



Cluster Maintenance

From the navigation pane, choose the options that are under the **Troubleshoot** menu.


- [Service Status, on page 1](#)
- [Admiral Alerts, on page 2](#)
- [Cluster Status, on page 11](#)
- [Data Backup and Restore, on page 14](#)
- [High Availability in Secure Workload, on page 33](#)
- [VM Information, on page 41](#)
- [Upgrading Cluster, on page 41](#)
- [Snapshots, on page 50](#)
- [Explore/Snapshot Endpoints Overview, on page 58](#)
- [Server Maintenance, on page 71](#)
- [Disk Maintenance, on page 78](#)
- [Cluster Maintenance Operations, on page 89](#)
- [Data Tap Admin - Data Taps, on page 92](#)

Service Status

In the left navigation pane, the **Troubleshoot > Service Status** page displays the health of all services that are used in your Cisco Secure Workload cluster along with their dependencies.

The graph view shows the health of the service, each node in the graph shows the health of the service and an edge represents dependency on other services. Unhealthy services are marked either red when the service is unavailable and orange when the service is degraded but available. A green node indicates that the service is healthy. For more debug information on these nodes, use tree view which has the **Expand All** button to show all child nodes in the dependency tree. “Down” indicates that the service is not functional and “Unhealthy” indicates that the service is not fully functional.

Figure 1: Service Status Page



Admiral Alerts

Admiral is an integrated alerting system. It processes alerts off the service health reported by [Service Status](#). Thus, users have a unified way of determining service/cluster health. Service Status shows the current (point in time) health of a service. The service is considered down when it reports as red on service status, otherwise it is considered as up. Uptime is the time when the service is reported as up. Admiral evaluates service health reported by service status over a period of time and raises an alert if the service uptime percentage falls below a certain threshold. This evaluation over a duration of time ensures that we reduce false positives and alert only on true service outages.

As services are different in their alerting needs, this percentage and time interval are fixed differently for different services.

Customers can use admiral notifications to be notified of these events. They are also visible on the [Investigate > Alerts](#) page under type PLATFORM.



Note Only a chosen subset of services have an admiral alert associated with them. If a service is not in the above subset, no admiral alert will be raised when it goes down. This subset of services with admiral alerts and their alerting threshold percentages and time intervals are fixed i.e. not user configurable.

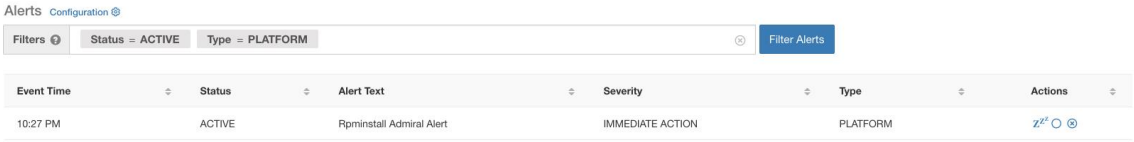
The following sections describe admiral alerts and notifications in more detail.

Lifecycle of Admiral Alert

Admiral checks for the uptime of services on service status. It raises an alert when this uptime becomes lower than the pre-configured threshold for alerting.

As an example, Rpminstall is a service which is used to install rpms during deploy, upgrades, patches etc. It is configured to generate an admiral alert if its uptime is less than 80% over one hour. If Rpminstall service goes down for a duration longer than the threshold specified above, an admiral alert for Rpminstall is generated with status ACTIVE.

Figure 2: Active Admiral Alert



The screenshot shows the Alerts Configuration page with filters set to Status = ACTIVE and Type = PLATFORM. A table displays one active alert:

Event Time	Status	Alert Text	Severity	Type	Actions
10:27 PM	ACTIVE	Rpminstall Admiral Alert	IMMEDIATE ACTION	PLATFORM	z ^z ○ ⊗

When the service recovers, its uptime percentage starts increasing. When the uptime goes higher than its threshold, the alert auto closes and its status moves to CLOSED. In the Rpminstall example described above, Rpminstall Admiral Alert will auto close when its uptime goes over 80% in one hour.



Note The close of alert will ALWAYS lag the service becoming normal. This is because admiral looks at service health over a duration of time. In the above example, since Rpminstall alert threshold is set to 80% of an hour of uptime, it will need to be up for at least 48 minutes (80% of one hour) before the alert will close.

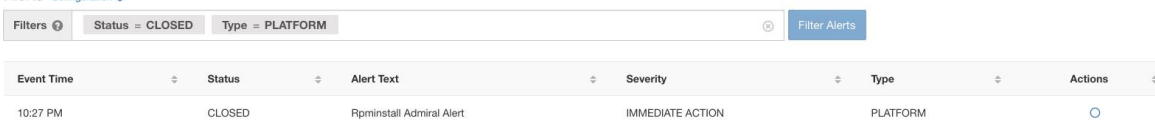
No action is required to close the alert. This ensures that all ACTIVE admiral alerts indicate a current underlying issue that needs attention.



Note No dedicated notification is generated when alerts close.

After an alert moves to CLOSED, it will no longer show under ACTIVE alerts. Closed alerts can still be seen on the UI using the filter Status=CLOSED as shown below:

Figure 3: Admiral Alert Auto Closes When Service Recovers



The screenshot shows the Alerts Configuration page with filters set to Status = CLOSED and Type = PLATFORM. A table displays one closed alert:

Event Time	Status	Alert Text	Severity	Type	Actions
10:27 PM	CLOSED	Rpminstall Admiral Alert	IMMEDIATE ACTION	PLATFORM	○

There are two kinds of admiral alerts:

- [Individual Admiral Alert](#)
- [Summary Admiral Alert](#)

Individual Admiral Alert

The alerts that are described in the previous section, alerts that are raised for individual services, fall under the Individual Admiral Alert category. The alert text always contains `<Service Name> Admiral Alert`. This makes it easy to filter individual alerts by service or by the **Admiral Alert** suffix.

Figure 4: Alert Text Filter for Individual Admiral Alerts

Event Time	Status	Alert Text	Severity	Type	Actions
10:14 PM	ACTIVE	Adm Admiral Alert	IMMEDIATE ACTION	PLATFORM	Z ¹ ○ ◎
7:04 PM	ACTIVE	Rpminstall Admiral Alert	IMMEDIATE ACTION	PLATFORM	Z ¹ ○ ◎
2:58 PM	ACTIVE	DataBackup Admiral Alert	IMMEDIATE ACTION	PLATFORM	Z ¹ ○ ◎

Summary Admiral Alert

The admiral generates daily Summary Alerts at midnight UTC. They contain a list of currently active alerts and all alerts closed within the last one day. This allows the user to see the overall cluster health reported by admiral in one place. This is also useful to see closed alerts which do not generate a dedicated notification otherwise. If the cluster is healthy and no alerts were closed within the last one day, no summary notifications are generated for that day. This is done to reduce unnecessary notifications and noise.

The Alerts Text in this case is always **Admiral Summary**. This makes it easy to filter summary alerts as shown in the following figure.

Figure 5: Admiral Summary Text Filter

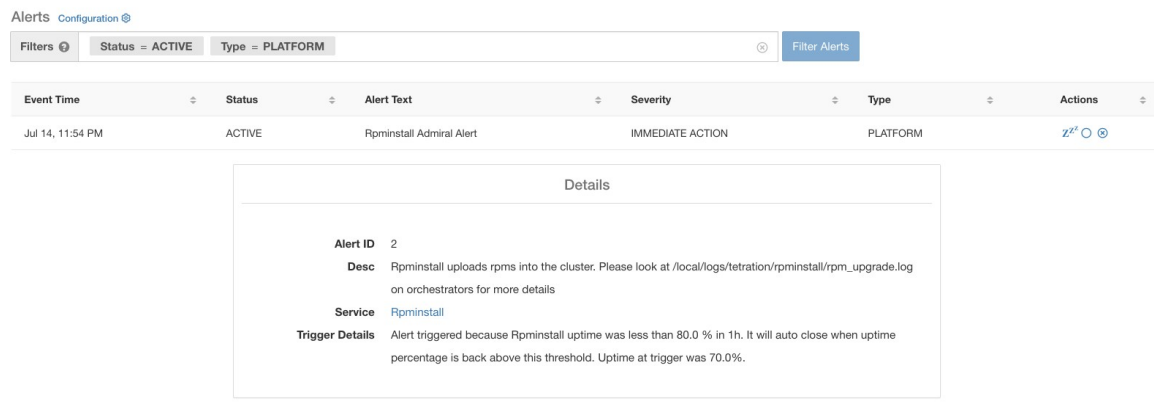
Event Time	Status	Alert Text	Severity	Type	Actions
5:04 PM	ACTIVE	Admiral Summary	LOW	PLATFORM	Z ¹ ○ ◎

Alert Details

Individual Alerts

On clicking the alert for an individual admiral alert, it expands to show fields useful for debugging and analyzing the alert.

Figure 6: Alert Details



Alerts Configuration

Filters Status = ACTIVE Type = PLATFORM Filter Alerts

Event Time	Status	Alert Text	Severity	Type	Actions
Jul 14, 11:54 PM	ACTIVE	Rpminstall Admiral Alert	IMMEDIATE ACTION	PLATFORM	z x o

Details

Alert ID 2

Desc Rpminstall uploads rpms into the cluster. Please look at /local/logs/tetration/rpminstall/rpm_upgrade.log on orchestrators for more details

Service [Rpminstall](#)

Trigger Details Alert triggered because Rpminstall uptime was less than 80.0 % in 1h. It will auto close when uptime percentage is back above this threshold. Uptime at trigger was 70.0%.

Table 1: Alert Details Field Descriptions

Field	Description
Alert ID	Unique ID for alerts. This helps unify a particular incidence of a service going down. As mentioned earlier, when the underlying uptime of the service being reported by the alert becomes normal, the alert auto closes. If the same service goes down again next, a new alert with a different Alert ID is generated. Thus the alert id helps unify each incident of the alert being raised.
Desc	The description field contains additional information about the service issue causing the alert.
Service	This contains a link taking the user to the service status page where the status of the service can be seen. User can also get more details on why the service is being marked down in the service status page.
Trigger Details	This contains the details on the trigger thresholds for the service. User can understand when to expect the alert to close after its underlying service is restored by looking at these thresholds. For example, Rpminstall threshold is mentioned as 80% uptime over one hour. Thus rpminstall service must be up for at least 48 minutes (80% of one hour) before the alert will auto close. This also shows the uptime value that is seen for the service when the alert was fired.

The following is a sample JSON Kafka output:

```
{
  "severity": "IMMEDIATE_ACTION",
  "tenant_id": 0,
  "alert_time": 1595630519423,
```

```

"alert_text": "Rpminstall Admiral Alert",
"key_id": "ADMIRAL_ALERT_5",
"alert_id": "/Alerts/5efcfd5497d4f474f1707c2/DataSource{location_type='TETRATION',
location_name='platform', location_grain='MIN',
root_scope_id='5efcfd5497d4f474f1707c2'}/66eb975f5f987fe9eaeafa81cee757c8b6dac5facc26554182d8112a98b35c4ab",

"root_scope_id": "5efcfd5497d4f474f1707c2",
"type": "PLATFORM",
"event_time": 1595630511858,
"Check /local/logs/tetration/rpminstall/rpm_upgrade.log on
orchestrators for more details\", \"Trigger Details\": \"Alert triggered because Rpminstall
uptime was less than 80.0 % in 1h. It will auto close when uptime percentage is back above
this threshold. Uptime at trigger was 65.0%. \"/>

```

All individual alerts follow the JSON Kafka format. The services (from service status) that are covered by admiral monitoring are listed in the following table:

Table 2: Services Covered by Admiral Monitoring

Service	Trigger Condition	Severity
KubernetesApiServer	Service Uptime falls below 90% in last 15 mins.	IMMEDIATE ACTION
Adm	Service Uptime falls below 90% in the last one hour.	IMMEDIATE ACTION
DataBackup	Service Uptime falls below 90% in the last 6 hours.	IMMEDIATE ACTION
DiskUsageCritical	Service Uptime falls below 80% in the last one hour.	IMMEDIATE ACTION
RebootRequired	Service Uptime falls below 90% in the last one hour.	IMMEDIATE ACTION
Rpminstall	Service Uptime falls below 80% in the last one hour.	IMMEDIATE ACTION
SecondaryNN_checkpoint_status	Service Uptime falls below 90% in the last one hour.	IMMEDIATE ACTION

For 8 or 39 RU physical clusters, the following services are also monitored:

Table 3: Services Covered by Admiral Monitoring for 8 or 39 RU Clusters

Service	Trigger Condition	Severity
DIMMFailure	Service Uptime falls below 80% in the last one hour.	IMMEDIATE ACTION
DiskFailure	Service Uptime falls below 80% in the last one hour.	IMMEDIATE ACTION
FanSpeed	Service Uptime falls below 80% in the last one hour.	IMMEDIATE ACTION

Service	Trigger Condition	Severity
ClusterSwitches	Service Uptime falls below 80% in the last one hour.	IMMEDIATE ACTION



Note Admiral relies on processing metrics that are generated by Service Status to generate alerts. If metric retrieval is not possible for a prolonged duration (For Eg: If service status is down), then an alert (TSDBOracleConnectivity) is raised notifying that service-based alert processing is off on the cluster.

Summary Alerts

Summary alerts are informational in nature and are always set to LOW priority. On clicking an admiral summary alert, it expands to show various fields containing summary information on admiral alerts.

Figure 7: Details of Admiral Summary Alert

Details	
Desc	Summary Of Alerts For Jul 14
Open	Service DataBackup with Alert ID 1.
Recently Closed	Service Rpminstall with Alert ID 3.
Service	Admiral
Summary ID	ADMIRAL SUMMARY Jul 14 20 23 13

Table 4: Admiral Summary Alert Field Descriptions

Field	Description
Desc	The description field contains the day for the daily summary.
Open	The open alerts indicate which alerts were active when the summary was generated.
Recently Closed	This contains alerts which closed within the last 24 hours i.e. during the day for which the summary was generated. Each alert's ID is also included. Since the alerts auto close, a given service could have gone down and created an alert, then become normal and alert auto closed. It could have done this multiple times in a day in which case recently closed will list each incident along with its unique alert id. However, this is not expected to happen often given that each service has to be up for a threshold time before its alert is closed. User can filter with Status = CLOSED to get more information on each incident.

Field	Description
Service	Service Status link for Admiral which is the service processing and generating the daily summary.
Summary ID	ID of the summary alert.

The following is a sample JSON Kafka output:

```
{
  "severity": "LOW",
  "tenant_id": 0,
  "alert_time": 1595721914808,
  "alert_text": "Admiral Summary",
  "key_id": "ADMIRAL_SUMMARY_Jul-26-20-00-04",
  "alert_id": "/Alerts/5efcfd5497d4f474f1707c2/DataSource{location_type='TETRATION',
location_name='platform', location_grain='MIN',
root_scope_id='5efcfd5497d4f474f1707c2'}/e95da4521012a4789048f72a791fb58ab233bbff63e6cbc421525d4272d469aa",

  "root_scope_id": "5efcfd5497d4f474f1707c2",
  "type": "PLATFORM",
  "event_time": 1595721856303,
  "alert_details": "{\"Desc\":\"Summary of alerts for Jul-26\\\", \"Recently
Closed\\\": \"None\\\", \"Open\\\": \" Service Rpminstall with Alert ID
5.\\\", \"Service\\\": \"Admiral\\\", \"Summary ID\\\": \"ADMIRAL_SUMMARY_Jul-26-20-00-04\\\"}\"
}
```

An example summary alert containing a service raising multiple alerts in a day is shown below:

Figure 8: Multiple Alerts

Details	
Desc	Summary Of Alerts For Jul 15
Open	Service DataBackup with Alert ID 1. Service Adm with Alert ID 7.
Recently Closed	Service Rpminstall with Alert ID 9. Service Rpminstall with Alert ID 10.
Service	Admiral
Summary ID	ADMIRAL SUMMARY Jul 15 20 19 30

User Actions

Since admiral alerts generate an individual notification only once per alert, including/excluding or snoozing specific alerts are not needed. Alerts auto close when the service becomes normal for threshold uptime as described above. There is a force close option available to forcibly close an alert. This should normally be used only to remove summary alerts from the UI as individual alerts close automatically.

Figure 9: Force Close Alert

Alerts Configuration					
Event Time	Status	Alert Text	Severity	Type	Force close an alert
5:04 PM	ACTIVE	Admiral Summary	LOW	PLATFORM	z x o

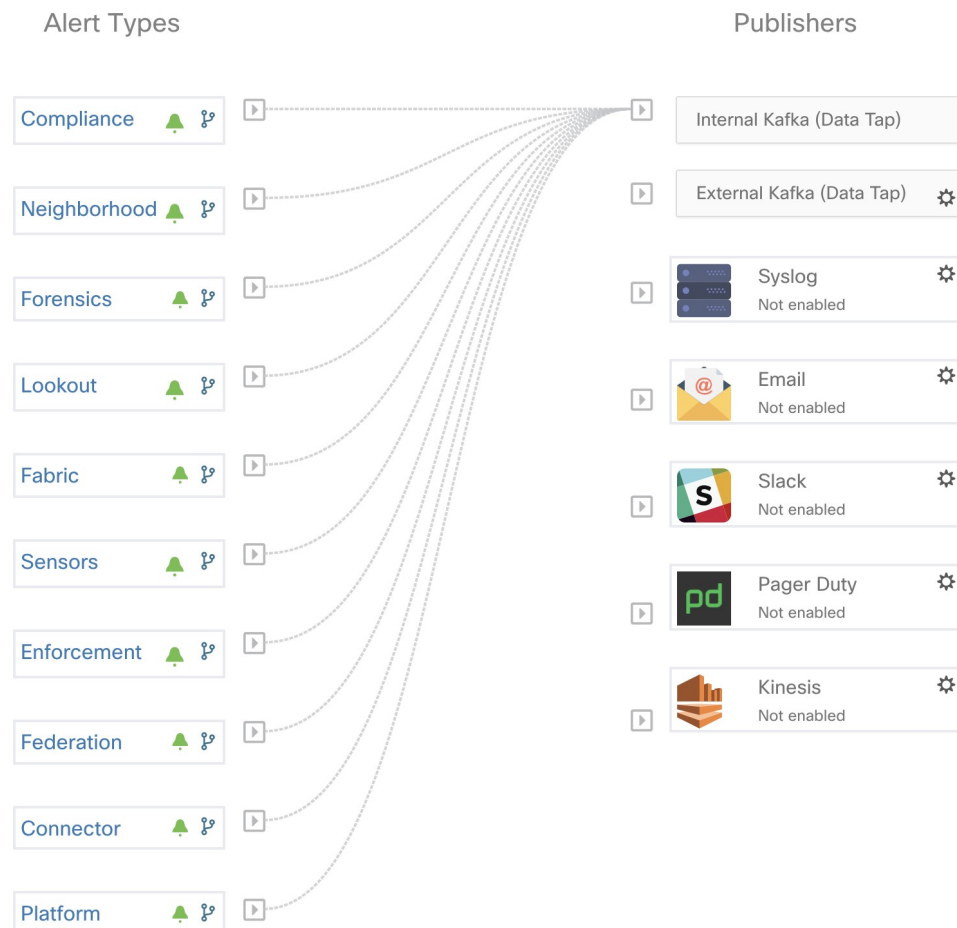
**Warning**

Individual alerts should not be closed forcefully. Doing so while the underlying service is still down or its uptime is below its expected threshold will lead to another alert getting raised for the same service on the next admiral processing iteration.

Admiral Notifications

Admiral Alerts are of Type PLATFORM. As such, these alerts can be configured to be sent to various publishers by appropriate connections for Platform Alerts via the configuration page `./configuration`. For convenience, the connection is turned on between Platform Alerts and Internal Kafka by default which allows admiral alerts to be seen on the Current Alerts page (go to **Investigate** > **Alerts**) without any manual configuration.

Figure 10: Platform Alerts Configuration



Admiral Alerts are also sent to the email address configured under **Platform** > **Cluster Configuration** > **Admiral Alert Email**.

Figure 11: Sample Admiral Email

There is a new admiral platform alert on your tetration cluster.

Service: Rpminstall

Start Time: 2020-07-14 23:09 UTC

Alert ID: 3

Description: Rpminstall uploads rpms into the cluster. Please look at /local/logs/tetration/rpminstall/rpm_upgrade.log for more details

This is an auto generated message about platform alerts on your cluster.

For more details, please go to [Alerts On Cluster](#)

Please make sure that you are on **Default Scope** to view the alerts.

Thus, users can receive admiral notifications even if they don't have the TAN edge appliance setup. This is similar to Bosun behavior in previous releases.

Figure 12: Admiral Email

cluster_state	Enabled till 2020-10-11 19:15:49 UTC
Cluster UUID ⓘ	8194c5ef-65df-8aa1-5963-d10514761b6f
Admiral Alert Email ⓘ	admiral@test.com ✉

These email notifications are generated based on the same triggers as the Current Alerts page. Thus, they are sent on alert creation and a daily summary email at midnight UTC. The daily summary email lists all active alerts and those closed within the last 24 hours.

Figure 13: Sample Summary Admiral Email

Daily summary of admiral platform alerts:

State:Active

Service: DataBackup

Start Time: 2020-07-14 21:58 UTC

Alert ID: 1

Description: The last successful checkpoint was over 48 hours ago.

State:Closed

Service: Rpminstall

Start Time: 2020-07-14 22:41 UTC

Alert ID: 2

Description: Rpminstall uploads rpms into the cluster. Please look at /local/logs/tetration/rpminstall/rpm_upgrade.log for more details

This is an auto generated message about platform alerts on your cluster.

For more details, please go to [Alerts On Cluster](#)

Please make sure that you are on **Default Scope** to view the alerts.

If there are no active alerts and no alerts closed within the last 24 hours, the summary emails are skipped to reduce email noise.

Cluster Status

The **Cluster Status** page under the **Troubleshoot** menu in the left navigation bar can be accessed by **Site Admin** users but the actions can be carried out by **Customer Support** users only. It shows the status of all the physical servers in the Cisco Secure Workload rack. Each row in the table represents a physical node with details such as its hardware and firmware configuration and CIMC IP address (if assigned). The detail view of the node can be viewed by clicking on the row. In this page, we can also change the CIMC password of the nodes and enable or disable external access to them. Orchestrator state is also displayed on the cluster status page to provide context for customer support.

Figure 14: Cluster Status

The screenshot shows the Cluster Status page for a cluster named 'Model: 8RU-PROD'. At the top, there are two buttons: 'CIMC/TOR guest password' and 'Change external access'. The Orchestrator State is 'IDLE'. Below the buttons, it says 'Displaying 6 nodes (0 selected)'. A table lists nodes with columns for checkboxes, State, Status, Switch Port, Serial, Uptime, and CIMC Snapshots. Two nodes are visible, both 'Commissioned' and 'Active'. The first node has Serial 'FCH2206V1NF' and Uptime '2mo 27d 13h 3m 47s'. The second node has Serial 'FCH2206V1ZF' and Uptime '2mo 27d 13h 2m 52s'. A detailed view for the first node is shown below the table, displaying various system details:

- Serial: FCH2206V1NF
- Private IP: 1.1.1.4
- CIMC IP: 10.13.4.12
- Status: Active
- State: Commissioned
- SW Version: 3.6.0.10.devel
- Hardware: 44 cores, 962G memory, 8 disks, 17.57T space, SSD
- Firmware: View Firmware Upgrade Logs
 - CIMC: 2.0(10e)
 - BIOS: 2.0.10e.0
 - Cisco 12G SAS Modular RAID Controller Slot HBA: 24.12.1-0205
 - UCS VIC 1225 10Gbps 2 port CNA SFP+ Slot 1: 4.1(3a)
 - Intel(R) I350 1 Gbps Network Controller Slot L: 0x80000E74-1.810.8
 - UCS VIC 1225 10Gbps 2 port CNA SFP+ Slot 2: 4.1(3a)
- Instances:
 - collector/Datamover-6
 - datanode-6
 - druidHistoricalBroker-4
 - enforcementCoordinator-3
 - orchestrator-2
 - redis-1
 - secondary/NameNode-1
- Disks Status:
 - 252:1 HEALTHY
 - 252:2 HEALTHY
 - 252:3 HEALTHY
 - 252:4 HEALTHY
 - 252:5 HEALTHY
 - 252:6 HEALTHY
 - 252:7 HEALTHY
 - 252:8 HEALTHY

Actions that affect all nodes

Changing CIMC password and enabling or disabling external CIMC access can be done using the **CIMC/TOR guest password** and **Change external access** options. The actions affect all nodes in the cluster.

