:nova and service:neutron and message:<uuid>`) in the **Search** field which is on top of the Dashboard. `<uuid>` is the number got from Horizon or nova list for the identifier of the instance VM.

Figure 6: Search Query Page

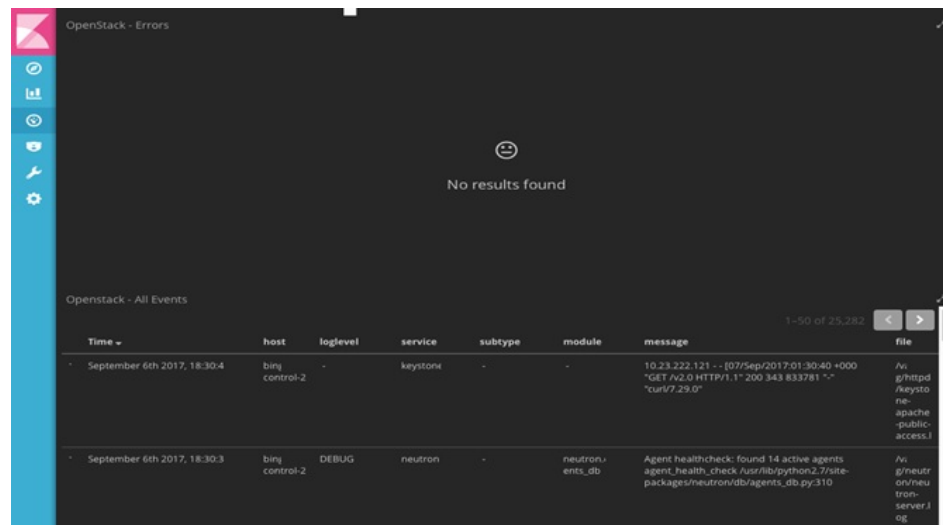


- For example, if the user wants to know the DHCP events of the Openstack Neutron add filters by clicking outer circle of pie chart:
 - On the OpenStack Dashboard, the Openstack - Events By Service panel has a pie chart with the inner section for the services and the outer sections for the service_subtypes. To add filters for selecting all the events in a service (for example, neutron), click on the inner section of the pie. To add filters for selecting the service_subtypes (for example, dhcp), click on the outer circle of the pie.

Figure 7: Events by Service



- You can scroll down the OpenStack Dashboard to see the OpenStack - Errors and the OpenStack - Events panel. The OpenStack - Errors panel displays the error messages. If there are no errors, the **No results found** message is displayed.