External CIMC Access Node Details

Clicking **Change external access** opens a dialog box that provides the status of external CIMC access and allows external access to CIMC to be enabled, renewed, or disabled.

Clicking **Enable** configures the cluster in the background to enable external CIMC access. It can take up to 60 seconds for the tasks to complete and external CIMC access to be fully enabled. When external CIMC access is enabled, a dialog box displays when access is set to automatically expire and **Enable** changes to **Renew** to reflect that you can renew external CIMC access. Renewing external CIMC access increases the expiry time by two hours from the current time.

If external CIMC access is enabled, the CIMC IP address in the node details (viewable by clicking on a row for a node) becomes a clickable link that allows you directly access the CIMC UI. You may need to reload the cluster status page to view the links.

Figure 15: External CIMC Access Node Details

The screenshot shows the detailed view of a node with Serial 'FCH2206V1NF'. A dialog box is open over the CIMC IP address '10.13.4.11', displaying the message: 'External access to CIMC UI is enabled'. The node details are the same as in Figure 14, but the Instance list is different:

- Instances:
 - adhocKafkaXL-1
 - collector/Datamover-5
 - datanode-5
 - druidHistoricalBroker-3
 - elasticsearch-3
 - namenode-1
 - orchestrator-1

The CIMC UI usually has a self-signed certificate, accessing the CIMC UI will likely result in an error in the browser indicating that the certificate is not valid. If you are using Google Chrome this may require you to type **thisisunsafe** without quotes when the invalid certificate error is shown in Google Chrome to bypass the certificate check and access the CIMC UI.

Within the CIMC UI, KVM access is only functional if the CIMC version is 4.1(1g) or later. After external CIMC access is enabled, it is automatically disabled in two hours time unless access is renewed or disabled.

Disabling external CIMC access configures the cluster in the background to disable external CIMC access. It can take up to 60 seconds for the task to complete and external CIMC access to be fully disabled.

Table 5: Physical Node Details

Field	Description
Status	<p>The Status field indicates the power status of the node. Possible values are:</p> <ul style="list-style-type: none"> • Active: The node is powered on. • Inactive: The node is not powered-on or connected.
State	<p>The State field indicates the cluster membership state for the node. Possible values are:</p> <ul style="list-style-type: none"> • New: The node is not yet part of the cluster. • Initialized: The node is part of the cluster. However, Secure Workload is not deployed on the node. • Commissioned: The node is up and running with SecureWorkload deployed on it. <p>The SW version field is also indicated and it turns red if an individual node does not have the same version as that of the whole cluster.</p> <ul style="list-style-type: none"> • Decommissioned: The node has been removed from the cluster for troubleshooting purposes. The node must be replaced with new hardware. A node can be decommissioned using the decommission action, see the following actions.
Switch Port	Refers to the switch port of the two switches on which the physical node is connected.
Uptime	Indicates the time for which the node has been running without a restart or shutdown.
CIMC Snapshots	Can be used to initiate a CIMC Tech Support collection and download a CIMC Tech Support.

Table 6: Cluster Remedial Actions

Action	Description
Commission	Select this action to integrate new nodes into the cluster. Only nodes with the state New are selectable for this action.
Decommission	Select this action to remove nodes that are part of the cluster. Only the nodes with state Commissioned or Initialized are selectable for this action.
Reimage	Select this action to redeploy the Secure Workload. This could erase all cluster data and is especially useful to upgrade the bare metal operating system from an older version to a new one. This step is required when a bare metal is decommissioned.
Firmware upgrade	Firmware information is available for the nodes where CIMC IP is reachable. This action is helpful to upgrade firmware on the nodes with older versions.
Power off	Select this action to power down the nodes. Note You cannot power down the nodes with the Inactive and Shutdown in progress status.

Firmware Upgrade Details

The Secure Workload on-premises cluster bundles a Unified Computing System (UCS) Cisco Integrated Management Controller (CIMC) Host Upgrade Utility (HUU) ISO. The firmware upgrade option on the Cluster Status page can be used to update a physical bare metal to the version of UCS firmware included in the HUU ISO that has been bundled in the Secure Workload RPMs.

A bare metal host can have the firmware update started on it when the status is *Active* or *Inactive* as long as the bare metal state is not *Initialized* or *SKU Mismatch*. Only one bare metal can have its UCS firmware that is updated at a time. To start the firmware update, the Secure Workload orchestrator state must be *Idle*. When the UCS firmware update is initiated, some of the UI functionality specific to the Cluster Status page may be temporarily impacted if the consul leader, active orchestrator, or active firmware manager (fwmgr) must be switched to other hosts - these switchovers should occur automatically. During the firmware update, the firmware details for the bare metal being updated will not be displayed and after the update it may take up to 15 minutes for the firmware details to display again in the Cluster Status page. Before starting the firmware update, check the Service Status page to verify that all services are healthy.

When you initiate a firmware update on a bare metal, fwmgr will verify that the update can continue, gracefully power down the bare metal if needed, then login to the CIMC on the bare metal and start the HUU-based firmware update. That HUU-based firmware update process involves booting the bare metal into the HUU ISO, doing the update, rebooting CIMC to activate the new firmware then booting the bare metal back into the HUU ISO to verify the update was completed. The overall update process can take 2+ hours for a G1 bare metal or 1+ hours for a G2 bare metal. When the firmware update process is initiated, the Service Status page may indicate that some services are unhealthy since a bare metal and all the virtual machines running on that bare metal are no longer active in the cluster. When the firmware update completes, it can take an extra 30

minutes for the bare metal to become active in the cluster again and more time may be needed for all services to become healthy again. If services do not recover within two hours after a firmware update, contact a customer service representative.

You can click a bare metal node in the Cluster Status page to expand details about the bare metal. When a firmware update is initiated, you can click the *View Firmware Upgrade Logs* button to view the status of the firmware update. The log contains the overall status of the firmware update and the status can be one of the following:

- **Firmware update has been triggered:** The firmware update was requested but has not started yet. During this status fwmgr will be checking to make sure the services required for the firmware update are functional and that CIMC can reach those services.
- **Firmware update is running:** The firmware update has been started. When a firmware update reaches this state, CIMC and HUU are in control of the update, and the Secure Workload cluster will report the status that it gets from CIMC about the update.
- **Firmware update has timed out:** This indicates that some process from the firmware update has exceeded the time that we expect it to complete. The overall firmware update process has a 240-minute time limit when it enters the *Firmware update is running* phase. During the firmware update CIMC may become unreachable when it reboots into the new version, this unreachable state has a timeout of 40 minutes before the firmware update is declared as timed out. When the firmware update has started, the monitoring of that update will time out after 120 minutes.
- **Firmware update has failed with an error:** This indicates that an error occurred and the firmware update has failed. CIMC usually does not give an indication of success or failure so this state usually indicates an error occurred before the firmware update actually running.
- **Firmware update has finished:** The firmware update finished without running into any errors or time outs. CIMC usually does not give an indication of success or failure, it is best to verify that the UCS firmware versions are updated when those details become available in the Cluster Status page - it can take up to 15 minutes for those details to become available.

Below the overall status in the *View Firmware Upgrade Logs* pop-up is an *Update progress* section that will contain timestamped log messages indicating the progress of the firmware update. When the *Rebooting Host In Progress* status is displayed in these log messages, CIMC is in control of the update and the cluster is monitoring that update - most log messages after this come directly from CIMC and are only added to the list of log messages if the status of the update changes.

Below the *Update progress* section of the *View Firmware Upgrade Logs* pop-up a *Component update status* section will be shown when CIMC starts providing individual component update statuses. This section summarizes the status of the update of the various UCS components on the bare metal.

Data Backup and Restore

Data backup and restore is a disaster recovery mechanism which copies data from Secure Workload cluster, connectors, and external orchestrators to an off-site storage. If a disaster occurs, data is restored from the off-site storage to a cluster of the same form-factor. You can also switch between different backup sites.

- Data backup and restore is supported for physical clusters—8 and 39 RU.
- Data can be backed up to any external object store compatible with the S3V4 API.

- Secure Workload requires sufficient bandwidth and storage to back up data. Slow network speeds and high latency can result in failed backups.
- Data storage limits are based on the selected type of backup.
 - For data backup using the continuous mode, we recommend 200 TB of storage for full backups, including flow data. To determine the actual storage space required, use the **Capacity Planner** option available on the Data Backup page. For more information, see [Use Capacity Planner, on page 20](#). Lack of storage space for multiple backups result in frequent deletion of old backups to be able to manage backups within the storage limit. There must be sufficient storage for at least one backup.
 - For lean mode backups, 1 TB of storage is sufficient because flow data, which constitutes most of the backup data, is not included in the backup.
- Data can only be restored to a cluster of compatible form-factor, running the same version as the primary. For example, you can restore data from an 8 RU cluster only to another 8 RU.

Data Backup

A schedule for data backup can be configured using the Data Backup section on the UI. The backups are triggered either once a day and at the scheduled time based on the configured settings or can be configured to run continuously. A successful backup is called a *checkpoint*. A checkpoint is a point in time snapshot of the cluster's primary datastores.

A successful checkpoint can be used to restore the data onto another cluster or the same cluster.

The cluster configuration data are always backed up for every checkpoint. Flow and other data contribute to the bulk of the data backed up. Therefore, if configured appropriately, only incremental changes are backed up. Incremental backups help reduce the amount of data pushed to the external storage, which avoids overloading the network. Optionally, a full backup can be triggered on a schedule for all data sources when incremental backup is configured. A full backup copies every object in a checkpoint, even if it is already copied and the object has not changed. This can add significant load on the cluster, on the network between the cluster and the object store, and the object store itself. A full backup may be necessary if there are corruptions in the objects or the object store has any unrecoverable hardware failures. Additionally, if the bucket provided for backup changes, a full backup is automatically enforced since a full backup is necessary before incremental backups will be useful.

Table 7: Cluster Data Backed Up in Different Modes

Secure Workload Cluster Data	Is the Data Backed Up in the Full Backup Mode?	Is the Data Backed Up in the Lean Mode?
Cluster configurations	Yes	Yes
RPMs used for imaging the cluster	Yes	Yes
Software agent deployment images	Yes	Yes
Flow database	Yes	No
Data required for automatic policy discovery	Yes	No

Secure Workload Cluster Data	Is the Data Backed Up in the Full Backup Mode?	Is the Data Backed Up in the Lean Mode?
Data to help with forensics such as file hashes, data leak models	Yes	No
Data to help with attack surface analysis	Yes	No
CVE databases	Yes	No

**Note**

- The secure connector information is not backed up or restored in the on-premise version of Secure Workload, but is backed up and restored in the SaaS version of Secure Workload.
- The virtual patch information of FMC connectors is not restored after restoring the backed-up data.

Prerequisites for Data Backup

- To obtain an activation key for the Data Backup and Restore (DBR) feature, send an email to ciscosecureworkload-licensing-support@cisco.com requesting a DBR activation key. Attach the cluster ID file in the email.

**Note**

The license entitlement is required only for the primary (active) cluster and not for the standby cluster.

- The access and secret keys for the object store are required. The Data backup and restore option does not work with the preauthenticated link for object store.
- Configure any policing to throttle the bandwidth that is used by the Secure Workload appliance to an object store. Policing with low bandwidth when the volume of data to be backed up is high can cause backup failures.
- Configure the cluster FQDNs and ensure that software agents can resolve the FQDNs.

**Note**

After you enable data backup and restore, only the current and later software agent versions are available for installation and upgrades. Versions earlier to the current cluster version remain hidden due to incompatibility.

Software Agent or Kafka FQDN Requirements

Software agents use IP addresses to get control information from the Secure Workload appliance. To enable data backup and restore and allow for seamless failover after disaster, agents must switch to using FQDN. Upgrading the Secure Workload cluster is not sufficient for this switch. Software agents support the use of FQDN starting Secure Workload version 3.3 and later. Therefore, to enable agent failover and to ensure that agents are ready for data backup and restore, upgrade the agents to version 3.3 or later.

If FQDNs are not configured, the default FQDNs are:

IP Type	Default FQDN
Sensor VIP	wss{{cluster_ui_fqdn}}
Kafka 1	kafka-1-{{cluster_ui_fqdn}}
Kafka 2	kafka-2-{{cluster_ui_fqdn}}
Kafka 3	kafka-3-{{cluster_ui_fqdn}}

The FQDNs can be changed on the **Platform > Cluster Configuration** page.

Figure 16: FQDNs or IP for Data Backup and Restore on Cluster Configuration Page

Field Name	Value
Cluster UUID	3b478c4d-6883-8861-c6e4-41bbea8d8d8
Admiral Alert Email	bugs-support@tetrationanalytics.com
CIMC Internal Network	10.13.4.0/25
CIMC Internal Network Gateway	10.13.4.2
Cluster Type	PHYSICAL
DNS Domain	cisco.com
DNS Resolver	172.21.106.115 172.21.106.116 172.26.230.8 172.26.230.9 171.70.168.183 173.36.131.10
Strong SSL Ciphers for Agent Connections	False
External IPs	
Leaf 1/2 Interconnect Network Mask	255.255.255.248
Internal Network	1.1.1.0/24
Kafka 1 FQDN	kafka-1-bean.tetrationanalytics.com
Kafka 1 IP	172.21.98.174
Kafka 2 FQDN	

Update the DNS record for the FQDNs with the IPs provided on the same page. The following table lists the mapping of IPs and FQDNs.

Field Name	Corresponding IP Field	Description
Sensor VIP FQDN	Sensor VIP	Update the FQDN to connect to cluster control plane
Kafka 1 FQDN	Kafka 1 IP	Kafka node 1 IP
Kafka 2 FQDN	Kafka 2 IP	Kafka node 2 IP
Kafka 3 FQDN	Kafka 3 IP	Kafka node 3 IP



Note FQDN for sensors VIP and Kafka hosts can only be changed before data backup and restore is configured. After the configuration, FQDN cannot be changed.

Object Store Requirements

The object store must provide a S3V4 compliant interface.



Note A few S3V4-compliant object stores do not support the DeleteObjects functionality. The DeleteObjects functionality is required to delete outdated checkpoint information. The lack of this functionality can lead to failures when attempting to delete outdated checkpoints from the storage and can cause the storage to run out of space.

- **Location**

The location of the object store is critical to the latency involved in backing up and restoring from the store. To improve restore time, ensure that the object store is located closer to the standby cluster.

- **Bucket**

Create a new and dedicated bucket for Secure Workload in the object store. Only the cluster should have *write* access to this bucket. The cluster will write objects and manage retention on the bucket. Provision at least 200 TB of storage for the bucket and obtain an access and secret key for the bucket. Data backup and restore in Secure Workload will not work with pre-authenticated links.



Note If you are using Cohesity as an object store, disable multi-part uploads while scheduling.

- **HTTPS**

The data backup option supports only HTTPS interface with the object store. This is to ensure that data in transit to the object store is encrypted and secure. If the storage SSL/TSL certificate is signed by trusted third-party CA, the cluster will use them to authenticate the object store. In case the object store uses self-signed certificate, the public key or the CA can be uploaded by selecting the **Use Server CA Certificate** option.

- **Server-side Encryption**

It is strongly recommended to turn ON server-side encryption for the bucket assigned to Secure Workload cluster. The cluster will use HTTPS to transfer data to object store. However, the object store should encrypt the objects to ensure that the data at rest is secure.


Configuration of Data Backup

To configure data backup in Secure Workload, perform the following:

1. **Planning:** The data backup option provides a planner to test the access to the object store, determine the storage requirement, and the backup duration needed for each day. This can be used to experiment before configuring a schedule.

To use data backup and restore calculators, navigate to **Platform > Data Backup**. If data backup and restore is not configured, this will navigate to the Data Backup landing page.

Figure 17: Backup Landing Page



To plan the data backup, use the following options:

- [Use Storage Planner, on page 19](#)
- [Use Capacity Planner, on page 20](#)



Note If you are unable to view the Data Backup option under Platform, ensure that you have the license to enable data backup and restore.

2. **Configuring and scheduling data backup:** Secure Workload will copy data to object store only in the configured time-window. While configuring backup for the first time, the pre-checks will run to ensure the FQDNs are resolvable and resolves to the right IP. After the initial validation, an update is pushed to registered software agents to switch to using FQDNs. Without FQDN, the agents cannot failover to another cluster after a disaster event. To support this, agents must be upgraded to the latest version supported by the cluster and all the agents should be able to resolve the sensor VIP FQDN. As of Secure Workload release 3.3 and later, only deep visibility and enforcement agents support data backup and restore and will switch to using FQDN.

To create a schedule and configure data backup, see [Configure Data Backup, on page 21](#).

Use Storage Planner

Procedure

- Step 1** To ensure that the storage is compatible with Secure Workload, perform one of the following actions:
 - On the **Data Backup** landing page, click **Storage Planning**.
 - From the **Planning** drop-down menu, choose **Storage**.

The **Storage Planning** page is displayed.
- Step 2** Enter the following details:
 - A name for the storage.
 - URL of an S3 compliant storage endpoint.
 - An S3 compliant bucket name configured on the storage.
 - (Optional for certain storage) Region of the S3 compliant storage.

- Access key to the storage.
- Secret key of the storage.

- Step 3** (Optional) If required, you can enable HTTP proxy.
- Step 4** (Optional) To use multi-part uploads of the backed data, enable **Use Multipart Upload**.
- Step 5** (Optional) If a CA certificate is required to authenticate the storage server, enable **Use Server CA Certificate** and enter the certificate details.
- Step 6** Click **Test**.

The storage validation will test:

- Authentication and access to the object store and bucket.
- Upload to and download from the configured bucket.
- Bandwidth checks.

The storage planning process can take about five minutes to complete.

Use Capacity Planner

Procedure

- Step 1** To plan the storage size and the backup window estimates, perform one of the following actions:
- On the **Data Backup** landing page, click **Capacity Planning**.
 - From the **Planning** drop-down menu, choose **Capacity**.
- The **Capacity Planning** page is displayed.
- Step 2** Enter the maximum bandwidth limit to back up the data.
- This bandwidth must at most be the policer configuration that will throttle data to the object store.
- Step 3** Registered software agents count is automatically populated. Based on forecasts, you can change the agents count.
- Step 4** (Optional) Enable **Lean Data Mode** to exclude the non-configuration data from being backed up. Using this option reduces the storage limitation by 75%.
- Step 5** The maximum storage configured for the storage bucket. This will automatically set the retention period for the backups.

After the required details are entered, the Estimated Backup Duration displays the time required to backup data of a day. This is an estimate based on typical agent load, estimated agents count, and the maximum bandwidth configured. The Estimated Maximum Storage displays the estimate of maximum storage required by Secure Workload to support specified retention and estimated agents count.

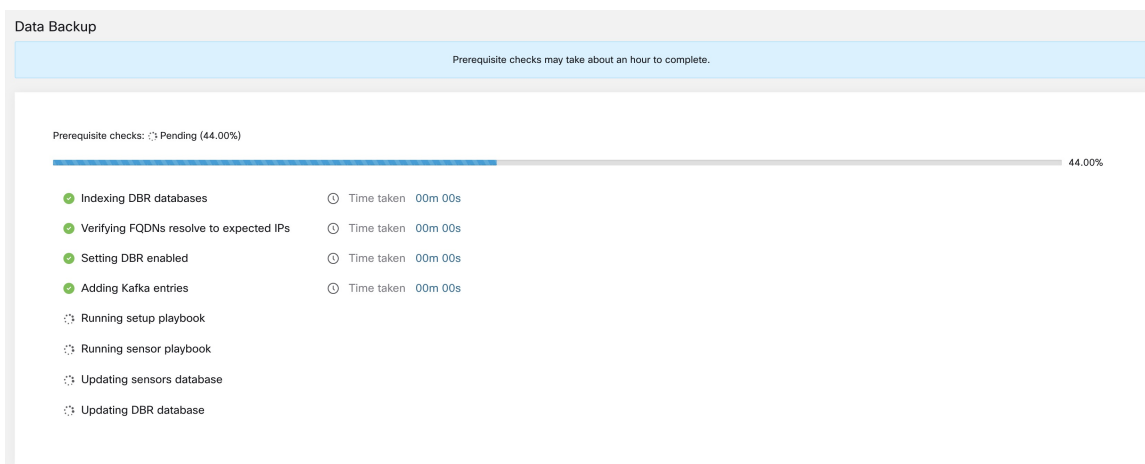
Configure Data Backup

Procedure

- Step 1** On the data backup landing page, click **Create new schedule**.
- Step 2** To confirm the prerequisite checks to run, check the **Approve** buttons and click **Proceed**.

The prerequisite check takes about 30 minutes to complete and are run only during the first time a schedule is configured.

Figure 18: Backup Prerequisites Run



- Step 3** To configure the storage, enter the following details and click **Test**.
- A name for the storage.
 - URL of an S3 compliant storage endpoint.
 - An S3 compliant bucket name configured on the storage.
 - (Optional for certain storage) Region of the S3 compliant storage.
 - Access key to the storage.
 - Secret key of the storage.
 - (Optional) Enable HTTP proxy, if required.
 - (Optional) To use multi-part uploads of the backed data, enable **Use Multipart Upload**.
 - (Optional) If a CA certificate is required to authenticate the storage server, enable **Use Server CA Certificate** and enter the certificate details.

Figure 19: Storage Configuration

1 Configure Storage 2 Configure Backup 3 Schedule Backup 4 Review

Name
Name is required.

URL
URL is required.

Bucket
Bucket is required.

Region

Access Key
Access Key is required.

Secret Key

Use HTTP Proxy

Use Multipart Upload

Use Server CA Certificate

Test

Cancel Next

Step 4

To configure the storage capacity, enter the following details:

- The maximum bandwidth limit to back up the data. This bandwidth must at most be the policer configuration that will throttle data to the object store.
- Registered software agents count is automatically populated. Based on forecasts, you can change the agents count.
- (Optional) Enable **Lean Data Mode** to exclude the non-configuration data from being backed up. Using this option reduces the storage limitation by 75%.
- The maximum storage configured for the storage bucket. This will automatically set the retention period for the backups.

Figure 20: Capacity Planning

CONFIGURE SCHEDULE

1 Configure Storage 2 Configure Backup 3 Schedule Backup 4 Review

Est. Observed Bandwidth 81 Mbps
Max. Bandwidth Limit 1000 Mbps

Est. Sensor Count 35

Lean Data Mode

Retention 8 days

Est. Backup duration 23 : 53

Est. Max Storage 182 TB

Cancel Previous Next

Step 5

To schedule the backup, enable the following:

- By default, **Set starting backup point from today** is enabled. This option will ignore all files created before midnight UTC on the day of configuration. In a working cluster, there could be high volume of data to be backed up on the first day and might overwhelm the cluster, network, and the object store. If you want to backup all existing data, disable this checkbox but note the impact on the network, object store and cluster.

Note All configuration data will be backed up irrespective of this option.

- Continuous backup - If enabled, the data will be backed up at 15 minutes after the previous backup is completed. This option allows for backups to be running continuously, instead of being scheduled at specific time. The **Time zone** and **Allowed Start backup window** options will not be available when Continuous backup is enabled.
- The next two options are used to configure schedule for the backup, if continuous backup is not used.
 - Time zone: Defaults to the web browser time zone
 - Allowed Start backup window: Time (in hour or minutes) when the backup will start. Time must be entered in the 24-hour format
 - Enable recurring full backup (not selected by default): If enabled, a schedule for full backup can be configured. By default, after the first full backup, all backups are incremental. Enabling this configuration will force a full backup at the specified schedule.

Figure 21: Schedule Backup

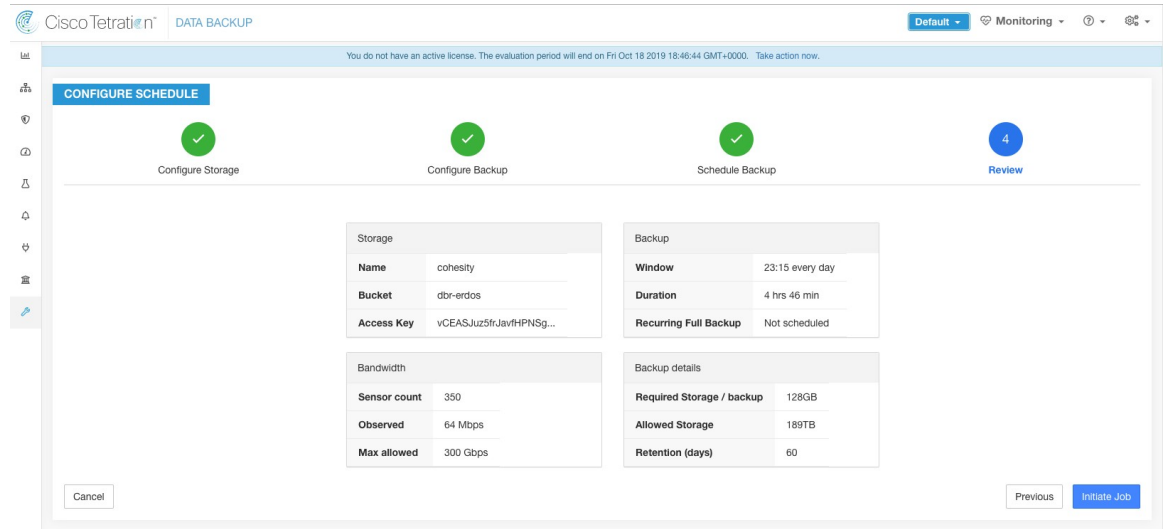
The screenshot shows a web interface titled "CONFIGURE SCHEDULE" with four steps: "Configure Storage" (completed), "Configure Backup" (completed), "Schedule Backup" (active), and "Review" (disabled). The "Schedule Backup" section contains the following options:

- Set starting backup point from today:**
- Continuous backup:**
- Timezone:** A dropdown menu showing "America/Los_Angeles".
- Allowed start backup window:** A form with "Every: Day" in a dropdown, "at: 0" in a text input, and "0" in another text input.
- Enable recurring full backup:**

Buttons for "Cancel", "Previous", and "Next" are visible at the bottom of the form.

Step 6 Review the configured backup schedule and settings, and then click **Initiate Job**.


Figure 22: Backup Configuration Review



Backup Status

After the configuration of data backup, backup is triggered everyday at a scheduled time, unless continuous mode is enabled. Status of the backups can be seen on the Data Backup dashboard by navigating to **Platform > Data Backup**.

Figure 23: Backup Status



Time since last successful checkpoint should be less than 24 hours + the time it takes to checkpoint. For example, if the checkpoint + backup takes around 6 hours, then the time since last successful checkpoint should be less than 30 hours.

The following graphs provide additional information:

- Checkpoint Duration: This graph shows the trendline for the amount of time the checkpoint takes.
- Upload Duration: This graph shows the trendline for how long it takes to upload the checkpoint to the backup.
- Checkpoint Size: This graph shows the trendline for the size of the checkpoint.
- Upload Bandwidth: This graph shows the trendline for the upload bandwidth.

The table shows all the checkpoints. Checkpoint labels can be edited and the labels will be available while choosing a checkpoint to restore data on the standby cluster.

A checkpoint transitions through multiple states and these are the possible states:

- Created/Pending: Checkpoint is just created and waiting to be copied

- Running: Data is getting actively backed up to external storage
- Success: Checkpoint is complete and is successful; can be used for data restore
- Failed: Checkpoint is complete and has failed; cannot be used for data restore
- Deleting/Deleted: An aged-out checkpoint is being deleted or is deleted

To change the schedule or the bucket, click on **Edit Schedule**. To complete the wizard, see the Configure Data Backup section.

To troubleshoot any errors during the creation of checkpoints, see [Troubleshooting: Data Backup and Restore, on page 31](#).

Deactivate Backup Schedule

Backups can be deactivated by clicking the **Deactivate Schedule** button. It is recommended to deactivate the backup schedule before making changes to the schedule. Deactivate a schedule only when no checkpoint is in progress. Running a test or disabling the schedule while a checkpoint is in progress may cause the checkpoint in progress to fail and the upload to be in an undefined state.

Object Store Retention

Secure Workload cluster manages the lifecycle of objects in the bucket. You must not delete or add objects to the bucket. Doing so may lead to inconsistencies and corrupt successful checkpoints. In the configuration wizard, the maximum storage to be used must be specified. Secure Workload will ensure the usage of bucket will stay within the configured limit. There is a storage retention service that ages out objects and deletes them from the bucket. After the storage usage reaches a threshold (80% of the bucket capacity), computed based on the configured maximum storage and incoming data rate, the retention will try to delete *un-preserved* checkpoints to reduce the usage to below the threshold. The retention will also keep a minimum of two successful checkpoints at any time and all the preserved checkpoints, whichever is more. If retention cannot delete any checkpoints to make space, *checkpoints will start failing*.

Preserve Checkpoints

As new checkpoints are created, old ones will age-out and are deleted. However, checkpoints can be preserved, preventing it from being deleted by retention. A preserved checkpoint will not be deleted. If there are multiple preserved checkpoints, at some point the storage will be insufficient for new objects and aged-out checkpoints cannot be deleted because they were preserved. As a best practice, preserve checkpoints on a need basis and update the Label for the checkpoint with the reason and validity as a reference. To preserve a checkpoint, click on the lock icon against the required checkpoint.

Restore Data

The data restore option is available under the **Platform** menu in the left navigation bar.

A cluster must be in the **DBR standby mode** to be restored using backed up data. Currently, a cluster can be set to standby mode **only during initial setup**.

Following combinations are allowed:

Primary Cluster SKU	Standby Cluster SKU
8RU-PROD	8RU-PROD, 8RU-M5
8RU-M5	8RU-PROD, 8RU-M5
39RU-GEN1	39RU-GEN1, 39RU-M5
39RU-M5	39RU-GEN1, 39RU-M5

Deploy Cluster in Standby Mode



Note Contact [Cisco Technical Assistance Center](#) to initiate data restore.

You can deploy a cluster in the Standby mode by configuring the recovery options in site information. While configuring site information during deployment, configure the restore details under the **Recovery** tab in the setup UI during deployment.

There are three modes (See the *Standby Deployment Modes* section) to deploy a standby cluster and for all the three modes, configure these settings:

- Set the **Standby Config** to **On**. You cannot change this configuration after it is set until the cluster is redeployed.
- Configure primary cluster name and FQDNs. You can change this configuration later.

Figure 24: Enable Standby Mode

Site Config

Complete this form to create or update the site config.

General

Email

L3

Network

Service

Security

UI

Advanced

Recovery

Continue Back

Standby Config On

Enable restore standby mode, Cluster will not functional until failed over.

Primary cluster site name

hui

Primary cluster site name

Sensor VIP FQDN

wsshui.tetrationanalytics.com

The fully qualified domain name that has been setup for WSS this cluster. This name should point to the cluster's sensor VIP. Sensors will connect to this FQDN when DBR is enabled. This takes effect only when DBR is enabled. Before changing this FQDN make sure it resolves to the sensor VIP IP address. Failure to resolve will prevent updating this field.

Kafka 1 FQDN

kafka-1-hui.tetrationanalytics.com

The fully qualified domain name that has been setup for kafka-1 instance in this cluster. This name should point to the cluster's Kafka instances. This FQDN will take effect only when DBR is enabled. Before changing this FQDN make sure it resolves to the corresponding kafka-1 IP address. Failure to resolve will prevent updating this field.

Kafka 2 FQDN

kafka-2-hui.tetrationanalytics.com

The fully qualified domain name that has been setup for kafka-2 instance in this cluster. This name should point to the cluster's Kafka instances. This FQDN will take effect only when DBR is enabled. Before changing this FQDN make sure it resolves to the corresponding kafka-2 IP address. Failure to resolve will prevent updating this field.

Kafka 3 FQDN

kafka-3-hui.tetrationanalytics.com

The fully qualified domain name that has been setup for kafka-3 instance in this cluster. This name should point to the cluster's Kafka instances. This FQDN will take effect only when DBR is enabled. Before changing this FQDN make sure it resolves to the corresponding kafka-3 IP address. Failure to resolve will prevent updating this field.

← Previous

- The rest of the deployment is the same as a regular deployment of a Secure Workload cluster.

- A banner is displayed on the Secure Workload UI after the cluster enters the standby mode.
- Primary cluster name and FQDNs can be reconfigured after the deployment to enable the standby cluster to track another cluster. This can be reconfigured later before failover is triggered from the **Cluster Configuration** page.

Standby Deployment Modes

- **Cold Standby:** There is no standby cluster. However, the primary cluster backs the data to S3. During a disaster, a new cluster (or the same cluster as the primary) must be provisioned, deployed in standby mode, and restored.
- **Warm Standby:** A standby cluster is operational and deployed in standby mode. It periodically fetches state from the S3 cluster and places it in the ready state to be operational if there is a disaster. During a disaster, log in to this new cluster and trigger a failover.
- **Luke Warm Standby:** Multiple primary clusters are backed by fewer standby clusters. The standby cluster is deployed in standby mode. Only after a disaster, the storage bucket information is configured, data is prefetched, and the cluster is restored.

Prefetch Cluster Data

Before the cluster can be restored, it must prefetch data. The checkpoint data is prefetched from the same storage bucket that is used for backing up data. Credentials must be provided for the backup service to download from the storage. If a storage is not set up for prefetch, the **Data Restore** tab will launch the setup wizard.



Note Standby cluster interacts only with the S3 storage. When the backup on Primary cluster is updated to use a different storage or bucket, the storage on standby cluster must be updated.

After the information is validated, storage is automatically configured for prefetch. The restore tab will display the prefetch status.

Figure 25: Prefetch Status

The screenshot shows the 'Data Restore' interface for a 'Tetration Cluster'. At the top, a notification states 'Cluster is in STANDBY mode, any changes made will be discarded once the cluster fail over.' The main area features a 'Tetration Cluster' icon with a red warning triangle, indicating a prefetch status issue. To the right, a 'Data Download Status' table shows the following information:

Data Download Status	
Restore to	N/A
Last successful data download	N/A
Last data download attempt	not_triggered
Last Prefetched Checkpoint	not_triggered

Below the status table, there is a 'SETTINGS' section with a table for storage configuration:

SETTINGS	
URL	... more
Access Key	... more
Bucket	...
Region	...

A 'Reconfigure Storage' button is located below the settings table. In the top right corner, there is a 'Restore Now' button.

The status page displays the following details:

- The upper left section has a graphic indicating readiness of various components for starting a restore. To check the data, hover over the components. The associated data is displayed in the upper right section.
 - **Bucket:** Displays the prefetch status. If the latest data is more than 45 minutes old, it shows up in red. Note that latest data being more than 45 minutes old is not a concern if the backup on the active takes more than 45 minutes for each checkpoint.
 - **DNS:** Displays the Kafka and WSS FQDN resolutions with respect to standby cluster IP addresses. During restore, if the FQDNs are not updated to standby cluster IP addresses, the agent cannot connect. After the FQDNs start resolving to the standby cluster, status will turn green.
 - **Agents:** Displays the number of software agents that have successfully switched over to the standby cluster. This is only relevant after a restore has been triggered.
- The upper right section displays the information relevant to the chosen graphic in the left section. Clicking **Restore Now** will initiate the restore process.
- The lower left section displays the prefetch storage settings that are in use.
- The lower right section displays a graph of prefetch delays.

A data prefetch updates several necessary components to ensure a fast restore. If a data prefetch is unable to complete, the reason for failure is displayed on the status page.

Common errors that can cause prefetch failures:

S3 Access Error: In this case the data from the storage could not be successfully downloaded. This may happen due to invalid credentials, a change in the storage policies, or temporary network issues.

Incompatible Cluster Versions: Data can be restored to a cluster running the same version (including the same patch version) of Secure Workload as the primary cluster. This can likely happen during upgrades when

only one of the cluster is upgraded. Or, during deploy when a different version is used for deploying. Deploying the clusters to a common version will resolve the issue.

Incompatible SKU Versions: Note down the allowed SKUs for standby clusters for the primary cluster. Only specific SKUs are allowed for restore of the primary cluster SKU.

Cluster Restore

A cluster restore can be triggered by clicking **Restore Now** in the upper right section of the **Restore Status** page. Before a restore action can be triggered, an acknowledgement is asked.

Cluster data is restored in two phases:

- **Mandatory Phase:** The data needed to restart services is restored first. The time taken by mandatory phase depends on the configuration, number of software agents installed, amount of data backed up, and flow metadata. During the mandatory phase, the UI is not accessible. **Working TA guest keys are required for any support during mandatory the phase, should such a need arise.**
- **Lazy Phase:** Cluster data (including flow data) is restored in the background and will not block cluster usage. The cluster UI is accessible and a banner with the completed percentage of restore is displayed. During this phase, the cluster is operational and data pipelines function normally and the flow searches are also available.

After the Mandatory Phase of the restore is complete and the UI is accessible, the changes in the cluster must be communicated to the software agents. In the DNS server used by the agents, the IP address associated with the cluster's FQDN must be updated, and the DNS entry should point to the restored cluster. A DNS lookup is triggered by the agents when the connection to the primary cluster is broken. Based on the updated DNS entry, the agents will connect to the restored cluster.

Recovery Time Objective and Recovery Point Objective

This section describes the Recovery Time Objective (RTO) and Recovery Point Objective (RPO) for the data backup and restore solution.

A backup initiated on the primary cluster requires some time to complete depending on the amount of data being backed up and the backup configuration. The different modes of backup defines the RPO for the solution.

- If scheduled, non-continuous backup is used and the backup is initiated once in a day. If a disaster occurs then the maximum time of lost data will be approximately 24 hours, plus time taken to copy the data to the backup storage. Therefore, the RPO is at least 24 hours.
- If continuous mode backup is used then a new backup is initiated 15 minutes after the previous backup. Each backup consumes a certain amount of time to create and then a certain amount of time to upload the data to the backup storage. The first backup is a full backup and the subsequent backups are incremental backups, the incremental backups do not take much time. If a disaster occurs, the amount of the data lost will be the sum of the time taken to create the backup and the time taken to upload the backup to the storage. Typically the RPO in this case will be approximately a few minutes to an hour.

When restoring a cluster, mandatory data is first prefetched from the storage, then mandatory restore phase is triggered. The UI is not available during the mandatory restore phase. After the mandatory restore is complete, the UI is available for usage. The rest of the data is restored in the lazy restore phase. RTO in this case is the time taken until the UI is available for usage after mandatory phase is complete. RTO depends on the standby deployment mode.

- **Cold Standby Mode:** In this mode, the cluster must be deployed first which takes approximately a few hours. The cluster must then be configured with the backup storage credentials. Since this is the first time the backup is uploaded into the standby cluster, there will be a lot of mandatory data that needs to be retrieved and processed. The time for prefetch is approximately tens of minutes (depending on the quantity of data backed up). The mandatory restore phase takes approximately 30 minutes to complete. Together this forms the RTO time of approximately a few hours, primarily due to the time taken to boot and deploy the cluster.
- **Luke Warm Standby Mode:** In this mode, the cluster is already deployed but the backup storage is not configured. The cluster must be configured with the backup storage credentials. Since this is the first time the backup is uploaded into the standby cluster, there will be a lot of mandatory data that needs to be retrieved and processed. The time for prefetch is approximately tens of minutes (depending on the quantity of data backed up). The mandatory restore phase takes approximately 30 minutes to complete. Together this forms the RTO time of approximately an hour to two hours, depending on the amount of data backed up and time to pull the data from backup storage.
- **Warm Standby Mode:** In this mode, the cluster is already deployed, the backup storage is configured, and prefetch is retrieving data from the storage. The cluster can now be restored, which will trigger the mandatory restore phase, which takes approximately 30 minutes to complete. This forms the RTO time of approximately 30 minutes. Note that there is some delay from when the backup is uploaded from the active to the storage to the time the backup is pulled by the standby. This is approximately a few minutes. If the latest backup from the active (prior to it experiencing a disaster event) has not been prefetched to the standby, you must wait for a few minutes for it to be retrieved.

Upgrade with Data Backup and Restore

When data backup and restore is enabled on the cluster, it is recommended to deactivate the schedule before starting the upgrade. See [Deactivate Backup Schedule](#). This ensures that a successful backup exists before upgrade is started and that no new backup is being uploaded. A schedule must be deactivated when a checkpoint is not in progress, to avoid creating a failed checkpoint.

Troubleshooting: Data Backup and Restore

S3 Configuration Checks are Unsuccessful

If the storage test is unsuccessful, identify the failure scenarios displayed on the right pane and ensure that:

- S3 compliant storage URL is correct
- The access and secret keys of the storage are correct
- Bucket on the storage exists and correct access (read/write) permissions are granted
- Proxy is configured if the storage needs to be accessed directly
- The multi-part upload option is disabled if you are using Cohesity

Error Scenarios of S3 Configuration Checks

The table lists the common error scenarios with resolution and is not an exhaustive list.

Table 8: Error Messages with Resolution during S3 Configuration Checks

Error Message	Scenario	Resolution
Not found	Incorrect bucket name	Enter correct name of the bucket configured on the storage
SSL connection error	SSL certificate expiry or verification error	Verify the SSL certificate
	Invalid HTTPS URL	<ul style="list-style-type: none"> • Re-enter correct HTTPS URL of the storage. • Resolve any failures during verification of SSL certificate.
Connection timed out	IP address of the S3 server is unreachable	Verify the network connectivity between the cluster and S3 server
Unable to connect to URL	Incorrect bucket region	Enter correct region of the bucket
	Invalid URL	Re-enter correct URL of the S3 storage endpoint
Forbidden	Invalid secret key	Enter correct secret key of the storage
	Invalid access key	Enter correct access key of the storage
Unable to verify S3 configuration	Other exceptions or generic errors	Try to configure the S3 storage after some time

Error Codes of Checkpoints

The table lists the common error codes of checkpoints and is not an exhaustive list.

Table 9: Error Codes of Checkpoints

Error Code	Description
E101: DB checkpoint failure	Unable to snapshot MongoDB oplogs
E102: Flow data checkpoint failure	Unable to snapshot Druid database
E103: DB snapshot upload failure	Unable to upload Mongo DB snapshot
E201: DB copy failure	Unable to upload Mongo snapshot to HDFS
E202: Config copy failure	Unable to upload Consul-Vault snapshot to HDFS
E203: Config checkpoint failure	Unable to checkpoint consul-vault data
E204: Config data mismatch during checkpoint	Cannot generate consul/vault checkpoint after maximum retry attempts

Error Code	Description
E301: Backup data upload failure	HDFS checkpoint failure
E302: Checkpoint upload failure	Copydriver failed to upload data to S3
E401: System upgrade during checkpoint	Cluster got upgraded during this checkpoint; checkpoint cannot be used
E402: Service restart during checkpoint	Bkpdriver restarted in the create state; checkpoint cannot be used
E403: Previous checkpoint failure	Checkpoint failed on previous run
E404: Another checkpoint in progress	Another checkpoint is in progress
E405: Unable to create checkpoint	Error in checkpoint subprocess
Failed: Completed	Some preceding checkpoint failed; likely an overlap of multiple checkpoints starting together.

High Availability in Secure Workload

Secure Workload provides high availability when there is a probability of services, nodes, and VMs failing. High availability provides recovery methods by ensuring minimum downtime and minimal intervention by the site administrator.

In Secure Workload, services are distributed across the nodes in a cluster. Multiple instances of services run simultaneously across the nodes. A primary instance and one or more secondary instances are configured for high availability across multiple nodes. When the primary instance of a service fails, a secondary instance of the service renders as primary and becomes active immediately.

Secure Workload Cluster Design

The key components of a Secure Workload cluster are:

- Bare metal servers that host multiple VMs, which in turn, host many services.
- Cisco UCS C-Series Rack Servers with Cisco Nexus 9300 Series switches that contribute to an integrated high-performance network.
- Hardware-based appliance models in either a small or large form factor to support a specific number of workloads:
 - Small form factor deployment with six servers and two Cisco Nexus 9300 switches.
 - Large form factor deployment with 36 servers and three Cisco Nexus 9300 switches.

Figure 26: Design of Secure Workload Cluster Design

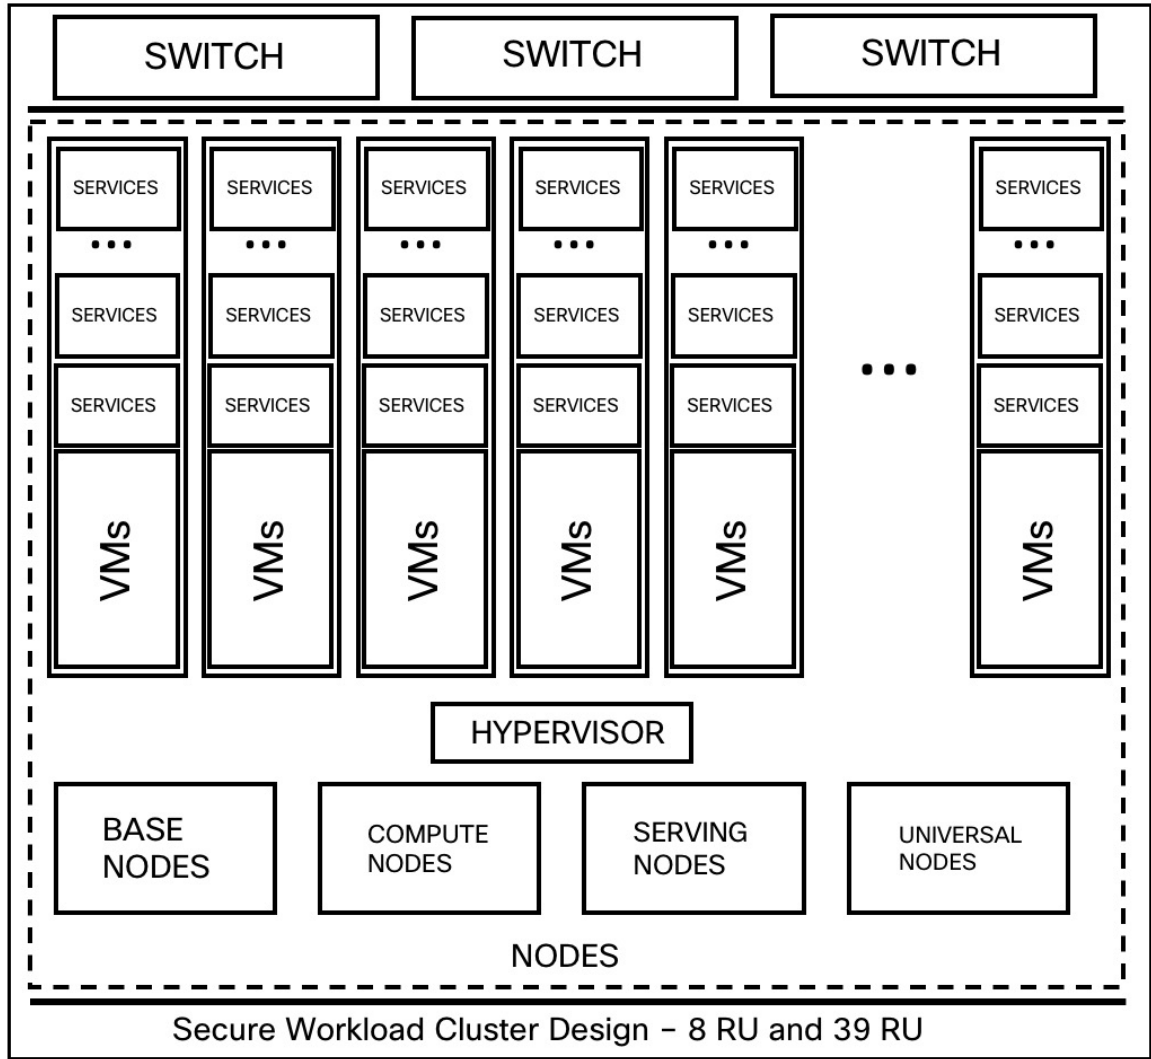


Table 10: Secure Workload Cluster Components

Attributes/ Form Factor	8 RU	39 RU
Number of nodes	6	36
Number of compute nodes	—	16
Number of base nodes	—	12
Number of serving nodes	—	8
Number of universal nodes	6	—
Number of VMs	50	106
Number of collectors	6	16

Attributes/ Form Factor	8 RU	39 RU
Number of network switches	2	3

Limitations of High Availability in Secure Workload

- In both the form factors (8RU and 39RU) of a cluster, if a failed node is hosting a Hadoop NameNode VM, manual intervention is required to fail over to a secondary namenode VM.



Note The failover is not automatic in Secure Workload Release 3.8.x and earlier.

- For a 2 VM or 3 VM service, such as orchestrators, Redis, MongoDB, Elasticsearch, enforcementpolicystore, AppServer, ZooKeeper, TSDB, Grafana, and so on, only a single VM failure is supported; a second VM failure renders the service inactive.

Impact and Recovery Details for Failure Scenarios

- There is no impact to the cluster operation at any point in time.
- There is no single point of failure. If any of the nodes or VMs within a cluster fail, it does not result in the failure of the entire cluster.
- There is minimal downtime in recovery from failure because of services, nodes, or VMs.
- There is no impact on the connections that are maintained by software agents to a Secure Workload cluster. The agents communicate with all the available collectors in the cluster. If a collector or VM fails, the software agents' connections to the other instances of the collectors ensure that the flow of data is not interrupted and there is no loss of functionality.
- The cluster services communicate with external orchestrators. When the primary instance of a service fails, the secondary instances take over to ensure the communication with external orchestrators is not lost.