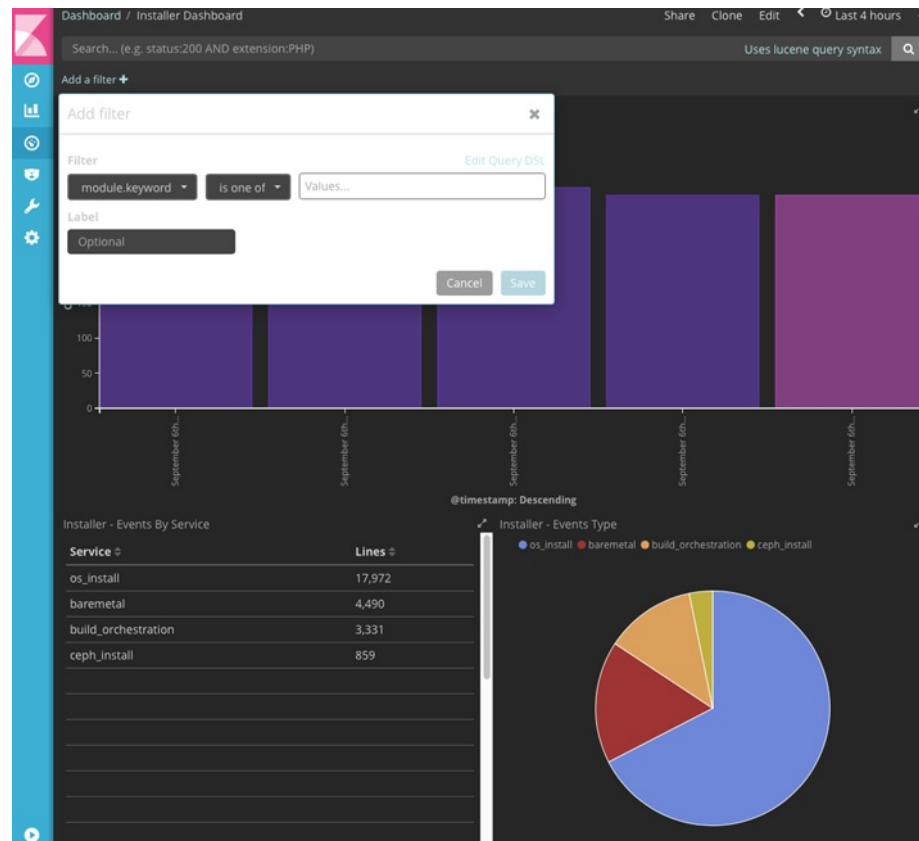


- Without knowing the Lucene Syntax, you can set the filter criteria in the **Search** field using the **Add a filter +** option.

Following are the steps to add a filter:


- Click Add a filter (+).
- Set the filter criteria by choosing appropriate label and operators from the drop-down lists, and entering keywords and click Save.

Figure 8: Add Filters Page



Set the filter criteria by choosing appropriate label and operators from the drop-down lists, and entering keywords.

Figure 9: Choosing Appropriate Labels



Rotation of the Cisco VIM Logs

Cisco VIM stores all logs in Elasticsearch. Elasticsearch indices are rotated on a periodic basis to prevent the disk space overflow by creating snapshots. The following lists show the Snapshots that are defined in `openstack_config.yaml`:

```
# vi ~/openstack-configs/openstack_config.yaml
...
# Elk rotation parameters
elk_rotation_frequency: "monthly" # Available: "daily", "weekly", "fortnightly", "monthly"
elk_rotation_size: 2 # Unit is in Gigabytes (float is allowed)
elk_rotation_del_older: 10 # Delete older than 10 units (where units depends on the
value set on elk_rotation_frequency)
...
```

You can change the frequency of the rotation by changing the values. For more information on how to set the Elasticsearch parameters through VIM API or CLI, refer to the section *Reconfiguring Passwords and OpenStack Configurations*.

Cisco VIM uses the open source Elasticsearch Curator tool to manage the Elasticsearch indices and snapshots. For more information about Elasticsearch handles snapshots, look at the official information on Elastic.co (Version 5.4) <https://www.elastic.co/guide/en/elasticsearch/client/curator/5.4/index.html>.

Snapshot Manager Tool for Elasticsearch

The `snapshot_mgr.sh` tool wraps up the Elasticsearch Curator APIs. This tool helps you to access the snapshots of the logs that are maintained by the Elasticsearch.

Run the following command to view the snapshot logs which is in the tools directory of the installer.

```
# ./tools/snapshot_mgr.py --help
usage: snapshot_mgr.py [options]
```

Snapshot Manager handles snapshot logs maintained by Elasticsearch

optional arguments:

```
-h, --help          show this help message and exit
--list              display all snapshots in Elasticsearch
--display GET_SS   get details of the snapshot called <GET_SS>
--create            create a snapshot
--restore RESTORE_SS restore snapshot named <RESTORE_SS>
--delete DELETE_SS delete the snapshot called <DELETE_SS>
--autodelete threshold_warning threshold_low threshold_high
autodelete snapshots until reach a disk space
threshold
```

Snapshot list gives you the details of the snapshot performed on the system like the UUID, the name the snapshot, end time of the snapshot, the state and the indices where it was snapshotted:

```
# ./snapshot_mgr.py --list
```

| indices_snapshotted | uuid | snapshot_name | time_snapshot_ended | state | failures |
|---------------------|------------------------|------------------------|---------------------|---------|----------|
| hostlogs-2018.03.02 | 6WGVUnKjQbGtZYzfC0yeEg | curator-20180304140002 | 2018-03-04 14:00:04 | SUCCESS | - |


```
| U4IVWJNnQW6PdFWxpRUC-A | curator-20180304150001 | 2018-03-04 15:00:04 | SUCCESS |
hostlogs-2018.03.03
| 5RxDuhnETC6TW4XSPDNZlw | curator-20180304160001 | 2018-03-04 16:00:24 | SUCCESS |
installer-2018.03.03, installer-2018.03.01, installer-2018.03.02, openstack-2018.03.02,
hostlogs-2018.03.04, installer-2018.03.04 | - |
| k2gZYwLeRPO98bJZslI2pw | curator-20180305040002 | 2018-03-05 04:00:32 | SUCCESS |
openstack-2018.03.03, hostlogs-2018.03.04, installer-2018.03.04
```