Types of Failure Scenarios

High availability supports the following failure scenarios:

- Services Failure
- VM Failure
- Node Failure
- Network Switch Failure

Services Failure

When a service fails on a node, another instance of that particular service picks up the functions of the failed service and continues to run.

Figure 27: Normal Operation




Figure 28: Failure Scenario of a Service

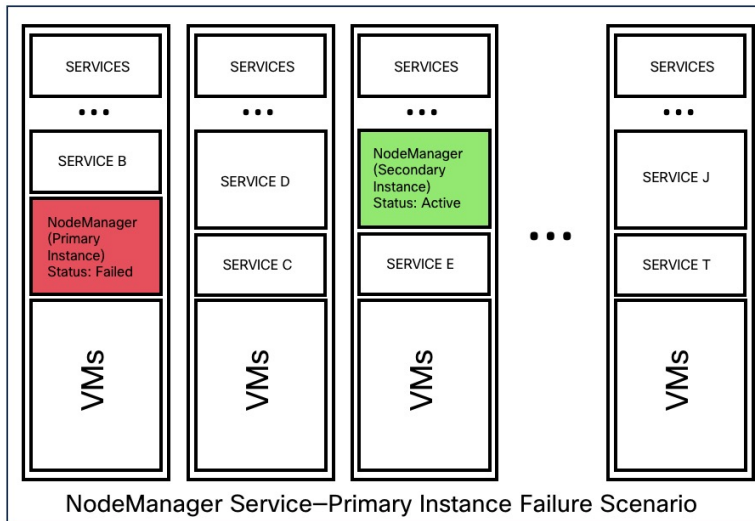


Table 11: Services Failure Impact and Recovery

Impact	No visible impact.
Recovery	<ul style="list-style-type: none"> • Minimal downtime for the UI or dependent services to continue to run from the secondary instances. • Recovery is automatic.

VM Failure

When one of the VMs fails, secondary VMs are available. The services on the secondary VMs pick up the services that the failed VM was running. Meanwhile, Secure Workload restarts the failed VM to recover it. For example, as illustrated in the [Figure: Failure Scenario of a VM](#), when a VM, in this instance, VM1, fails, the services running on it also fails. The secondary VMs continue to be operational and the secondary instances pick up the services that the failed VM was running.

Figure 29: Normal Operation


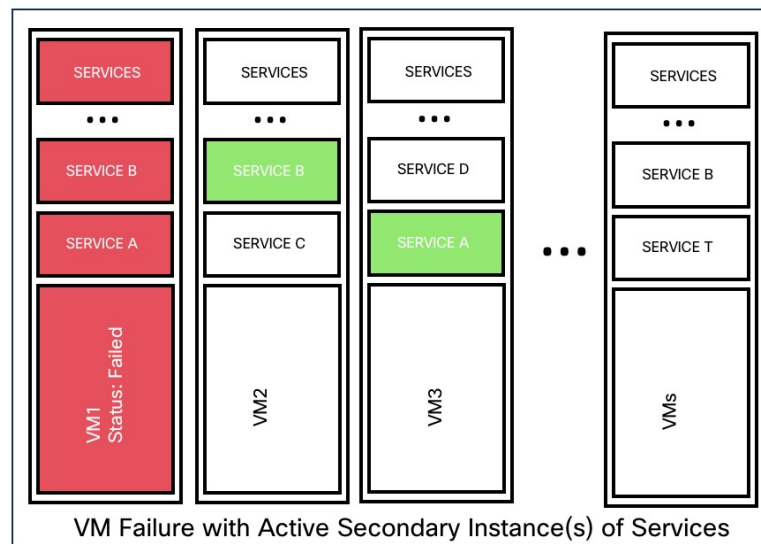


Figure 30: Failure Scenario of a VM



For services provided by symmetric VMs, such as `collectordatamovers`, `datanode`, `nodemanager`, and `druidHistoricalBroker` VMs, multiple VMs can fail but the applications will continue to function at reduced capacity.

Table 12: Symmetric VM Types

Service Type	Total VMs	Number of VM Failures Supported
Datanode	6	4
DruidHistorical	4	2
CollectorDataMover	6	5
NodeManager	6	4
UI/ AppServer	2	1



Note The nonsymmetric VM types tolerate only one VM failure before the corresponding services are rendered unavailable.

Table 13: VM Failure Impact and Recovery

Impact	No visible impact.
Recovery	<ul style="list-style-type: none"> Minimal downtime for the UI or dependent services to continue to run from the secondary instances on other VMs. Recovery is automatic. However, if a VM remains inactive, contact Cisco Technical Assistance Center to troubleshoot the issue. In a few instances, you may have to replace the bare metal.

Node Failure

Figure 31: Normal Operation




Figure 32: Failure Scenario of a Node




Table 14: Number of Node Failures Tolerated

Node Failures	8 RU	39 RU
Number of node failures that are tolerated for high availability	1	1*

* In 39 RU clusters, single node failure is always tolerated. A second node failure might be allowed as long as the two failed nodes do not host VMs for a 2 VM or 3 VM service, such as orchestrators, Redis, MongoDB, Elasticsearch, enforcementpolicystore, AppServer, ZooKeeper, TSDB, Grafana, and so on. In general, the second node failure results in a critical service becoming unavailable because of two VMs being affected.



Caution We recommend that you immediately restore the failed node because the failure of a second node will most likely result in an outage.

Table 15: Node Failure Impact and Recovery

Impact	No impact in the functionality of the cluster. However, contact Cisco Technical Assistance Center to replace the failed node immediately. Failure of a second node will most likely result in an outage.
Recovery	<ul style="list-style-type: none"> Minimal downtime. If a node fails, we recommend that you contact Cisco Technical Assistance Center for assistance to remove the faulty node and replace it with another node.

Network Switch Failure

The switches in Secure Workload always remain active. In the 8RU form-factor deployment, there is no impact if a switch fails. In the 39RU form-factor deployment, the clusters experience half the input capacity if a switch fails.



Note The switches in the Secure Workload cluster do not have the recommended port density to support the VPC configuration for public networks.

Figure 33: Normal Operation




Figure 34: Failure Scenario of a Switch

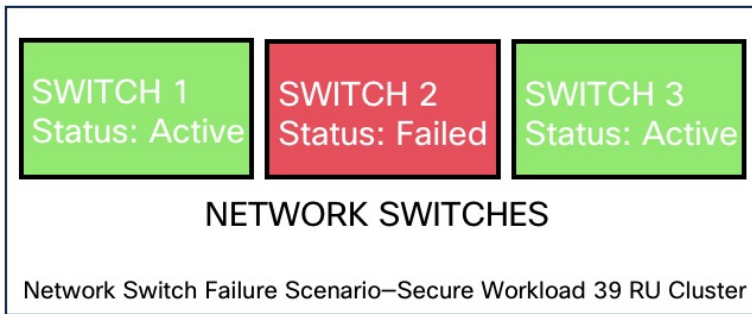


Table 16: Number of Switch Failures Tolerated

Form Factor	8 RU	39 RU
Number of switch failures that are tolerated for high availability	1 Note If two or more switches fail, it is likely to have an impact on the entire functionality of the cluster.	1 Note A single switch failure results in half input capacity. Two or more failures are likely to impact the entire functionality of the cluster.

Table 17: Network Switch Failure Impact and Recovery

Impact	<ul style="list-style-type: none"> • A faulty switch or network card on a bare metal causes loss of network connectivity within the cluster. • There is no impact in the functionality of a cluster because of a single switch failure. However, two or more failures are likely to impact the entire functionality of the cluster. • Connectivity issues to multiple VMs on a cluster, or intermittent and prolonged connectivity problems result in unpredictable behaviour within the cluster.
Recovery	<ul style="list-style-type: none"> • Recovery is automatic. • Contact Cisco Technical Assistance Center for assistance with faulty switches or network cards on bare metals.

VM Information

The **Virtual Machine** page under the **Troubleshoot** menu displays all virtual machines that are part of the Cisco Secure Workload cluster. It displays their deployment status during cluster bring up or upgrade (if any) and also public IPs. Note that all VMs in the cluster are not part of a public network therefore they may not have a public IP.

Upgrading Cluster

To access upgrade options, choose **Platform > Upgrade/Reboot/Shutdown** in the left navigation bar.

There are two types of upgrade. This section describes the **full** upgrade process. During this upgrade all VMs in the cluster except for Orchestrator-VMs are shut down, new VMs are deployed, and the services are re-provisioned. All the data within the cluster are persisted during this upgrade. Except a downtime of around 2 hours during this upgrade.

Initiating Upgrade

To initiate an upgrade, choose **Platform > Upgrade/Reboot/Shutdown** in the left navigation bar. You can upgrade, patch upgrade, shutdown, or reboot the cluster.

To initiate a full upgrade, click **Send Upgrade Link**. During the full upgrade process, the VMs are powered off, except the orchestrator VMs, and upgrade them and redeploy them. This results in 2+ hours of cluster downtime. Patch upgrade minimizes the downtime, but just updating the services that must be patched and will not result in VM restarts. The downtime is usually in the order of few minutes. To initiate Patch Upgrade, click Send Patch Upgrade Link. Use Send Reboot Link to initiate cluster reboot after a power down. Clicking on either of these links generate an email with a link in it and will send it to the user who initiated the upgrade.

Figure 35: Initiate a Full Upgrade

Hello Site Admin!

We received a request that you intend to upgrade the cluster "50". You can do this through the link below.

[Upgrade 50](#)

The above link expires by Mar 26 09:29:50 pm (PDT).

If you didn't request this, please ignore this email.

Upgrade will not be triggered until you actually click the above link.

Cisco TetrationOS Software, Version 2.2.1.34.devel

TAC Support: <http://www.cisco.com/tac>

Copyright (c) 2015-2018 by Cisco Systems, Inc.

All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. Cisco products are covered by one or more patents.

Before sending the email, orchestrator runs several verification checks to make sure the cluster is upgradable. The checks include:

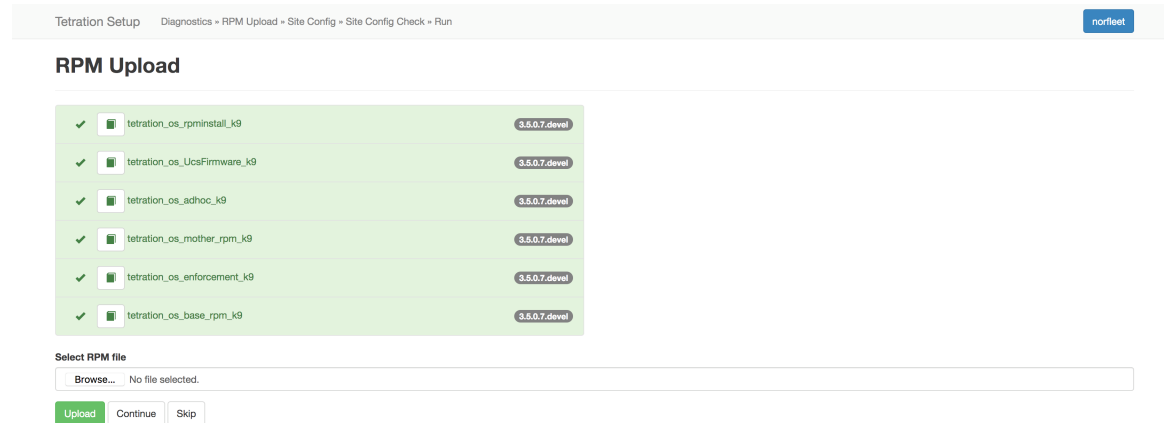
- Checks to see there are no decommissioned nodes.
- Checks each bare metal to make sure there are no hardware failures. This covers:
 - Drive failure
 - Drive predicted Failure.
 - Drive missing.
 - StorCLI failures
 - MCE log failures
- Checks to ensure we have all the BMs in commissioned state. Nothing fewer than 36 servers for 39RU and 6 for 8RU.

If there are any of these failures, an upgrade link will not be sent and you see 500 error with information like HW failure, or missing host and check orchestrator logs for more info. In this scenario, use explore to tail -100 on /local/logs/tetration/orchestrator/orchestrator.log in the host orchestrator.service.consul. This provides detailed information about which one of the 3 checks caused the failure. This usually requires fixing the hardware and recommissioning the node. After that is done we can restart upgrade by clicking on "Send Upgrade Link".

Upload RPMs

Click on the link in the email will connect to the setup UI in the cluster. Setup UI is a operations UI that will be used for deploy/upgrade of the cluster. The initial page will show the list of RPMs that are currently installed in the cluster. This is also the upload page to upload all the RPMs

Figure 36: RPM Upload



Upload the RPMs in the order that is shown on setup UI. The order is

1. tetration_os_rpminstall_k9
2. tetration_os_UcsFirmware_k9
3. tetration_os_adhoc_k9
4. tetration_os_mother_rpm_k9
5. tetration_os_enforcement_k9
6. tetration_os_base_rpm_k9



Note For Secure Workload Virtual clusters deployed on vSphere, be sure to also upgrade the tetration_os_ova_k9 RPM and do not upload the tetration_os_base_rpm_k9.

Uploading any other order will result in upload failure. Until all the RPMs are uploaded in the correct order Continue button will be disabled.

Logs for each upload can be seen by clicking on the Log symbol on the left of every RPM. Also uploads that failed will be marked RED in color.

Figure 37: RPM Upload Log

Tetration Setup Diagnostics > RPM Upload > Site Config > Site Config Check > Run notTest

RPM Upload

✓	tetration_os_rpminstall_k9	3.5.0.7.dev
✓	tetration_os_UcsFirmware_k9	3.5.0.7.dev
✓	tetration_os_adhoc_k9	3.5.0.7.dev
✓	tetration_os_mother_rpm_k9	3.5.0.7.dev
✓	tetration_os_enforcement_k9	3.5.0.7.dev
✓	tetration_os_base_rpm_k9	3.5.0.7.dev

Select RPM file

Browse... tetration_os_enforcement_k9-3.5.0.8.dev.rpm

Upload Continue Skip

verifying RPM...

RPM downloaded

RPM install failed

Site Information

The next step is to update the site information. Not all site information fields are updateable. Only the following fields can be updated:

- SSH public Key
- Sentinel Alert Email (for Bosun)
- CIMC Internal Network
- CIMC Internal Network Gateway
- External Network



Note Do not change the existing external network, you can add additional networks by appending to the existing ones. Changing or Removing existing network will make the cluster unusable.

- DNS Resolvers
- DNS Domain
- NTP Servers
- SMTP Server
- SMTP Port
- SMTP Username (Optional)
- SMTP Password (Optional)

- Syslog Server (Optional)
- Syslog Port (Optional)
- Syslog Severity (Optional)

**Note**

- The syslog server severity ranges from critical to informational. Severity needs to be set to warning or higher (informational) for bosun alerts.
- From 3.1 version, **External syslog via setup UI is not supported**. Users will have to configure TAN Appliance to export data to syslog. Refer to [External syslog tunneling moving to TAN](#) for more details.
- Secure Workload supports secure SMTP communication with mail servers that support SSL/TLS communication via the STARTTLS command. The standard port for servers that support secure traffic is usually 587/TCP, but many servers also accept secure communication on the standard 25/TCP port.

Secure Workload does not support the SMTPS protocol for communicating with external mail servers.

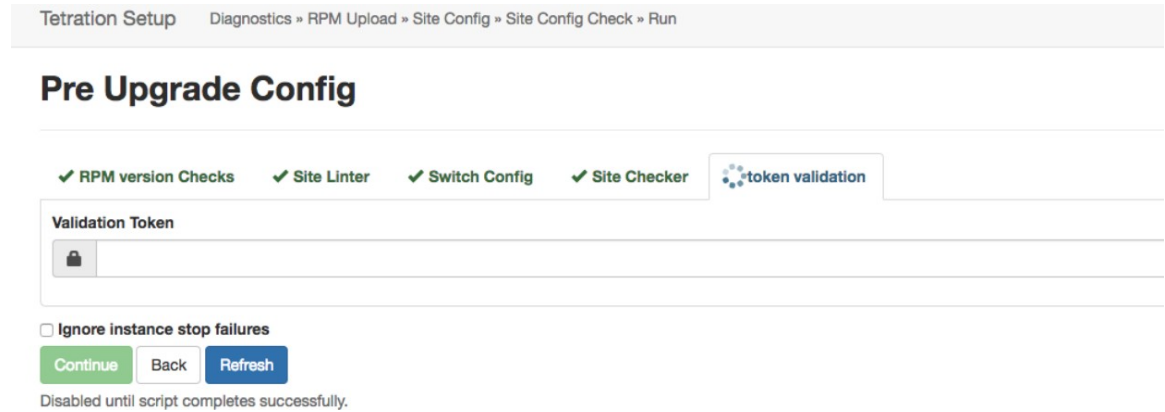
Rest of the fields are not updatable. If there are no changes, click on Continue to trigger the Pre-Upgrade Checks, else update the fields and then click on Continue.

Pre Upgrade Checks

Before we start upgrades we do few checks on the cluster and ensure things are in order before we start upgrading:

1. RPM version checks - checks to ensure all the RPMs are uploaded and the version is correct. It doesn't check if the order was correct, just checks if it was uploaded. Note Order checks are done as a part of upload itself.
2. Site Linter - Does Site Info Linting
3. Switch Config - Configures the Leafs/Spine switches
4. Site Checker - Does DNS, NTP and SMTP server checks. Sends an email at the end with a token, the email is sent to the primary site admin account. If any of the services - DNS, NTP or SMTP is not usable, this step will fail.
5. Token Validation - Enter the token sent in the email and hit Continue.

Figure 38: Pre Upgrade Checks



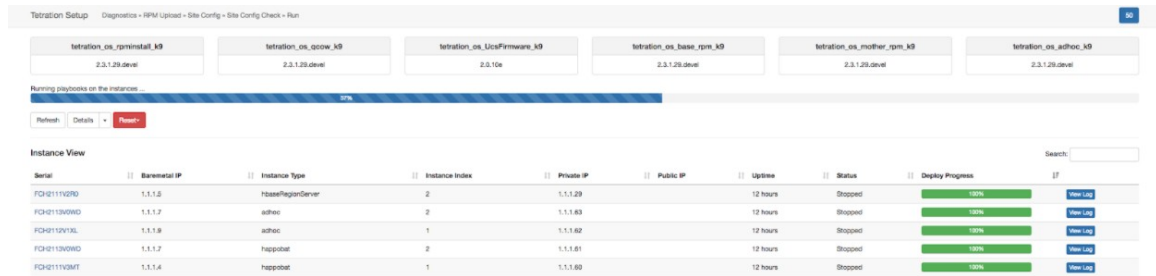
Upgrading the Cluster

After you complete the pre-upgrade step, and enter the token received in the 'verify token email', click **Continue** to start the upgrade. Avoid checking the option 'Ignore Stop Failures'. This is a recovery option for upgrade failures when certain services do not shut down. Using this option will shut down the VMs that can create failures when the services come up.



Note Use this option under supervision.

Figure 39: Upgrading the Cluster



On clicking on “Continue” - Upgrade will start.

Procedure

- Step 1** On the top right clicking on the cluster name will show the site information used.
- Step 2** Below that will have all tetration_os RPMs and their versions.
- Step 3** The global upgrade bar will show the upgrade progress. It will be blue in color while things are in progress, green when done and red when it fails. Right above the progress bar will show the current status of upgrade.
- Step 4** There are 3 buttons:
 - a) Refresh - Refreshes the page.

- b) Details - Click **Details**, this shows all the steps that have completed during this upgrade. Click the arrow next to it to show the logs that can be opened.
- c) Reset - This will have an option to Reset Orchestrator State. This Option will cancel the upgrade and take you back to the start. DO NOT use this unless the upgrade had failed and few minutes have passed after upgrade had failed to let all the processes reach completion before restarting upgrade.
- d) Restart - When an upgrade fails, click **Restart** to restart the cluster and begin a fresh upgrade. This can help resolve any pending cleanup operations or issues that may be blocking the upgrade processes.

Step 5

On the instance view, every individual VMs deploy status is tracked. The columns include:

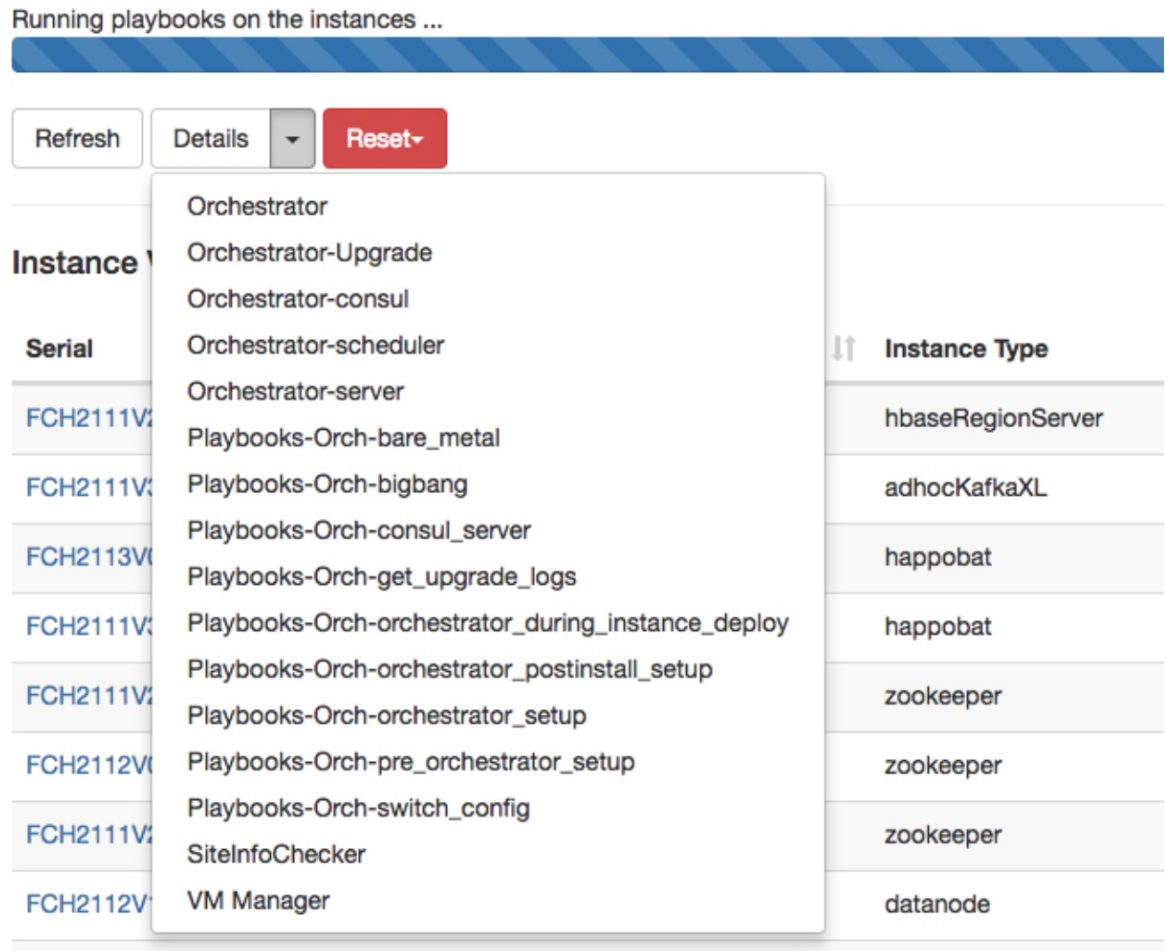
- a) Serial - Baremetal Serial that hosts this VM
 - b) Baremetal IP - the Internal IP assigned to this Baremetal
 - c) Instance Type - the type of VM
 - d) Instance Index - Index of the VM - there are multiple VMs of the same type for high-availability.
 - e) Private IP - the Internal IP assigned to this VM
 - f) Public IP - the routable IP assigned to this VM - not all VMs have this.
 - g) Uptime - Uptime of the VM
 - h) Status - Can be Stopped, Deployed, Failed, Not Started or In Progress.
 - i) Deploy Progress - Deploy Percentage
 - j) View Log - button to view the deploy status of the VM
-

Logs

There are two type of logs:

1. VM deployment logs - these logs can be seen by clicking on “View Log” button.
2. Orchestration Logs. These can be seen by clicking on the arrow next to the details button. It will show up:

Figure 40: Logs



Each of the links will point to the logs.