To view the details of the individual snapshot run the display option command.:

```
# ./tools/snapshot_mgr.py --display curator-20180304140002
{ 'duration_in_millis': 1944,
  'end_time': '2018-03-04T14:00:04.019Z',
  'end_time_in_millis': 1520172004019,
  'failures': [],
  'indices': ['hostlogs-2018.03.02'],
  'shards': { 'failed': 0, 'successful': 5, 'total': 5},
  'snapshot': 'curator-20180304140002',
  'start_time': '2018-03-04T14:00:02.075Z',
  'start_time_in_millis': 1520172002075,
  'state': 'SUCCESS',
  'uuid': '6WGVUnKjQbGtZYzfC0yeEg',
  'version': '6.0.0',
  'version_id': 6000099}
```

To create a snapshot run the following command:

```
# ./tools/snapshot_mgr.py --create
Executing: curl PUT
http://localhost:9200/_snapshot/es_backup/3a9b90c2979b46bf9c7b3f9223074d5d?wait_for_completion=true
-d
{ 'indices': 'installer-*,hostlogs-*,openstack-*,vmtp-*', 'ignore_unavailable': 'true',
  'include_global_state': 'false' }
Response: {u'snapshot': {u'uuid': u'BSznQj1SQ9mjxxxk9swTirQ', u'duration_in_millis': 46496,
  u'start_time':
  u'2018-03-06T16:37:49.774Z', u'shards': {u'successful': 35, u'failed': 0, u'total': 35},
  u'version_id': 6000099,
  u'end_time_in_millis': 1520354316270, u'state': u'SUCCESS', u'version': u'6.0.0',
  u'snapshot': u'3a9b90c2979b46bf9c7b3f9223074d5d', u'end_time': u'2018-03-06T16:38:36.270Z',

  u'indices': [u'installer-2018.03.06', u'vmtp-2018.03.02', u'hostlogs-2018.03.06',
  u'hostlogs-2018.03.05',
  u'installer-2018.03.05', u'openstack-2018.03.05', u'openstack-2018.03.06'],
  u'failures': [], u'start_time_in_millis': 1520354269774}}
```

Run the following command to delete a snapshot:

```
# ./tools/snapshot_mgr.py --delete 3a9b90c2979b46bf9c7b3f9223074d5d
Executing: curl DELETE
http://localhost:9200/_snapshot/es_backup/3a9b90c2979b46bf9c7b3f9223074d5d -d None
Response: {u'acknowledged': True}
```

Restore the indices of a snapshot back to the Elasticsearch database by using the restore option. Run the following command to restore:

```
# ./snapshot_mgr.py --restore curator-20180306050001
Executing: curl POST
http://localhost:9200/hostlogs-2018.03.04,installer-2018.03.05,installer-2018.03.04,
openstack-2018.03.04,hostlogs-2018.03.05,openstack-2018.03.02/_close -d None
```

Remote NFS Backup for Elasticsearch Snapshots

Cisco VIM 2.4, supports remote NFS backup of the Elasticsearch snapshots. This allows you to empty the disk space in the Elasticsearch snapshots. You can use the snapshot manager tool to manually create, list, show, and delete snapshots.

Remote NFS backup of the Elasticsearch snapshots.feature can be configured by adding the following section to the `setup_data.yaml` configuration file:

```
ES_REMOTE_BACKUP: # Set if Elasticsearch backups can use a remote host
  service: 'NFS' # Set if an remote NFS server is used
  remote_host: <ip_addr> # IP of the NFS server
  remote_path: /root/es_remote # Path to location of the backups in the remote server
```

Important considerations about the remote NFS directory on the remote server (specified by the `remote_path` config option):

- This directory allows the `elasticsearch` user (pid number 2020) and group `mercury` (pid 500) to read, and write. Otherwise, Curator cannot copy the snapshots to the remote NFS directory.
- It is good if the folder is empty and is used only by Cisco VIM.
- Cisco VIM does not delete the information in this directory after `unbootstrap`.

You can enable or disable this feature by running `reconfigure`. Also it can change the `remote_host` ip or the `remote_path` by doing `reconfigure`.

Network Performance Test with NFVBench

NFVBench is a network performance benchmarking tool integrated with Cisco VIM. For more details, refer to NFVBench section of *Chapter 1* in the admin guide for details.