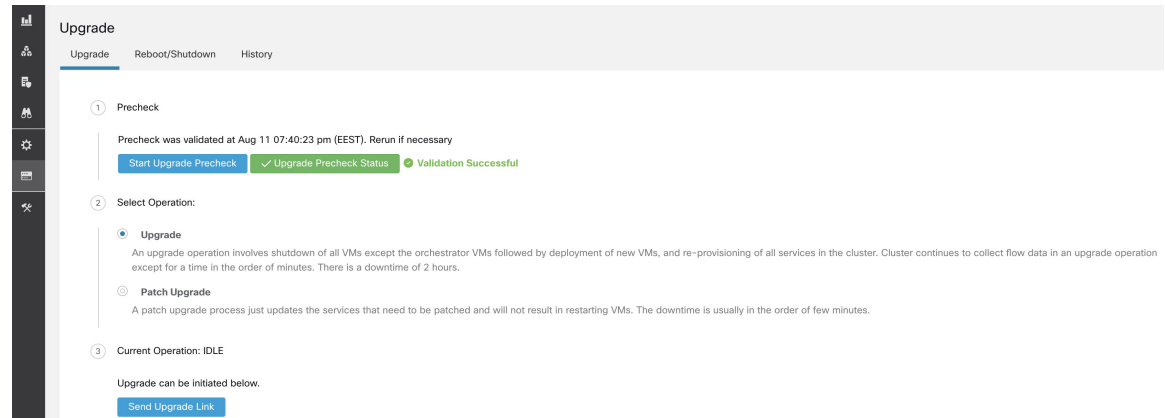
Procedure

- Step 1** Orchestrator - Orchestrator log - this is the first place to track progress. Any failures will point to another log to look at.
- Step 2** Orchestrator-Upgrade - NOP for 2.3
- Step 3** Orchestrator-consul - consul logs that runs on primary orchestrator
- Step 4** Orchestrator-Scheduler - VM scheduler logs - which VM got placed on which baremetal and the scheduling log.
- Step 5** Orchestrator-server - HTTP server logs from orchestrator
- Step 6** Playbooks-* - all the playbook logs that run on orchestrator.

Running Pre-Upgrade Checks any time

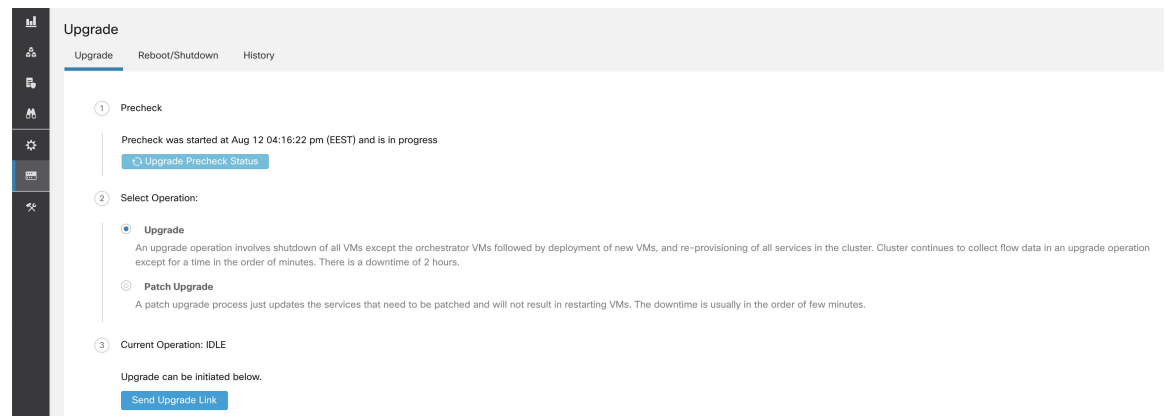
Occasionally, after scheduling an upgrade and while initiating an upgrade, there might be a hardware failure or cluster is not ready to be upgraded. This might require to be fixed before proceeding with upgrades. Instead of waiting until an upgrade window, Pre-Upgrade checks can be initiated any time. These checks can be run any number of times and any time except when an upgrade/patch/reboot is initiated. To run Pre-Upgrade Checks any time, go to the Upgrade Page.

Figure 41: Running Pre-Upgrade Checks any time steps



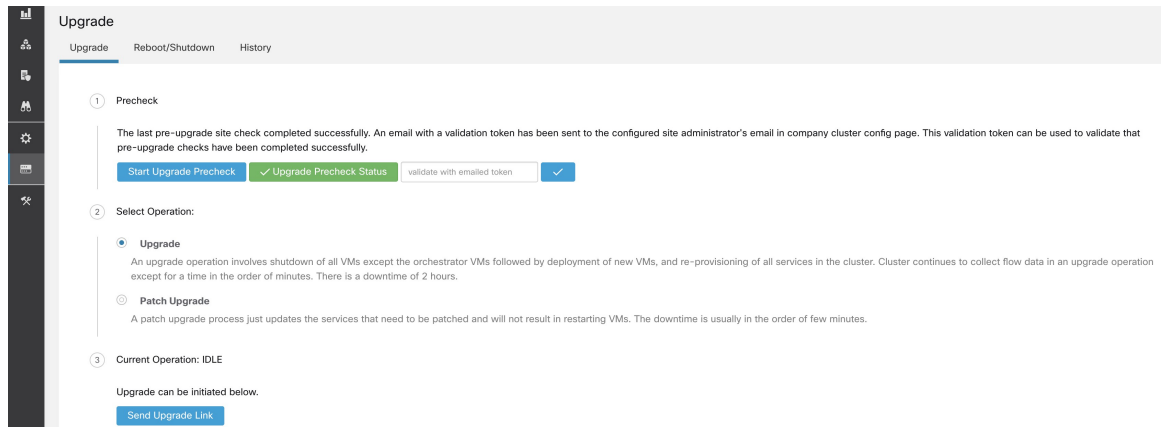
Click on the Start Upgrade Precheck. This will initiate the pre-upgrade checks and will transition to running state:

Figure 42: Running Pre-Upgrade Checks any time steps



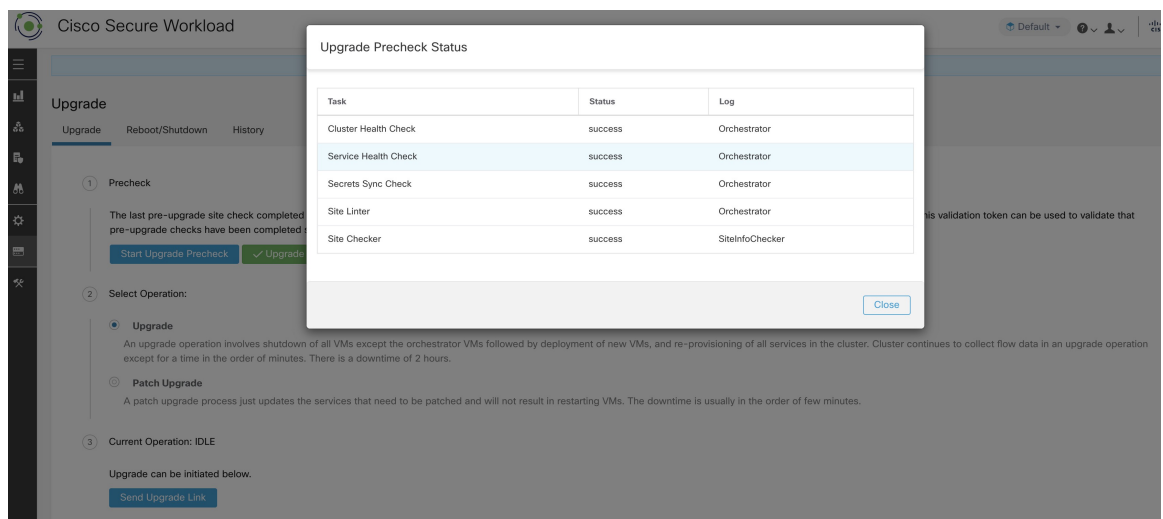
During this time orchestrator runs all the pre-upgrade checks. Once all the checks pass, an email will be sent to the user who initiated the check with an email token. Enter the token to complete the pre-upgrade checks.

Figure 43: Running Pre-Upgrade Checks any time steps



If there are any failures during pre-upgrade checks it will transition to failed state and will show which task failed. Any time the status can be checked and will show up in a new dialog box.

Figure 44: Running Pre-Upgrade Checks any time steps



Data Backup and Restore (DBR)

If **DBR** is enabled on the cluster, also see [Upgrade with Data Backup and Restore](#).

Snapshots

Accessing the Snapshot Creation User Interface

Users with **Customer Support** role can access the snapshot tool by selecting **Troubleshoot > Snapshots** from the navigation bar at the left side of the window.

The Snapshot tool can be used to create a Classic Snapshot or a Cisco Integrated Management Controller (CIMC) technical support bundles. Clicking on the Create Snapshot button on the Snapshot file list page loads a page to choose a Classic Snapshot or a CIMC Snapshot (technical support bundle). The option to choose a CIMC Snapshot is disabled on Secure Workload Software Only (ESXi) and Secure Workload SaaS.

Clicking on the Classic Snapshot button loads the Snapshot tool runner user interface:

Figure 45: Snapshot tool runner

Clicking on the CIMC Snapshot button loads the CIMC Technical Support tool runner user interface:

Figure 46: CIMC Technical Support runner

Creating a Snapshot

Selecting Create Snapshot with the default options, the Snapshot tool collects:

- Logs
- State of Hadoop/YARN application and logs
- Alert history
- Numerous TSDB statistics

It is possible to override the defaults and specify certain options.

- logs options
 - max log days - number of days of logs to collect, default 2.

- max log size - maximum number of bytes per log to collect, default 128kb.
- hosts - hosts to get logs/status from, default all.
- logfiles - regex of logs to be fetched, default all.
- yarn options
 - yarn app state - application states (RUNNING, FAILED, KILLED, UNASSIGNED, etc) to get information for, default all.
- alerts options
 - alert days - the number of days worth of alert data to collect.
- tsdb options
 - tsdb days - the number of days worth of tsdb data to collect, increasing this can create very large Snapshots.
- fulltsdb options
 - fulltsdb - a JSON object that can be used to specify startTime, endTime fullDumpPath, localDumpFile and nameFilterIncludeRegex to limit which metrics are collected.
- comments - can be added to describe why or who is collecting the snapshot.

After selecting Create Snapshot, a progress bar for the snapshot is displayed at the top of the Snapshot file list page. When the snapshot completes, it can be downloaded using the Download button on the Snapshots file list page. Only one snapshot can be collected at a time.

Creating a CIMC Technical Support Bundle

On the CIMC Snapshot (technical support bundle) page, select the serial number of the node the CIMC Technical Support Bundle should be created for and click the Create Snapshot button. A progress bar for the CIMC Technical Support Bundle collection will appear in the Snapshot file list page and the comments section will reflect that the CIMC Technical Support Bundle collection has been triggered. Once the CIMC Technical Support Bundle collection is complete, the file can be downloaded from the Snapshot file list page.

Using a Snapshot

Untarring a snapshot creates a `./clustername_snapshot` directory that contains the logs for each machine. The logs are saved as text files that contain the data from several directories from the machines. The Snapshot also saves all the Hadoop/TSDB data that was captured in JSON format.

Figure 47: Using a Snapshot

```
~/Downloads/tet-snapshot $ ls -lhrGg
total 93840
drwxr-xr-x@ 42 staff 1.4K Mar 30 15:24 zookeeper-3
drwxr-xr-x@ 42 staff 1.4K Mar 30 15:24 zookeeper-2
drwxr-xr-x@ 42 staff 1.4K Mar 30 15:24 zookeeper-1
drwxr-xr-x@ 1691 staff 56K Mar 30 15:23 yarn
drwxr-xr-x@ 42 staff 1.4K Mar 30 15:24 tsdbBosunGrafana-3
drwxr-xr-x@ 42 staff 1.4K Mar 30 15:24 tsdbBosunGrafana-2
drwxr-xr-x@ 42 staff 1.4K Mar 30 15:24 tsdbBosunGrafana-1
-rw-r--r--@ 1 staff 45M Mar 30 15:22 tsdb.json
-rw-r--r--@ 1 staff 4.8K Mar 30 15:19 tet_snapshot_manifest.json
-rw-r--r--@ 1 staff 34K Mar 30 15:24 snapshot_report.log
drwxr-xr-x@ 42 staff 1.4K Mar 30 15:24 secondaryNamenode-1
drwxr-xr-x@ 42 staff 1.4K Mar 30 15:24 resourceManager-2
drwxr-xr-x@ 42 staff 1.4K Mar 30 15:24 resourceManager-1
drwxr-xr-x@ 42 staff 1.4K Mar 30 15:23 redis-3
drwxr-xr-x@ 42 staff 1.4K Mar 30 15:23 redis-2
drwxr-xr-x@ 42 staff 1.4K Mar 30 15:23 redis-1
drwxr-xr-x@ 41 staff 1.4K Mar 30 15:21 orchestrator-3
drwxr-xr-x@ 41 staff 1.4K Mar 30 15:21 orchestrator-2
drwxr-xr-x@ 41 staff 1.4K Mar 30 15:21 orchestrator-1
drwxr-xr-x@ 42 staff 1.4K Mar 30 15:24 nodemanager-9
drwxr-xr-x@ 42 staff 1.4K Mar 30 15:24 nodemanager-8
drwxr-xr-x@ 42 staff 1.4K Mar 30 15:24 nodemanager-7
drwxr-xr-x@ 42 staff 1.4K Mar 30 15:24 nodemanager-6
drwxr-xr-x@ 42 staff 1.4K Mar 30 15:24 nodemanager-5
drwxr-xr-x@ 42 staff 1.4K Mar 30 15:24 nodemanager-4
drwxr-xr-x@ 42 staff 1.4K Mar 30 15:24 nodemanager-3
drwxr-xr-x@ 42 staff 1.4K Mar 30 15:24 nodemanager-2
drwxr-xr-x@ 42 staff 1.4K Mar 30 15:24 nodemanager-10
drwxr-xr-x@ 42 staff 1.4K Mar 30 15:24 nodemanager-1
drwxr-xr-x@ 42 staff 1.4K Mar 30 15:24 namenode-1
drwxr-xr-x@ 42 staff 1.4K Mar 30 15:23 mongodbArbiter-1
drwxr-xr-x@ 42 staff 1.4K Mar 30 15:23 mongodb-2
drwxr-xr-x@ 42 staff 1.4K Mar 30 15:23 mongodb-1
```

When opening the packaged index.html in a browser, there are tabs for:


- Terse list of alert state changes.

Figure 48: Terse list of alert state changes

Alerts	Dashboard	Hadoop	Logs
Fri Oct 23 2015 16:29:51 GMT-0700 (PDT): adm.checkMissingAdmNightlyMetric: 1			
Fri Oct 23 2015 16:29:51 GMT-0700 (PDT): sys.diskUsageIsMoreThan90Percent: 1			
Fri Oct 23 2015 16:29:51 GMT-0700 (PDT): druid.checkMissingMetrics: 1			
Fri Oct 23 2015 16:29:51 GMT-0700 (PDT): pipeline.flowsWithNoEPGIsHigh: 1			
Fri Oct 23 2015 16:29:51 GMT-0700 (PDT): adm.checkMissingMachineInfoMetric: 1			
Fri Oct 23 2015 16:35:51 GMT-0700 (PDT): druid.checkMissingMetrics: 0			
Fri Oct 23 2015 16:44:51 GMT-0700 (PDT): druid.checkMissingMetrics: 1			
Fri Oct 23 2015 16:49:51 GMT-0700 (PDT): druid.checkMissingMetrics: 0			
Fri Oct 23 2015 16:59:51 GMT-0700 (PDT): druid.checkMissingMetrics: 1			
Fri Oct 23 2015 17:04:51 GMT-0700 (PDT): druid.checkMissingMetrics: 0			
Fri Oct 23 2015 17:14:51 GMT-0700 (PDT): druid.checkMissingMetrics: 1			
Fri Oct 23 2015 17:24:52 GMT-0700 (PDT): pipeline.BDPipelineRuntimeSecsIsOverThreshold: 1			
Fri Oct 23 2015 17:49:52 GMT-0700 (PDT): pipeline.BDPipelineRuntimeSecsIsOverThreshold: 0			
Fri Oct 23 2015 18:49:37 GMT-0700 (PDT): druid.checkMissingMetrics: 0			
Fri Oct 23 2015 18:59:37 GMT-0700 (PDT): druid.checkMissingMetrics: 1			
Fri Oct 23 2015 19:04:52 GMT-0700 (PDT): druid.checkMissingMetrics: 0			
Fri Oct 23 2015 19:29:37 GMT-0700 (PDT): druid.checkMissingMetrics: 1			
Fri Oct 23 2015 19:34:52 GMT-0700 (PDT): druid.checkMissingMetrics: 0			

- Reproduction of grafana dashboards.

Figure 49: Reproduction of grafana dashboards

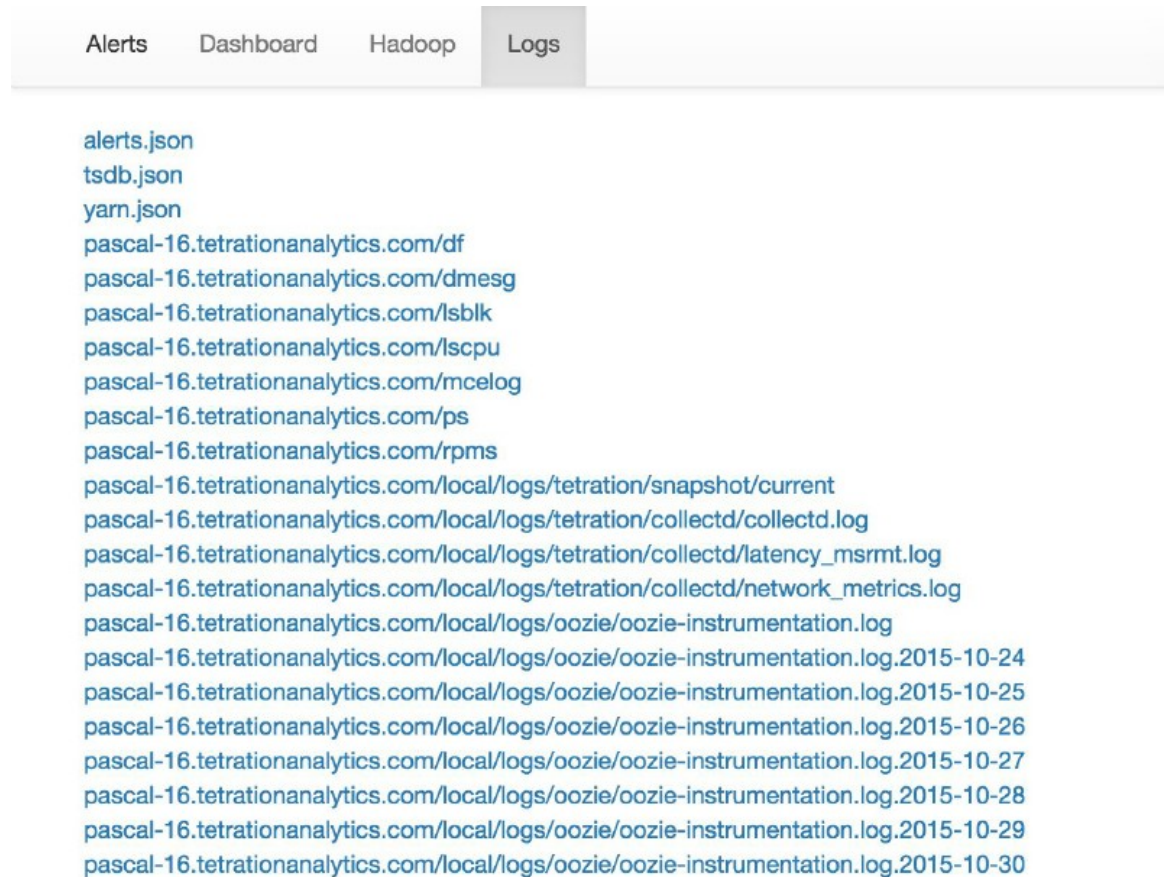


- Reproduction of the Hadoop Resource Manager front end that contains jobs and their state. Selecting a job displays the logs for the job.

Figure 50: Reproduction of the Hadoop Resource Manager

Alerts Dashboard Hadoop Logs					
RUNNING FAILED All jobs					
state	id	name		applicationType	elapsedTime
RUNNING	application_1442528378995_192995	com.tetration.pipeline.PipelineMain		SPARK	948440504
RUNNING	application_1442528378995_107366	com.tetration.pipeline.ActiveFlow		SPARK	2419532064
RUNNING	application_1442528378995_107368	com.tetration.pipeline.UberBidirCopier		SPARK	2419507170
RUNNING	application_1442528378995_107367	com.tetration.retention.RetentionMain		SPARK	2419512413
RUNNING	application_1442528378995_107369	com.tetration.pipeline.UberMachineInfoCopier		SPARK	2420352532
RUNNING	application_1442528378995_256357	attacks-index-generator-Optional.of([2015-11-02T23:21:00.000Z/2015-11-02T23:22:00.000Z])		MAPREDUCE	10483
RUNNING	application_1442528378995_256356	aggregated_flows-index-generator-Optional.of([2015-11-02T23:21:00.000Z/2015-11-02T23:22:00.000Z])		MAPREDUCE	10178
RUNNING	application_1442528378995_256355	hosts-index-generator-Optional.of([2015-11-02T23:22:00.000Z/2015-11-02T23:23:00.000Z])		MAPREDUCE	10513
RUNNING	application_1442528378995_256348	aggregated_flows-index-generator-Optional.of([2015-11-02T23:19:00.000Z/2015-11-02T23:20:00.000Z])		MAPREDUCE	115046
RUNNING	application_1442528378995_256354	sensor_stats-index-generator-Optional.of([2015-11-02T23:22:00.000Z/2015-11-02T23:23:00.000Z])		MAPREDUCE	10721
RUNNING	application_1442528378995_256351	aggregated_flows-index-generator-Optional.of([2015-11-02T23:20:00.000Z/2015-11-02T23:21:00.000Z])		MAPREDUCE	60209
RUNNING	application_1442528378995_256344	aggregated_flows-index-generator-Optional.of([2015-11-02T23:18:00.000Z/2015-11-02T23:19:00.000Z])		MAPREDUCE	164729
FINISHED	application_1442528378995_253998	attacks-index-generator-Optional.of([2015-11-02T13:32:00.000Z/2015-11-02T13:33:00.000Z])		MAPREDUCE	47868
FINISHED	application_1442528378995_253997	sensor_stats-index-generator-Optional.of([2015-11-02T13:33:00.000Z/2015-11-02T13:34:00.000Z])		MAPREDUCE	24514

- List of all logs collected.

Figure 51: List of all logs collected.

Using the Snapshot Service for Debugging and Maintenance

The snapshot service can be used to run service commands, but it requires Customer Support privileges.


Using the Explore tool (**Troubleshoot > Maintenance Explorer**), you can hit arbitrary URIs within the cluster:

Figure 52: Using the Snapshot Service for Debugging and Maintenance Example

The Explore tool only appears for users with Customer Support privileges.

The snapshot service runs on port 15151 of every node. It listens only on the internal network (not exposed externally) and has POST endpoints for various commands.

Figure 53: Using the Snapshot Service for Debugging and Maintenance Example



The URI you must hit is **POST** `http://<hostname>:15151/<cmd>?args=<args>`, where args are space separated and URI encoded. It does **not** run your command with a shell. This would avoid allowing anything to be run.

Endpoints of a snapshot are defined for:

- **snapshot 0.2.5**

- ls

- svstatus, svrestart - runs **sv status, sv restart** Example: `1.1.11.15:15151/svrestart?args=snapshot`

- hadoopls runs `hadoop fs -ls <args>`

- hadoopdu - runs `hadoop fs -du <args>`

- ps Example: `1.1.11.31:15151/ps?args=eafux`

- du

- ambari - runs `ambari_service.py`

- monit

- MegaCli64 (`/usr/bin/MegaCli64`)

- service

- hadoopfsck - runs `hadoop -fsck`

- **snapshot 0.2.6**

- makecurrent - runs `make -C /local/deploy-ansible current`

- netstat

- **snapshot 0.2.7 (run as uid “nobody”)**

```

-cat
-head
-tail
-grep
-ip -6 neighbor
-ip address
-ip neighbor

```

There is another endpoint, POST /runsinged, which will run shell scripts signed by Secure Workload. It runs `gpg -d` on the POSTed data. If it can be verified against a signature, it will run the encrypted text under a shell. This means importing a public key on each server as part of the ansible setup and the need to keep the private key secure.

Run Book

Users with Customer Support privileges can use Run Book by selecting **Troubleshoot > Maintenance Explorer** from the navigation bar at the left side of the window. Select **POST** from the drop-down menu. (Otherwise you will receive Page Not Found errors when running commands.)

Using the snapshot REST endpoint to restart services:

- **druid: 1.1.11.17:15151/service?args=supervisord%20restart**

-druid hosts are all IPs .17 through .24; .17, .18 are coordinators, .19 is the indexer, and .20-.24 are brokers

- **hadoop pipeline launchers:**

-1.1.11.25:15151/svrestart?args=activeflowpipeline

-1.1.11.25:15151/svrestart?args=adm

-1.1.11.25:15151/svrestart?args=batchmover_bidir

-1.1.11.25:15151/svrestart?args=batchmover_machineinfo

-1.1.11.25:15151/svrestart?args=BDPipeline

-1.1.11.25:15151/svrestart?args=mongo_indexer

-1.1.11.25:15151/svrestart?args=retentionPipeline

- **policy engine**

-1.1.11.25:15151/svrestart?args=policy_server

- **wss**

-1.1.11.47:15151/svrestart?args=wss

Explore/Snapshot Endpoints Overview

To run any endpoint, you will need to go to the **Troubleshoot > Maintenance Explorer** page from the navigation bar at the left side of the window.

You can also view each endpoint overview in the explore page by running a **POST** command on any host as `<end-point>?usage=true`.

For example: `makecurrent?usage=true`

GET commands

Endpoint	Description
<code>bm_details</code>	<ul style="list-style-type: none"> Displays the baremetals information
<code>endpoints</code>	<ul style="list-style-type: none"> Lists all the endpoints on the host
<code>members</code>	<ul style="list-style-type: none"> Displays the current list of consul members, along with their status
<code>port2cime</code>	<ul style="list-style-type: none"> Lists the IPs that the port is connected to Should be run on the orchestrator hosts only
<code>status</code>	<ul style="list-style-type: none"> Displays the status of the snapshot service on the host
<code>vm_info</code>	<ul style="list-style-type: none"> Displays the VM information of the location Should be run on the Baremetal hosts only Run endpoint as <code>vm_info?args=<vmname></code>

POST commands

Endpoint	Description
bm_shutdown_or_reboot	<ul style="list-style-type: none"> Gracefully shutdown or reboot a baremetal host by first shutting down all the virtual machines on that host then issuing a shutdown or reboot command to the bare metal. You can also get the shutdown or reboot status using this endpoint. To get the shutdown or reboot status of a node use: <code>bm_shutdown_or_reboot? query=serial=FCH2308V0FH</code> To start a graceful bare metal shutdown use: <code>bm_shutdown_or_reboot? method=POST</code> and set the body to a JSON object that describes the host serial number. For example: <code>{"serial": "FCH2308V0FH"}</code> To start a graceful bare metal reboot use: <code>bm_shutdown_or_reboot? method=POST</code> and set the body to a JSON object that describes the host serial number and include a reboot key set to 'true'. For example: <code>{"serial": "FCH2308V0FH", "reboot": true}</code>
cat	<ul style="list-style-type: none"> wrapper command for unix 'cat' command
cimc_password_random	<ul style="list-style-type: none"> Randomizes the CIMC password. Should be run on the orchestrator hosts only
cleancmdlogs	<ul style="list-style-type: none"> Clears the logs in <code>/local/logs/tetration/snapshot/cmdlogs/snapshot_cleancmdlogs_log</code>
clear_sel	<ul style="list-style-type: none"> Clears the system event logs Should be run on the Baremetal hosts only

Endpoint	Description
cluster_fw_upgrade	<ul style="list-style-type: none"> • This is a BETA feature. • Run a UCS firmware upgrade across the whole cluster. • After this completes successfully each bare metal will need to be rebooted to activate the BIOS and other component firmware. • Run as: cluster_fw_upgrade • This endpoint will kick off and monitor the firmware upgrade and update the log file when a stage of the upgrade has been started or completed. • To get the status of the upgrade, use the cluster_fw_upgrade_status end- point.
cluster_fw_upgrade_status	<ul style="list-style-type: none"> • This is a BETA feature. • Get the status of the full cluster UCS firmware upgrade. • Run as cluster_fw_upgrade_status
cluster_powerdown	<ul style="list-style-type: none"> • Powers down the cluster • USE WITH CAUTION, BRINGS THE CLUSTER DOWN • Run endpoint as cluster_powerdown?args==start
collector_status	<ul style="list-style-type: none"> • Displays the status of the collector • Should be run on the collector hosts only
consul_kv_export	<ul style="list-style-type: none"> • Displays k-v pairs from consul in JSON format • Should be run on the orchestrator hosts only
consul_kv_recurse	<ul style="list-style-type: none"> • Displays k-v pairs from consul in tabular format • Should be run on the orchestrator hosts only
df	<ul style="list-style-type: none"> • wrapper command for unix ‘df’ command
dig	<ul style="list-style-type: none"> • wrapper command for unix ‘dig’ command
dmesg	<ul style="list-style-type: none"> • wrapper command for unix ‘dmesg’ command

Endpoint	Description
dmidecode	<ul style="list-style-type: none"> • wrapper command for unix ‘dmidecode’ command
druid_coordinator_v1	<ul style="list-style-type: none"> • Displays the druid stats.
du	<ul style="list-style-type: none"> • wrapper command for unix ‘du’ command
dusorted	<ul style="list-style-type: none"> • wrapper command for unix ‘dusorted’ command
externalize_change_tunnel	<ul style="list-style-type: none"> • Changes the collector IP that will be used to tunnel the CIMC UI • Runas: externalize_change_tunnel?method=POST • Pass {“collector_ip” : “<IP>”} in the Body • Should be run on the orchestrator hosts only
externalize_mgmt	<ul style="list-style-type: none"> • Displays the current status of externalizing the CIMC UI’s for each server • Displays the address and time remaining for externalization • Should be run on the orchestrator hosts only
externalize_mgmt_read_only_password	<ul style="list-style-type: none"> • Changes the read only password (ta_guest) for both the switch and CIMC UI • Changes only when they are externalized • Run as: externalize_mgmt_read_only_password?method=POST • Pass {“password” : “<password>”} in the Body • Should be run on the orchestrator hosts only
fsock	<ul style="list-style-type: none"> • wrapper command for unix ‘fsock’ command • Should be run on Baremetal host only

Endpoint	Description
get_cimc_techsupport	<ul style="list-style-type: none"> • INPUT Internal IP address of BM. • Retrieves the CIMC techsupport. • Once it is completed it will be available for down- load from the snapshots page in the UI. • This can be run from any host on the cluster and requires the baremetal internal ip address as an argument. • Example: get_cimc_techsupport?args=1.1.0.9
syslog_endpoints	<ul style="list-style-type: none"> • Controls the syslog configurations for 1 or more of the ucs servers. • Run the command with -h to get full list of pa- rameters
grep	<ul style="list-style-type: none"> • wrapper command for unix ‘grep’ command
hadoopbalancer	<ul style="list-style-type: none"> • Distributes HDFS data uniformly across all nodes • Should be run on hosts that have hdfs for exam- ple launcherhost
hadoopdu	<ul style="list-style-type: none"> • Prints the directory utilization of hdfs • Should be run on hosts that have hdfs for exam- ple launcherhost
hadoopfsck	<ul style="list-style-type: none"> • Runs hadoop fsck and reports the state of the pro- vided hdfs file system • It also takes “-delete “as an argument to clear cor- rupt or missing blocks • Before deleting make sure all the DataNodes are up else we might lose data • Should be run on the launcher hosts only • To report state run as: hadoopfsck?args=/raw • To delete corrupt files run as: hadoopfsck?args=/raw -delete
hadoopls	<ul style="list-style-type: none"> • Lists the Hadoop File System • Should be run on hosts that have hdfs for exam- ple launcherhost

Endpoint	Description
hbasehbk	<ul style="list-style-type: none"> • Checks for consistency and table integrity problems and repairing a corrupted HBase • Should be run on the HBase hosts only • To identify inconsistency, run as: hbasehbk?args=-details • To repair a corrupted HBase, run as: hbasehbk?args=-repair • Output written to: /local/etcd/raft/snapshot/mdlog/snapshot_hbasebk_log.txt • Repair with caution
hdfs_safe_state_recover	<ul style="list-style-type: none"> • Removes HDFS from safe state • Required if HDFS is in READ_ONLY_STATE due full capacity and space has been cleared • Should be run on the launcher hosts only • Run as:hadoopfs-rm'{{ hdfs_safe_state_marker_location }}/HDFS_READ_ONLY'
initctl	<ul style="list-style-type: none"> • wrapper command for unix 'initctl' command
head	<ul style="list-style-type: none"> • wrapper command for unix 'head' command
internal_haproxy_status	<ul style="list-style-type: none"> • Prints the internal haproxy status and stats • Should be run on the orchestrator hosts only
ip	<ul style="list-style-type: none"> • wrapper command for unix 'ip' command
ipmifru	<ul style="list-style-type: none"> • Prints Field Replaceable Unit (FRU) Information • Should be run on the Baremetal hosts only
ipmilan	<ul style="list-style-type: none"> • Prints the LAN configuration • Should be run on the Baremetal hosts only
ipmisel	<ul style="list-style-type: none"> • Prints System Event Log (SEL) entries • Should be run on the Baremetal hosts only

Endpoint	Description
ipmisensorlist	<ul style="list-style-type: none"> Prints the IPMI sensor information Should be run on the Baremetal hosts only
jstack	<ul style="list-style-type: none"> Prints Java stack traces of Java threads for a given Java process or core file
ls	<ul style="list-style-type: none"> wrapper command for unix 'ls' command
lshw	<ul style="list-style-type: none"> wrapper command for unix 'lshw' command
lsuf	<ul style="list-style-type: none"> wrapper command for unix 'lsuf' command
lvdisplay	<ul style="list-style-type: none"> wrapper command for unix 'lvdisplay' command
lvs	<ul style="list-style-type: none"> wrapper command for unix 'lvs' command
lvscan	<ul style="list-style-type: none"> wrapper command for unix 'lvscan' command
makecurrent	<ul style="list-style-type: none"> Resets/fastforwards the pipeline processing the marker to the current timestamps Should be run on the orchestrator nodes only Run endpoint as makecurrent?args=start
mongo_rs_status	<ul style="list-style-type: none"> Displays the mongo replication status Should be run on either the mongodb or the enforcementpolicystore hosts
mongo_stats	<ul style="list-style-type: none"> Displays the mongo stats Should be run on either the mongodb or the enforcementpolicystore hosts
mongodump	<ul style="list-style-type: none"> Dumps the collections from the database Should be run on either the mongodb or the enforcementpolicystore hosts Run as: mongodump?args=<collection>[-db DB]
monit	<ul style="list-style-type: none"> wrapper command for unix 'monit' command
namenode_jmx	<ul style="list-style-type: none"> Displays the primary namenode jmx metrics
ndisc6	<ul style="list-style-type: none"> wrapper command for unix 'ndisc6' command

Endpoint	Description
netstat	<ul style="list-style-type: none"> • wrapper command for unix 'netstat' command
ntpq	<ul style="list-style-type: none"> • wrapper command for unix 'ntpq' command
orch_reset	<ul style="list-style-type: none"> • Resets orchestrator state to IDLE • Run after commissioning or decommissioning failure • Should be run on the orchestrator.service.consul host only • Do not use this command without consulting customer support
orch_stop	<ul style="list-style-type: none"> • Stops the orchestrator primary and trigger a switchover • Should be run on the orchestrator.service.consul host only • USE WITH CAUTION
ping	<ul style="list-style-type: none"> • wrapper command for unix 'ping' command
ping6	<ul style="list-style-type: none"> • wrapper command for unix 'ping6' command
ps	<ul style="list-style-type: none"> • wrapper command for unix 'ps' command
pv	<ul style="list-style-type: none"> • wrapper command for unix 'pv' command
pvs	<ul style="list-style-type: none"> • wrapper command for unix 'pvs' command
pvdisplay	<ul style="list-style-type: none"> • wrapper command for unix 'pvdisplay' command
rdisc6	<ul style="list-style-type: none"> • wrapper command for unix 'rdisc6' command
rebootnode	<ul style="list-style-type: none"> • Reboots the node • Should be run on the Baremetal hosts only
recover_rpmdb	<ul style="list-style-type: none"> • Recovers a corrupt RPMDB on a node • Can be run on Baremetals or VMs

Endpoint	Description
recoverhbase	<ul style="list-style-type: none"> Recovers Hbase and TSDB Service Should be run on orchestrator hosts only Should be run when HDFS is Healthy
recovervm	<ul style="list-style-type: none"> Try to recover VM via stop/fsck/start Should be run on orchestrator hosts only Run endpoint as recovervm?args=<vmname>
restartservices	<ul style="list-style-type: none"> Stops and starts all non UI services Should be run on the orchestrator.service.consul host only USE WITH CAUTION Run endpoint as restartservices?args=-start
runsigned	<ul style="list-style-type: none"> Runs the signed script provided by cisco Follow the steps provided in the script guidelines
service	<ul style="list-style-type: none"> wrapper command for unix 'service' command
smartctl	<ul style="list-style-type: none"> Run the smartctl executable Should only be run on a bare metal node
storcli	<ul style="list-style-type: none"> wrapper command for unix 'storcli' command
sudocat	<ul style="list-style-type: none"> wrapper for 'cat' command that works only under /var/log or /local/logs
sudogrep	<ul style="list-style-type: none"> wrapper for 'grep' command that works only under /var/log or /local/logs
sudohead	<ul style="list-style-type: none"> wrapper for 'head' command that works only under /var/log or /local/logs
sudols	<ul style="list-style-type: none"> wrapper for 'ls' command that works only under /var/log or /local/logs

Endpoint	Description
sudo tail	<ul style="list-style-type: none"> • wrapper for ‘tail’ command that works only under /var/log or /local/logs
sudo zgrep	<ul style="list-style-type: none"> • wrapper for ‘zgrep’ command that works only under /var/log or /local/logs
sudo zcat	<ul style="list-style-type: none"> • wrapper for ‘zcat’ command that works only under /var/log or /local/logs
sv restart	<ul style="list-style-type: none"> • Restarts the service mentioned, run command assvrestart?args=<servicename>
sv status	<ul style="list-style-type: none"> • Prints the status of the service mentioned, run as svstatus?args=<servicename>
switch info	<ul style="list-style-type: none"> • Get the information about the cluster switches
switch_namenode	<ul style="list-style-type: none"> • Manually fail over namenode from primary or secondary • Should be run on the orchestrator.service.consul host only • Run while recommission or decommission of namenode hosts • Run endpoint as switch_namenode?args=-start
switch_secondarynamenode	<ul style="list-style-type: none"> • Manually fail over secondarynamenode from secondary to primary • Should be run on the orchestrator.service.consul host only • Run while recommission or decommission of namenode hosts • Run endpoint as switch_secondarynamenode?args=-start

Endpoint	Description
switch_yarn	<ul style="list-style-type: none"> Manually fail over resource manager from primary or secondary or vice versa Should be run on the orchestrator.service.consul host only Run while recommission or decommission of resource manager hosts Run endpoint as switch_yarn?args=-start
tail	<ul style="list-style-type: none"> wrapper command for unix 'tail' command
toggle_chassis_locator	<ul style="list-style-type: none"> Toggle a chassis locator on a physical bare metal specified by the node serial number. Run from any node as: toggle_chassis_locator?method=POST Set the body to a JSON object that describes the host serial number (only one serial number is supported at a time), for example: {"serials": ["FCH2308V0FH"]}
tnp_agent_logs	<ul style="list-style-type: none"> Create a snapshot with all log files provided by Load Balancer agents registered as External Orchestrators Should be run on the launcherhost hosts
tnp_datastream	<ul style="list-style-type: none"> Create a snapshot with policy stream data consumed by Load Balancer policy enforcement agents registered as External Orchestrators Should be run on the orchestrator hosts In order to download policy status stream data run endpoint as tnp_datastream?args=-ds_type datasink
ui_haproxy_status	<ul style="list-style-type: none"> Prints the haproxy stats and status for external haproxy
uptime	<ul style="list-style-type: none"> wrapper command for unix 'uptime' command
userapps_kill	<ul style="list-style-type: none"> Kills all the running user application Should be run on the launcherhost hosts only
vgdisplay	<ul style="list-style-type: none"> wrapper command for unix 'vgdisplay' command

Endpoint	Description
vgs	<ul style="list-style-type: none"> • wrapper command for unix 'vgs' command
vmfs	<ul style="list-style-type: none"> • Lists the file system on a VM • Should be run on the Baremetal hosts only • Run endpoint as vmfs?args=<vmname>
vminfo	<ul style="list-style-type: none"> • Prints the VM information • Should be run on the Baremetal hosts only • Run endpoint as vminfo?args=<vmname>
vmlist	<ul style="list-style-type: none"> • Lists of all the VM on a baremetal • Should be run on the Baremetal hosts only • Run endpoint as vmlist?args=<vmname>
vmreboot	<ul style="list-style-type: none"> • Reboots the VM • Should be run on the Baremetal hosts only • Run endpoint as vmreboot?args=<vmname>
vmshutdown	<ul style="list-style-type: none"> • Gracefully shutdown the VM • Should be run on the Baremetal hosts only • Run endpoint as vmshutdown?args=<vmname>
vmstart	<ul style="list-style-type: none"> • Starts the VM • Should be run on the Baremetal hosts only • Run endpoint as vmstart?args=<vmname>
vmstop	<ul style="list-style-type: none"> • Force shutdown the VM • Should be run on the Baremetal hosts only • Run endpoint as vmstop?args=<vmname>
yarnkill	<ul style="list-style-type: none"> • Kills a running Yarn application • Should be run on the launcherhost hosts only • Run endpoint as yarnkill?args=<application id> • To kill all the applications run as yarnkill?args=ALL

Endpoint	Description
yarnlogs	<ul style="list-style-type: none"> Dumps the last 500 mb of yarn application logs Should be run on the launcherhost hosts only Run endpoint as yarnlogs?args=<application id> <job user>
zcat	<ul style="list-style-type: none"> wrapper command for unix 'zcat' command
zgrep	<ul style="list-style-type: none"> wrapper command for unix 'zgrep' command

Server Maintenance

Server maintenance involves replacement of any faulty server component like Hard Disk, Memory or replacement of the entire server itself.



Note If there are multiple servers on the cluster that need maintenance then do server maintenance on them one at a time. Decommissioning multiple servers at the same time can lead to loss of data.

The **Cluster Status** page (accessed from the **Troubleshoot** menu in the left navigation bar) is used to perform all the steps involved in server maintenance. It can be accessed by all users but the actions can be carried out by **Customer Support** users only. It shows the status of all the physical servers in Cisco Secure Workload rack.

Figure 54: Server Maintenance

Model: 8RU-PROD

[CIMC/TOR guest password](#) [Change external access](#)

Orchestrator State: IDLE


Displaying 6 nodes (0 selected)

<input type="checkbox"/>	State ↑↓	Status ↑↓	Switch Port ↑	Serial ↑↓	Uptime ↑↓	
<input type="checkbox"/>	Commissioned	Active	Ethernet1/1	FCH2206V1NF	2mo 27d 18h 25m 47s	Select action <input checked="" type="checkbox"/> + Commission <input checked="" type="checkbox"/> Decommission <input checked="" type="checkbox"/> Reimage <input checked="" type="checkbox"/> Firmware upgrade <input checked="" type="checkbox"/> Power off <input checked="" type="checkbox"/> Reboot
<input type="checkbox"/>	Commissioned	Active	Ethernet1/2	FCH2206V1ZF	2mo 27d 18h 24m 52s	
<div style="border: 1px solid #ccc; padding: 5px;"> <p>Serial: FCH2206V1ZF</p> <p>Private IP: 1.1.1.4 CIMC IP: 10.13.4.12 Status: Active State: Commissioned SW Version: 3.6.0.10_devel Hardware: 44 cores, 96GB memory, 8 disks, 17.57T space, SSD Firmware: View Firmware Upgrade Logs</p> <ul style="list-style-type: none"> CIMC: 2.0(10e) BIOS: 2.0.10e.0 Cisco 12G SAS Modular RAID Controller Slot HBA: 24.12.1-0205 UCS VIC 1225 10Gbps 2 port CNA SFP+ Slot 1: 4.1(3a) Intel(R) I350 1 Gbps Network Controller Slot L: 0x80000E74-1.810.8 UCS VIC 1225 10Gbps 2 port CNA SFP+ Slot 2: 4.1(3a) <p>Instances</p> <ul style="list-style-type: none"> collectorDatamover-6 datanode-6 druidHistoricalBroker-4 enforcementCoordinator-3 orchestrator-2 redis-1 secondaryNamenode-1 <p>Disks Status</p> <ul style="list-style-type: none"> 252:1 HEALTHY 252:2 HEALTHY 252:3 HEALTHY 252:4 HEALTHY 252:5 HEALTHY 252:6 HEALTHY 252:7 HEALTHY 252:8 HEALTHY </div>						
<input type="checkbox"/>	Commissioned	Active	Ethernet1/3	FCH2206V1N1	2mo 27d 18h 25m 35s	+ ↓
<input type="checkbox"/>	Commissioned	Active	Ethernet1/4	FCH2133V2LN	2mo 27d 18h 26m 52s	+ ↓

Steps involved in server or component replacement

Figure 55: Server Maintenance steps

Server State Transition Diagram



- Determine the server that requires maintenance** : This can be done using the server *Serial* number or the *Switchport* the server is connected to , from the *Cluster Status* page. Note the CIMC IP of the server to be replaced. it would be shown in the server box on the *Cluster Status* page
- Check for actions for special VMs** : From the server box find out the VMs or instances present on the server and check if any special actions need to be carried out for those VMs. The next section lists out Actions for VMs during server maintenance.
- Decommission the server** : Once any pre-decommission actions are performed, use the **Cluster Status** page to decommission the server. Even if the server has failed and appears *Inactive* on the page , we still have to perform all the server maintenance steps. Decommission steps can be performed even if the server is powered off

Figure 56: Server Maintenance steps

Displaying 7 nodes (3 non-Active) (0 selected) Select action

<input type="checkbox"/>	State	Status	Switch Port	Serial	Uptime
<input type="checkbox"/>	Commissioned	Active	Ethernet1/1	FCH2036V224	15d 5h 8m
<input type="checkbox"/>	Commissioned	Active	Ethernet1/2	FCH2036V10Z	15d 5h 8m 33s
<input type="checkbox"/>	New	Active	Ethernet1/3	FCH2033V31K	15d 5h 8m 28s
<input type="checkbox"/>	Decommissioned	Shutdown in progress	Ethernet1/4	FCH2038V0Y5	15d 5h 8m 32s

Serial: FCH2038V0Y5 Switch Port: Ethernet1/4

Private IP: 1.1.1.4
 CIMC IP: 10.16.238.14
 Status: Shutdown in progress
 State: Decommissioned
 SW Version: 3.0.3.31225.deepai.tet.mrpm.build [△](#)
 Hardware: 44 cores, 1T memory, 8 disks, 19.32T space, SSD
 Firmware: [View Firmware Upgrade Logs](#)

- CIMC: 2.0(10e)
- Cisco 12G SAS Modular Raid Controller: 24.9.1-0018
- UCS VIC 1225 10Gbps 2 port CNA SFP+: 4.1(1g) [△](#)
- Intel(R) I350 1 Gbps Network Controller: 0x80000B15-1.808.2
- BIOS: C220M4.2.0.10e.0.0620162104 [△](#)

Shutdown Status:

Shutdown Errors:

4. **Perform server maintenance** : After the node is marked *Decommissioned* on the **Cluster Status** page perform any post decommission special actions for the VMs. Any component or server replacement can be carried out now. If the entire server is replaced, then change the CIMC IP of the new server to be same as that of the replaced server. The CIMC IP for each server is available on the **Cluster Status** page
5. **Reimage after component replacement** : Reimage the server after the component replacement using the **Cluster Status** page. Reimage takes about 30 mins and requires cimc access to servers. The Server is marked *NEW* after reimage is completed.
6. **Replacing entire server** : If the entire server is replaced, then the server would appear in *NEW* state on the **Cluster Status** page. The s/w version for the server can be seen on the same page. If the s/w version is different from the s/w version of the cluster then reimage the server.

Figure 57: Server Maintenance steps

Displaying 7 nodes (3 non-Active) (0 selected) Select action

<input type="checkbox"/>	State	Status	Switch Port	Serial	Uptime
<input type="checkbox"/>	Commissioned	Active	Ethernet1/1	FCH2036V224	15d 5h 8m
<input type="checkbox"/>	Commissioned	Active	Ethernet1/2	FCH2036V10Z	15d 5h 8m 33s
<input type="checkbox"/>	New	Active	Ethernet1/3	FCH2033V31K	15d 5h 8m 28s

Serial: FCH2033V31K Switch Port: Ethernet1/3

Private IP: 1.1.1.5
 CIMC IP: 10.16.238.13
 Status: Active
 State: New
 SW Version: 3.0.3.31225.deepai.tet.mrpm.build [△](#)
 Hardware: 44 cores, 1T memory, 8 disks, 19.32T space, SSD
 Firmware: [View Firmware Upgrade Logs](#)

Instances

- collectorDatamover-3
- datanode-1
- druidHistoricalBroker-1
- enforcementCoordinator-1
- enforcementPolicyStore-3
- happobot-2
- hbaseRegionServer-2
- orchestrator-3
- resourceManager-2
- zookeeper-1

7. **Commission the server** : After the server is marked *NEW* we can kick off the commissioning of the node from the **Cluster Status** page. This step will provision the VMs on the server. Commissioning of a server takes about 45 mins. The server will be marked *Commissioned* after commissioning completes.

Figure 58: Server Maintenance steps

Displaying 6 nodes (0 selected) Select action

<input type="checkbox"/>	State	Status	Switch Port	Serial	Uptime
<input type="checkbox"/>	Commissioned	Active	Ethernet1/1	FCH2110V1ZY	1d:15h:27m:39s
<input type="checkbox"/>	Commissioned	Active	Ethernet1/2	FCH2048V2WZ	4h:15m:41s
<input type="checkbox"/>	Initialized	Active	Ethernet1/3	FCH2048V2VY	10m:40s
<div style="border: 1px solid #ccc; padding: 5px;"> <p>Serial: FCH2048V2VY Switch Port: Ethernet1/3</p> <p>Private IP: 1.1.1.4</p> <p>CIMC IP: 172.26.230.178</p> <p>Status: Active</p> <p>State: Initialized</p> <p>SW Version: 2.3.1.24.devel</p> <p>Hardware: 44 cores, 1T memory, 8 disks, 19.32T space, SSD</p> <p>Firmware: View Firmware Upgrade Logs</p> <p>Instances</p> <ul style="list-style-type: none"> • collectorDatamover-3 • datanode-1 • druidHistoricalBroker-1 • enforcementCoordinator-1 • enforcementPolicyStore-3 • hbaseRegionServer-2 • orchestrator-3 • resourceManager-2 • zookeeper-1 </div>					
<input type="checkbox"/>	Commissioned	Active	Ethernet1/4	FCH2049V00C	1d:15h:27m:45s
<input type="checkbox"/>	Commissioned	Active	Ethernet1/5	FCH2048V2W0	1d:15h:28m:46s
<input type="checkbox"/>	Commissioned	Active	Ethernet1/6	FCH2049V008	1d:15h:28m:31s

Actions for VMs during server maintenance

Some of the VMs require special actions during the server maintenance procedure. These actions could be pre-decommission, post-decommission or post-commission.

1. **Orchestrator primary** : This is a pre-decommission action. If the server undergoing maintenance has primary orchestrator on it, then POST `orch_stop` command to `orchestrator.service.consul` from explore page before doing decommission. This will switch the primary orchestrator.


Figure 59: Server Maintenance steps

Maintenance Explorer

POST

If you try to decommission a server with primary orchestrator, you will see the following error


Figure 60: Server Maintenance steps



To determine the orchestrator primary run the explore command "primaryorchestrator" on any host.

2. **Namenode** : If the server undergoing maintenance has namenode VM on it, then POST `switch_namenode` on `orchestrator.service.consul` from explore page after decommission and then POST `switch_namenode` on `orchestrator.service.consul` after commission. This is both post-decommission and post-commission action.
3. **Secondary namenode** : If the server undergoing maintenance has secondarynamenode VM on it, then POST `switch_secondarynamenode` on `orchestrator.service.consul` from explore page after decommission and then POST `switch_secondarynamenode` on `orchestrator.service.consul` after commission. This is both post-decommission and post-commission action.
4. **Resource manager primary** : If the server undergoing maintenance has resourcemanager primary on it, then POST `switch_yarn` on `orchestrator.service.consul` from explore page. This is both post-decommission and post-commission action.
5. **Datanode** : The cluster tolerates only one Datanode failure at a time. If multiple servers having Datanode VMs need servicing, then do server maintenance on them one at a time. After each server maintenance wait for the chart under Monitoring | hawkeye | hdfs-monitoring | Block Sanity Info, Missing blocks and Under replicated counts to be 0.

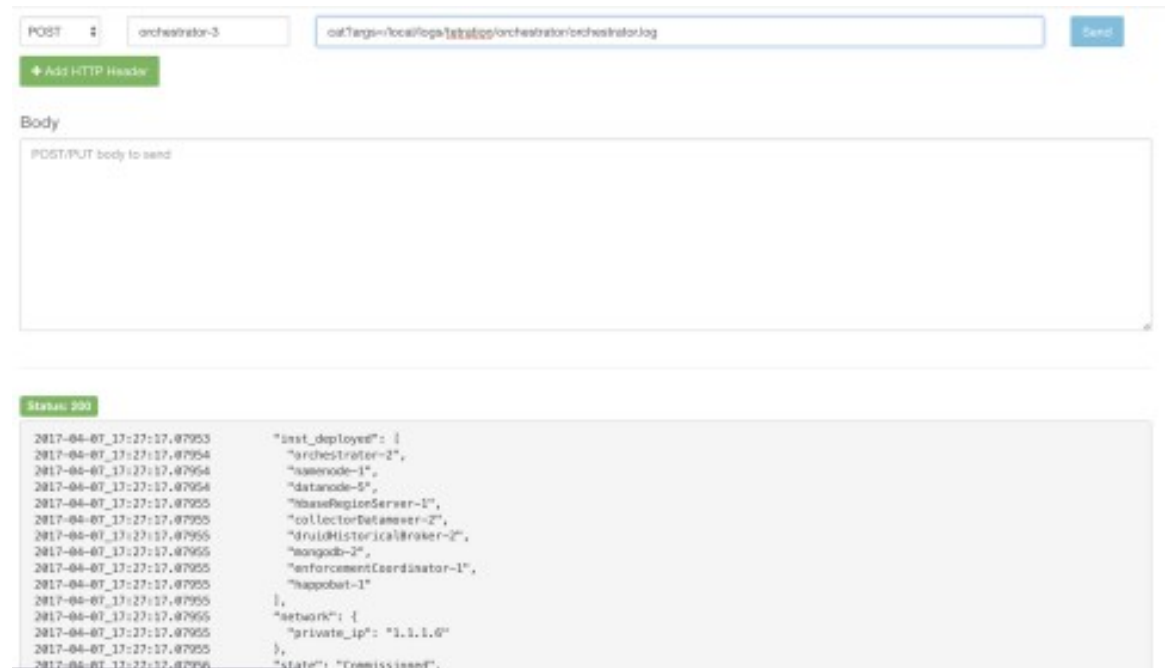
Figure 61: Server Maintenance steps



Troubleshooting server maintenance

1. **Logs** : All the server maintenance logs are part of the orchestrator log. The location is `/local/logs/tetration/orchestrator/orchestrator.log` on `orchestrator.service.consul`.

Figure 62: Server Maintenance log



2. Decommission :

- a. This step deletes the VMs/instances on the server.
- b. It then deletes the entry of these instances in backend consul tables.
- c. This step takes about 5 mins.
- d. The server will be marked *Decommissioned* once the step completes.



Note Decommissioned does not mean the server is powered off. Decommissioning only deletes the Secure Workload content on the server.

- e. If the server is powered off it will be marked **Inactive**. We can still run Decommission on this server from the cluster status page. But the VMs deletion step will not run since the server is powered off. Make sure this server does not join back the cluster in decommissioned state. It needs to be reimaged and added back to the cluster.

3. Reimage :

- a. This step installs the Secure Workload base OS or Hypervisor OS on the server.
- b. It also formats the hard drives and installs few Secure Workload libraries on the server.
- c. Reimage runs a script called **mjlnir** to initiate the server imaging. mjlnir run takes about 5 mins after which the actual imaging begins. Imaging takes about 30 mins. The logs during imaging can be seen only on the console of the server being reimaged. The user can use `ta_dev` key to check for additional info regarding the reimage, like `/var/log/nginx` logs during pxe boot up, `/var/log/messages` to check for dhcp ip and pxe boot configs.

- d. Reimage requires CIMC connectivity from the orchestrator. The easiest way to check for cimc connectivity is to use explore page and POST `ping?args=<cimc ip>` from `orchestrator.service.consul`. **Remember** to change the CIMC IP incase the server is replaced and set the cimc password to the default password
- e. Also cimc network should have been set in site info when the cluster is deployed so that the switches get configured with the correct routes. In case the cluster cimc connectivity is not set correctly you will see the following result in the orchestrator logs.

4. Commissioning:

- a. Commissioning schedules the VMs on the server and runs playbooks in the VMs to install Secure Workload software
- b. it takes about 45 mins for commissioning to finish.
- c. The workflow is similar to deploy or upgrade.
- d. The Logs will indicate any failures during commissioning
- e. The server on the cluster status page will be marked initialized during commissioning and marked commissioned only after the step completes

Baremetal Exclude (bmexclude)

If a hardware failure is detected upon restart of a cluster after power shutdown, currently the cluster gets stuck in a state where we can neither run Reboot workflow to get services stable nor run Commission workflow as down services result in commissioning failure. This feature is expected to help in such scenarios by allowing user to reboot (upgrade) with a bad hardware, after which regular RMA process for the failed baremetal can be performed.

User is expected to use a post to explore endpoint with serial of the baremetal to be excluded.

1. Action: POST
2. Host: `orchestrator.service.consul`
3. Endpoint: `exclude_bms?method=POST`
4. Body: `{“baremetal”: [“BMSERIAL”]}`

Orchestrator performs few checks to determine if the exclusion is feasible. In which case, it will setup few consul keys and return success message indicating which baremetal and VMs will be excluded in the next reboot/upgrade workflow. If the baremetals include certain vms, they can't be excluded as described in the Limitation section below, the explore endpoint will reply back with the message indicating why the exclusion is not possible. After successful post on the explore endpoint, user can initiate reboot/upgrade through main UI and proceed with reboot as usual. At the end of the upgrade, we remove the exclude bm list. If there is a need to run upgrade/reboot again with exclude BMs, users are expected to post to the `bmexclude explore` endpoint again.

Limitations We don't allow following VMs to be excluded currently. 1. namenode 2. secondaryNamenode 3. mon- godb 4. mongodbArbiter

Disk Maintenance

Disk Maintenance involves replacement of any faulty hard disks from one or more servers. Orchestrator monitors the health of the disks as reported by bmmgr on every server in the cluster. If there are any faulty disks, it is indicated on a banner on the **Cluster Status** page (From the navigation pane, choose **Troubleshoot**).

The banner displays the number of disks that are in UNHEALTHY state. Click *here* on the banner, which takes you to the disk replacement wizard. As a user, you can only access the disk replacement page, however, with the help of the wizard, **Customer Support** can perform all the steps required for disk maintenance.

Figure 63: Faulty Disk Banner

The screenshot shows the Cisco Tetrating interface for Cluster Status. At the top, there is a navigation pane with 'Default' and 'Monitoring' tabs. Below the navigation, a banner indicates 'You do not have an active license. The evaluation period will end on Mon Aug 03 2020 05:04:13 GMT+0000. Please notify admin.' The main content area shows 'Model: 8RU-PROD' and 'Orchestrator State: IDLE'. A red banner below this states 'There are 3 unhealthy disks in the appliance. You can replace them. Please check here'. Below the banner, a table displays 6 nodes (0 selected) with columns for State, Status, Switch Port, Serial, Uptime, and CIMC Snapshots.

State	Status	Switch Port	Serial	Uptime	CIMC Snapshots
Commissioned	Active	Ethernet1/1	FCH2148V1EU	16d 11h 22m 40s	
Commissioned	Active	Ethernet1/2	FCH2148V1N9	16d 11h 22m 40s	
Commissioned	Active	Ethernet1/3	FCH2148V1NG	16d 11h 24m 4s	
Commissioned	Active	Ethernet1/4	FCH2148V1EP	16d 11h 20m 15s	
Commissioned	Active	Ethernet1/5	FCH2148V1N2	16d 11h 22m 18s	
Commissioned	Active	Ethernet1/6	FCH2148V1NE	16d 11h 21m 54s	

Disk Replacement Wizard

Before you start the replacement process, keep the replacement disks available.

The Disk Replacement Wizard shows the details of the failed disks. These details include the size, the type, the make and the model for every disk that needs replacement. The page also shows the slot ID and lists all the VMs that use each of these disks.

Figure 64: Cluster Status-Disk Replacement

Cisco Tetration | CLUSTER STATUS - DISK REPLACEMENT

Default Monitoring ?

1 Prerequisites 2 Decommission Drives 3 Replace Drives 4 Commission Drives

Drive Replacement Process

- Decommission all the disks that are in **UNHEALTHY** status.
- Replace all the disks one by one in the physical appliance.
- Commission all the replaced disks together in the final step.

Before you begin

- Keep the **replacement disks** with following configuration in hand.
 - 2 disks of type 1.454 TB SSD INTEL SSDSC2EB016T7K
 - 1 disk of type 3.492 TB SSD SAMSUNG MZ7LM3T8HMLP-00003

Node Serial: **FCH2148V1EP**

Enclosure:Slot	Status	Affected VMs
252:3	UNHEALTHY	druidHistoricalBroker-4

Node Serial: **FCH2148V1N9**

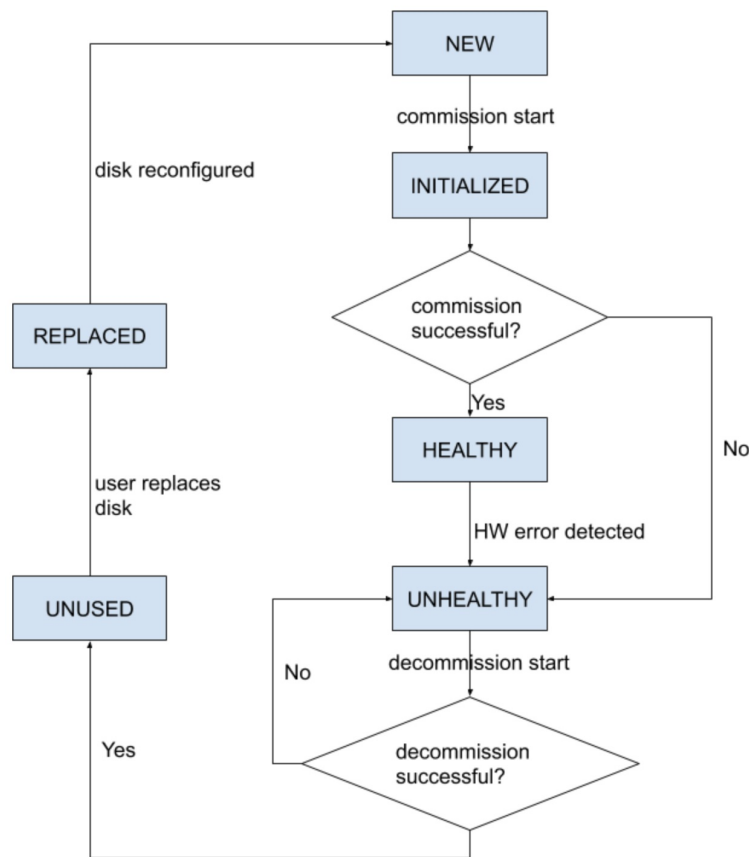
Enclosure:Slot	Status	Affected VMs
252:1	UNHEALTHY	druidCoordinator-1, orchestrator-2, enforcementPolicyStore-1, enforcementCoordinator-3, redis-1, secondaryNamenode-1, datanode-6, collectorDatamover-6, tsdbBosunGrafana-1
252:7	UNHEALTHY	datanode-6

> Proceed to Decommission

Disk Status Transitions

In the cluster, Hard Disks can have six states. **HEALTHY**, **UNHEALTHY**, **UNUSED**, **REPLACED**, **NEW**, and **INITIALIZED**. After you deploy or upgrade the cluster, the status of every disk in the cluster is **HEALTHY**. Based on various error detections, the status of one or more disks can become **UNHEALTHY**.

Figure 65: Disk Status Transitions



The first step of the disk replacement process is decommission where all the VMs that use these disks are removed from the cluster. The status of disks that are decommissioned become UNUSED. After decommission, the replacement disks should be inserted in their appropriate slots. Users confirm that the disks are replaced, which will be the backend's signal to reconfigure the newly added disks. This will change the status to REPLACED and after the next hardware scan these replaced disks' status will change to NEW. This transition can take 2–3 minutes.

Once all the disks have been replaced and reconfigured, you can deploy all the VMs that were removed as part of the decommission process. The start of commission changes the disk status to INITIALIZED. A successful commission makes all disks' status HEALTHY. A failure in this step makes the status UNHEALTHY again so that we start the recovery from decommission again.

Requirement Prechecks

Before any of the decommission or commission step can take place, a requirement precheck must be performed. Backend performs various checks all of which must pass before user can proceed with the decommission or commission step. Any failed checks will be reported on the disk replacement wizard with the failure detail and suggested corrective action, which must be taken before the needed step can proceed.

Example of such pre check are: namenode and secondaryNamenode cannot be decommissioned together. only one datanode can be decommissioned at one time. namenode is healthy before commissioning.

Figure 66: Disk Replacement Prechecks

The screenshot displays the Cisco Tetration interface for 'CLUSTER STATUS - DISK REPLACEMENT'. The 'Decommission Drives' step is active, showing instructions for decommissioning unhealthy drives and a table of selected disks.

Decommissioning Unhealthy Drives

1. Prechecks should be run successfully before decommission. You can re-run these prechecks after addressing any precheck failures.
2. Decommission step is not necessary if there is no disk with **UNHEALTHY** status.
3. In case of decommission failure, you have to run prechecks again before attempting decommission.

Select Disks

Select unhealthy disks for decommission

Selected 2 disks

Serial	Enclosure:Slot	Status	Affected VMs
FCH2148V1EP	252:3	UNHEALTHY	druidHistoricalBroker-4
FCH2148V1N9	252:7	UNHEALTHY	datanode-6

Prechecks

Start Prechecks

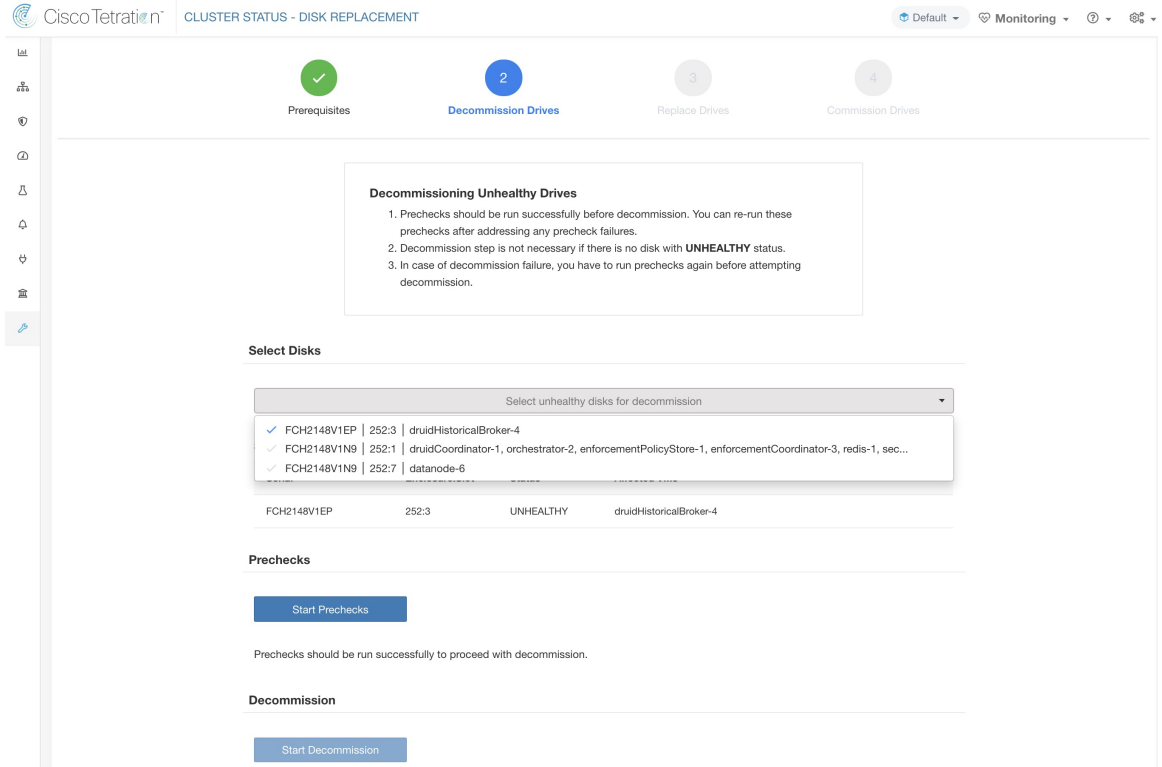
✓ Prechecks were successful at May 5 05:17:05 pm (PDT).

Decommission

Start Decommission

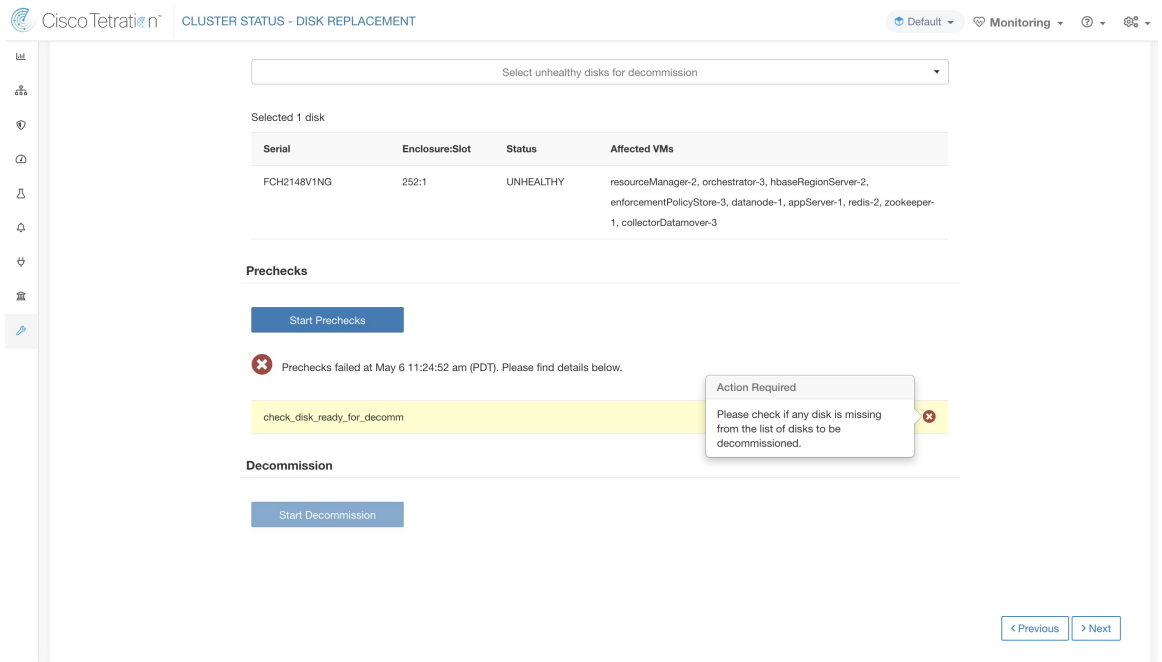
User can select any set of failed disks to be decommissioned together and start the decommission precheck. Changing the set of failed disk requires a rerun of the precheck. Same prechecks are checked again before the task (decommission/commission) starts to ensure that there are no new precheck failure between last precheck run and the start of the decommission task.

Figure 67: Unhealthy Disks for Decommission



Upon any failed precheck, a detailed message can be seen by clicking on the failure message and a suggested action will be shown in a pop-over when pointer hovers over the red cross button.

Figure 68: Suggested Action in Pop-Over for Failed Precheck



Decommission Disk

Once the prechecks pass, user can proceed to decommission disk. The progress of decommission will be shown on the disk replacement wizard. Once progress of decommission reaches 100%, all the decommissioned disk status changes to UNUSED.

Figure 69: Monitoring Disk Decommission Progress

The screenshot displays the Cisco Tetration interface for 'CLUSTER STATUS - DISK REPLACEMENT'. The main content area is titled 'Select Disks' and features a dropdown menu labeled 'Select unhealthy disks for decommission'. Below this, a table lists 'Selected 2 disks' with the following data:

Serial	Enclosure:Slot	Status	Affected VMs
WZP233016TN	134:2	UNHEALTHY	datanode-14
WZP233016TN	134:5	UNHEALTHY	datanode-14

Below the table, there are two sections: 'Prechecks' with a 'Start Prechecks' button, and 'Decommission' with a 'Start Decommission' button. A progress indicator shows 'Decommission is in progress.' with a 2% progress bar. A terminal-style window displays the following JSON output:

```
Running Requirements Check:
Starting Decommission: {'serials': [], 'disks': [{'u'slot': 2, u'serial': u'WZP233016TN', u'enclosure': 134}, {u
```

Navigation buttons for '< Previous' and '> Next' are located at the bottom right of the interface.

Replace Disk

Figure 70: Reconfigure Newly Added Disks

Replace Unused Drives

1. Use **disk locator on/off** to identify the exact location of the disk on physical appliance.
2. Once a disk is physically replaced, notify that it has been replaced using **Replace** button.
3. Proceed to **commission** step after all the disks are notified as replaced

Note

- After decommissioning, status of unhealthy drives changes to **UNUSED**.
- After a disk is notified as replaced, the status of the disk changes to **REPLACED**.
- **Serial numbers, size and model** of all disks are also provided for identification.

Turn Off All Node Locators Turn Off All Disk Locators

Node Serial: FCH2148V1EP Switch Port: Ethernet1/4

Enclosure:Slot	Disk Serial	Model	Status	Locator On/Off	Replaced?
252:3	PHDV745600DW1P6EGN	1.454 TB SSD INTEL SSDSC2BB016T7K	UNUSED	🔦	Replace

Node Serial: FCH2148V1N9 Switch Port: Ethernet1/2

Enclosure:Slot	Disk Serial	Model	Status	Locator On/Off	Replaced?
252:2	PHDV745600J81P6EGN	1.454 TB SSD INTEL SSDSC2BB016T7K	UNUSED	🔦	Replace
252:7	S3LJNX0J400526	3.492 TB SSD SAMSUNG MZ7LM3T8HMLP-00003	UNUSED	🔦	

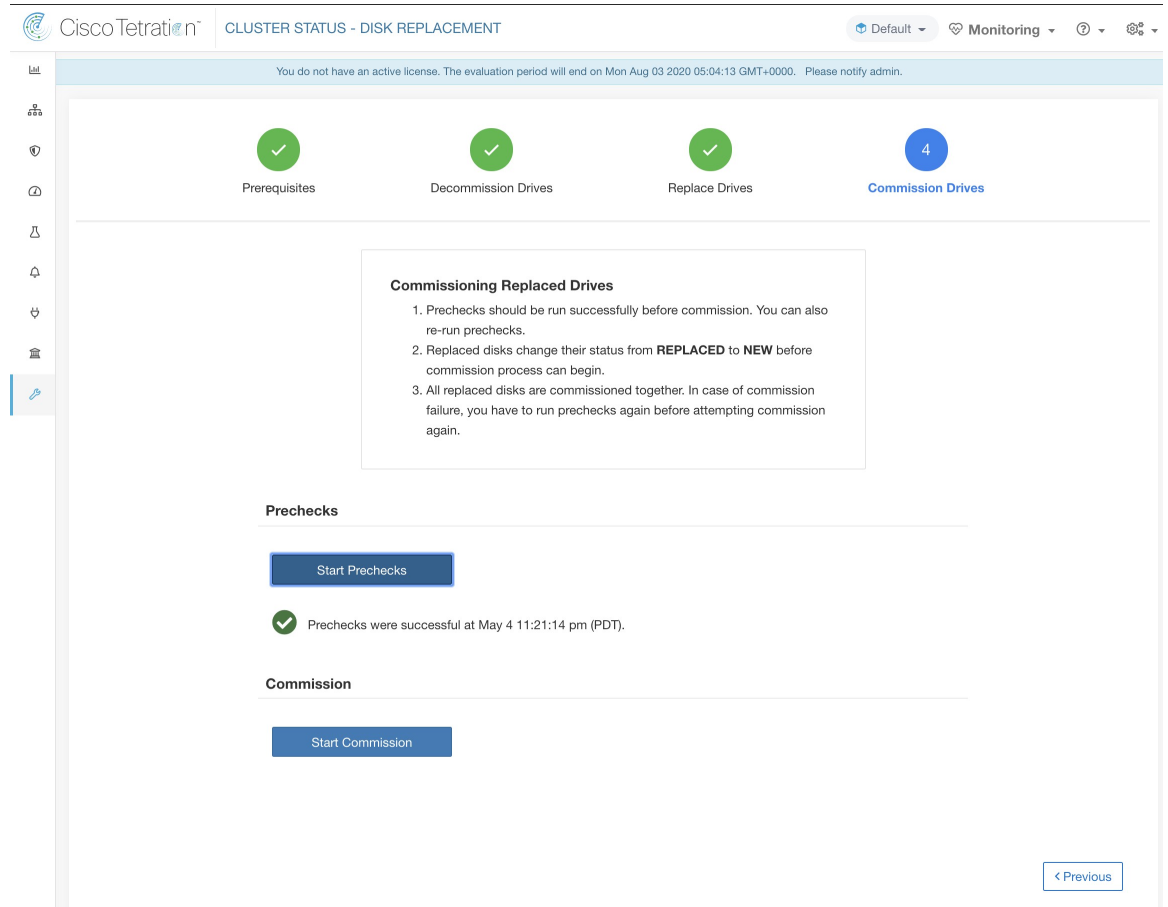
After disk decommission, user is expected to physically replace the disks. To help in this process, we have added disk and server locator LED access on the replace page. There are buttons to switch off all the server and disks locator LEDs to take care any other process that might have left the locators on.

Disks can be physically replaced in any order but they must be reconfigured in smallest to largest slot numbers for a given server. This order is enforced through both UI and the backend. UI has replace button active for disk with the lowest slot number with status UNUSED.

Commission Disk

When all the disks are replaced, we proceed to commission. Like decommission, we must run a set of prechecks before we can continue to commission.

Figure 71: Prechecks Before Commission



Progress of commission is monitored on the disk commission page. At the end successful commission, the status all disks change to HEALTHY.

Figure 72: Commission Progress**Prechecks**[Start Prechecks](#)

Prechecks should be run successfully to proceed with commission.

Commission[Start Commission](#)

Commission is in progress.

82%

```
Starting Commission: {'serials': [], 'disks': [{'u'slot': 3, u'serial': u'FCH2148V1EP', u'enc
All Orchestrator Nodes brought up and Consul Quorum formed
Baremetal IP assignment done. Running pre-deploy playbook
Pre-deploy playbook done.
IDL parsed, Running instance bring up
Stack Manager brought the instances UP
Generating ansible vars, generating ansible tar.gz and setting up to support Service Manager
Running playbooks on the instances
```

[< Previous](#)**Recovery from failure during commission**

A failure after VMs have been redeployed, can be recovered via resume. In such failures, a *Resume Commission* button appears on the disk commission page, which can be clicked to continue commission by restarting the post deploy playbooks.

Figure 73: Resume Commission

Prechecks

Start Prechecks

Prechecks should be run successfully to proceed with commission.

Commission

Start Commission Resume Commission

✘ Last commission attempt has failed.

Failed ORC-1015 Cluster certs playbook failed, check Playbooks-Orch-cluster_certs log - All instances are fully deployed, Running post instance bringup playbooks

```
Running Requirements Check:
Starting Commission: {'serials': [], 'disks': [{u'slot': 3, u'serial': u'FCH2126V0NS', u'enclosure': 252}, {u'slot':
Initial playbook to kick start deploy started
All Orchestrator Nodes brought up and Consul Quorum formed
Baremetal IP assignment done. Running pre-deploy playbook
Pre-deploy playbook done.
IDL parsed, Running instance bring up
Stack Manager brought the instances UP
Generating ansible vars, generating ansible tar.gz and setting up to support Service Manager
Running playbooks on the instances
ORC-1015 Cluster certs playbook failed, check Playbooks-Orch-cluster_certs log - All instances are fully deployed, Rur
```

In case of any failure before the VMs have been redeployed, the disks that were being commissioned will have their status changed to UNHEALTHY. That will require us to restart the replacement process from the decommission of UNHEALTHY disks.

Additional disk failures during commission

In case of any other disks than the ones that are being replaced fails while disk commission is in progress, notice of this failure will be displayed on the disk replacement wizard after the ongoing commission process finishes, either in success or failure.

In cases of resumable failures, user will have two options in what next steps to take.

1. They can try to resume and complete current commission and perform the disk replacement process for the new failures later.
2. Alternatively, they can start decommission of newly failed disk and perform commission of all the disks together.

This second path will be the only path available in cases of non-resumable failures. If the post deploy failure is caused due to the newly failed disks, the second path will again be only way forward, even though we will have resume button available.

Troubleshooting

Logs

1. All the disk commission/decommission logs are part of orchestrator logs. Starting debug point should be /lo- cal/logs/tetration/orchestrator/orchestrator.log on orchestrator.service.consul.

2. Details of any failure during disk replace/reconfigure action can be found on the bmmgr log on the server in consideration. The log location on the server would be `/local/logs/tetration/bmmgr/bmmgr.log`

Limitations

1. Disk containing server's root volumes can't be replaced using this procedure. Such disk failure must be corrected using server maintenance process.
2. Disk commissioning can happen only when all servers are active and in commissioned state. See special handling section below to that describes how to proceed in the cases where a combination of disk and server replacement is needed.

Special handling

Disk and Server Replacement together

In the case of failure scenarios where a disk and a server needs to be commissioned together, user is expected to decommission and replace all the disks that can be decommissioned. Commission of those disk would be prevented by the precheck that ensure that

1. All non healthy disks have the status of NEW
2. All servers are in the *Commissioned* state with status *Active*

Figure 74: Ensure the all servers are commissioned and active before disk commission

Cisco Tetration | CLUSTER STATUS - DISK REPLACEMENT

Default

Prerequisites Decommission Drives Replace Drives Commission Drives

Commissioning Replaced Drives

1. Prechecks should be run successfully before commission. You can also re-run prechecks.
2. Replaced disks change their status from **REPLACED** to **NEW** before commission process can begin.
3. All replaced disks are commissioned together. In case of commission failure, you have to run prechecks again before attempting commission again.

Prechecks

Start Prechecks

Prechecks failed at May 13 06:49:53 pm (PDT). Please find details below.

All Nodes are Commissioned Check

Nodes ['WZP232913LX:(State: New, Status: Active)'] state/status is not (State: Commissioned, Status: Active)

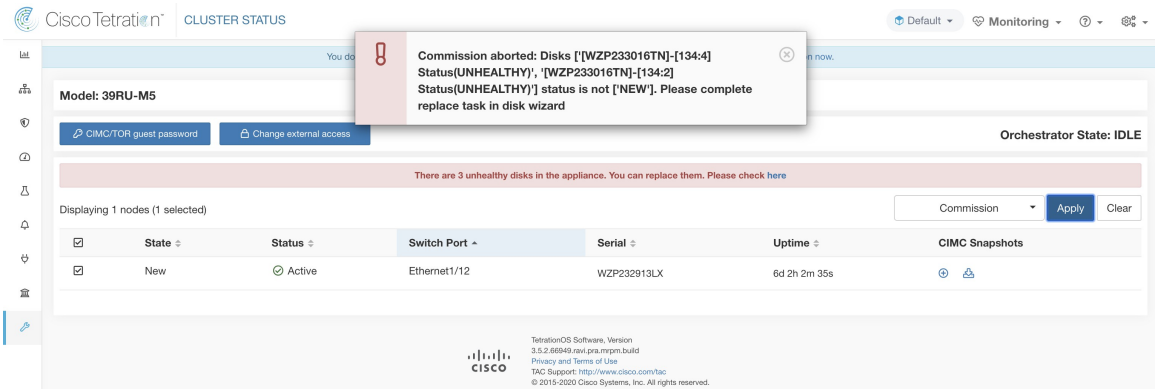
Commission

Start Commission

Once all the UNHEALTHY disks are in the NEW state, the faulty server is expected to be decommission/reimaged/commission back using the server maintenance procedure.

Now server commission will be prevented if there are any disk without status HEALTHY or NEW. A successful server commission will also make the status of all disks HEALTHY.

Figure 75: Ensure the All Faulty Disks Are in NEW State Before Server Commission



Cluster Maintenance Operations

This section describes the maintenance operations that affect the entire cluster.

Shut Down the Secure Workload Cluster

Shutting down the cluster stops all running Secure Workload processes, and powers down all individual nodes. Perform the following steps to shut down the cluster.

Initiate Cluster Shutdown

Procedure

- Step 1** From the navigation pane, choose **Platform > Upgrade/Reboot/Shutdown**.
- Step 2** Click the **Reboot/Shutdown** tab.
- Step 3** Select **Shutdown** and click **Send Shutdown Link**. The shutdown link is delivered to the email address.

Figure 76: Shutdown email

Hello Site Admin!

We received a request that you intend to shutdown the cluster "98". You can do this through the link below.

[Shutdown 98](#) (For best results, please use [Google Chrome](#))

The above link expires by Jul 22 08:34:30 pm (PDT).

If you didn't request this, please ignore this email.

Shutdown will not be triggered until you actually click the above link.

Step 4 On the **Cluster Shutdown** page, click **Shutdown**.

Important You cannot cancel the shutdown after clicking the **Shutdown** button.

Cluster Shutdown Progress

After you initiate the cluster shutdown, the progress of the shutdown and the status are displayed.

Figure 77: Cluster Shutdown Progress

The screenshot shows the 'Cluster Shutdown Progress' page. At the top, there are five buttons for different components: 'tetration_os_rpminstall_k9', 'tetration_os_UcsFirmwar...', 'tetration_os_adhoc_k9', 'tetration_os_mother_rp...', and 'tetration_os_base_rpm_k9', each with a version number '3.3.1.19.devel'. Below these is a progress bar for 'Pre setup for cluster shutdown ...' at 12%. There are 'Refresh' and 'Details' buttons. The 'Instance View' table below shows the following data:

Serial	Baremetal IP	Instance Type	Instance Index	Private IP	Public IP	Uptime	Status	Deploy Progress
FCH2132V1RJ	1.1.1.5	zookeeper	2	1.1.1.23		an hour	Deployed	100%
FCH2133V2J6	1.1.1.8	enforcementPolicyStore	3	1.1.1.48		an hour	Deployed	100%
FCH2133V2J6	1.1.1.8	collectorDatamover	3	1.1.1.36	172.29.154.106	an hour	Deployed	100%
FCH2133V2J6	1.1.1.8	happobat	2	1.1.1.64		an hour	Deployed	100%
FCH2133V1CR	1.1.1.7	appServer	1	1.1.1.10	172.29.154.102	an hour	Deployed	100%

If an error occurs in the initial shutdown prechecks, the progress bar will turn red and click the resume button to restart shutdown after fixing the errors.

After prechecks are completed, VMs are stopped. As the VMs progressively stop, the progress is displayed. The page is similar to the VM stop under upgrades. For more information, see the upgrades section on each field. Stopping all the VMs can take up to 30 minutes.

Figure 78: Stopping VMs

The screenshot shows the 'Cluster Shutdown Progress' page. At the top, there are five buttons for different components: 'tetration_os_rpminstall_k9', 'tetration_os_UcsFirmwar...', 'tetration_os_adhoc_k9', 'tetration_os_mother_rpm...', and 'tetration_os_base_rpm_k9', each with a version number '3.3.1.9.devel'. Below these is a progress bar for 'Stopping all VMs ...' at 15%. There are 'Refresh' and 'Details' buttons. The 'Instance View' table below shows the following data:

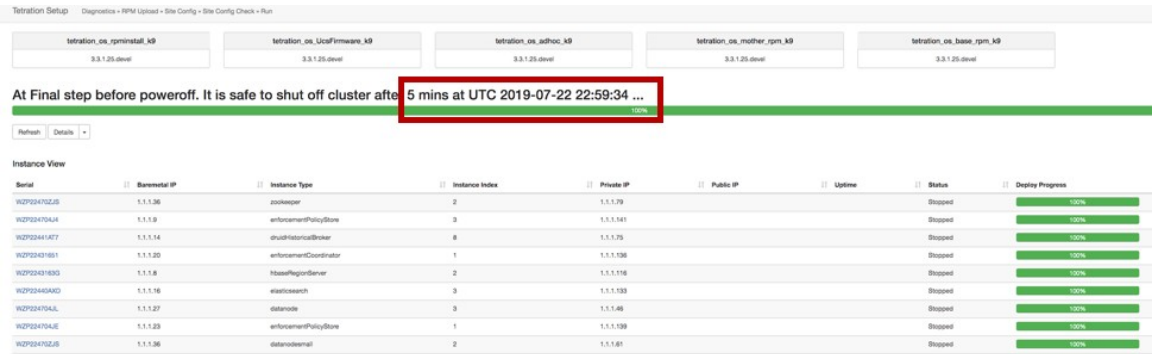
Serial	Baremetal IP	Instance Type	Instance Index	Private IP	Public IP	Uptime	Status	Deploy Progress
FCH2132V1RJ	1.1.1.5	zookeeper	2	1.1.1.23		a day	In Progress	66%
FCH2133V2J6	1.1.1.8	enforcementPolicyStore	3	1.1.1.48		a day	Stopped	100%
FCH2133V2J6	1.1.1.8	collectorDatamover	3	1.1.1.36	172.29.154.106	a day	In Progress	50%
FCH2133V2J6	1.1.1.8	happobat	2	1.1.1.64		a day	Stopped	100%

When the cluster is ready to be shut down, the progress bar will go to a 100% and indicate the time after which it is safe to power off the cluster. See the highlighted in the following screenshot.



Note Do not power off the cluster before waiting for the time displayed on the progress bar.

Figure 79: Shutdown 100 Percent



Reboot the Secure Workload Cluster

To recover the cluster after shutdown, power on the bare metals. When all the individual bare metals are up, the UI becomes accessible. After logging into the cluster, reboot the cluster to render the cluster operational.



Note You *must* reboot the cluster after a shutdown to render it operational.

Initiate Cluster Reboot

Procedure

- Step 1** From the navigation pane, choose **Platform > Upgrade/Reboot/Shutdown**.
- Step 2** Click the **Reboot/Shutdown** tab.
- Step 3** Select **Reboot** and click **Send Reboot Link**.

Click the link that you receive on your email ID to reboot the cluster. On the setup UI page, initiate the cluster reboot. During the reboot, a restricted upgrade operation is performed.

View History of Cluster Maintenance Jobs

To view the previously run cluster maintenance jobs:

1. Navigate to **Platform > Upgrade/Reboot/Shutdown**, and then click the **History** tab.
The cluster operation column lists the cluster tasks such as deploy, upgrade, reboot, or shutdown.
2. To download logs of the cluster jobs, click **Download Logs**.

Data Tap Admin - Data Taps

1. Data Taps
2. Managed Data Taps

Data Taps



Note Cisco Secure Workload Currently supports writing to Kafka Brokers 0.9.x, 0.10.x, 1.0.x and 1.1.x for Datataps

To push any alerts out from Secure Workload cluster, user needs to use a configured data taps. Data Tap Admin users are the only ones who can configure and activate new/existing data taps. Users can only view data taps that belong to their **Tenant**.

Figure 80: Available Data Taps

Data Tap Admin - Data Taps							+ New Data Tap
Name	Topic	Description	Kafka Broker	Type	Status	Actions	
DataTap1	default-datatap1-topic01	The First Data Tap	b4kafka3.tetrationanalytics.com:9092	External	Active		

To manage data taps, click **Manage > Data Tap Admin** in the navigation bar at the left side of the window.

Recommended Kafka Config







While configuring Kafka cluster, Secure Workload recommends to use the ports from 9092, 9093 or 9094 since, these are the ports Secure Workload opens for outgoing traffic for Kafka. The following are the recommended settings for Kafka Brokers:

```
broker.id=<incremental number based on the size of the cluster>
auto.create.topics.enable=true
delete.topic.enable=true
listeners=PLAINTEXT://:9092
port=9092
default.replication.factor=2
host.name=<your_host_name>
advertised.host.name=<your_adversited_hostname>
num.network.threads=12
num.io.threads=12
socket.send.buffer.bytes=102400
socket.receive.buffer.bytes=102400
socket.request.max.bytes=104857600
log.dirs=<directory where logs can be written, ensure that there is sufficient space to
hold the kafka journal logs>
num.partitions=72
num.recovery.threads.per.data.dir=1
log.retention.hours=24
log.segment.bytes=1073741824
log.retention.check.interval.ms=300000
log.cleaner.enable=false
zookeeper.connect=<address of zookeeper ensemble>
zookeeper.connection.timeout.ms=18000
```

Data Tap Admin Section

Data Tap Admins can navigate to **Manage > Data Tap Admin > Data Taps** page to view and configure all available data taps. The data taps are configured per **Tenant**.

Figure 81: All Available Data Taps

Name	Topic	Description	Kafka Broker	Type	Status	Actions
DataTap1	default-datatap1-topic01	The First Data Tap	b4kafka3.tetrationanalytics.com:9092	External	Active	  
DataExport	DataExportTopic-610881bf497d4f7bd287a224	DataTap Managed by Tetration	172.21.156.186:443	Internal	Active	
Alerts	topic-610881bf497d4f7bd287a224	DataTap Managed by Tetration	172.21.156.186:443	Internal	Active	
Policy Stream 1 ALPHA	Policy-Stream-1	Tetration Network policy for Tenant1	172.21.156.186:443	Internal	Active	

Adding New Data Tap

Data Tap Admins can click on the  to add new data tap

Figure 82: Adding New Data Tap

New Data Tap

Name

Description

Kafka Broker

Enter Topic Name here

Topic



Note Changing any Data Tap values will require settings to be validated.

Deactivating a Data Tap

To temporarily prevent messages from leaving Secure Workload a Data Tap Admin can deactivate a data tap. Any messages to that data tap will not be sent. The data tap can reactivated at any time.

Figure 83: Deactivating a Data Tap

Name	Topic	Description	Kafka Broker	Type	Status	Actions
DataTap1	default-datatap1-topic01	The First Data Tap	b4kafka3.tetrationanalytics.com:9092	External	Active	
DataTap2	default-datatap2-topic02	The Second Data Tap	b4kafka3.tetrationanalytics.com:9093	External	Active	

Deleting a Data Tap

Deleting a datatap will delete any Secure Workload Apps instances that depend on that app. For example, if a user has specified that Compliance alerts should be sent to DataTap A (in the alerts Secure Workload app), and an admin deletes DataTap A, then the Alerts app will no longer list DataTap A as an alert output.

Managed Data Taps

Managed Data Taps (MDT) are Data Taps hosted within the Secure Workload cluster. It is completely secure in terms of authentication, encryption and authorization. To send and receive messages from MDTs, clients need to be authenticated, and data sent over the wire is encrypted, and only authorized users can read/write messages from/to Secure Workload MDT. Secure Workload provides Client certificates to be downloaded from the UI. Secure Workload uses Apache Kafka 1.1.0 as the messages broker, and, recommends clients to use secure clients compatible with the same version.

MDTs are automatically created upon the creation of root scope. Every root scope has an Alerts MDT created. To pull any alerts out from the Secure Workload cluster, user needs to use the Alerts MDT. Data Tap Admin users are the only ones who can download the certificates. Users can only view MDT that belong to their **root scope**.

Figure 84: List of configured Data Taps

Name	Topic	Description	Kafka Broker	Type	Status
Alerts	topic-610881bf497d4f7bd287a224	DataTap Managed by Tetration	172.21.156.186:443	Internal	Active
b4kafka3	default-b4kafka3-preparedemo	Cisco Building 4 Kafka Instance	b4kafka3.tetrationanalytics.com:9092	External	Active


All Secure Workload App alerts are sent to MDT by default, but can be changed to other Data Taps. There are two choices for downloading the certs:

1. JKS (Java Keystore format). JKS format works well with Java Client
2. Certs. Regular certs are easier to use with Go Clients.

Figure 85: Download

Name	Topic	Description	Kafka Broker	Type	Status	Actions
Alerts	topic-610881bf497d4f7bd287a224	DataTap Managed by Tetration	172.21.156.186:443	Internal	Active	Download Client Certificate
DataExport	DataExportTopic-610881bf497d4f7bd287a224	DataTap Managed by Tetration	172.21.156.186:443	Internal	Active	
DataTap1	default-datatap1-topic01	The First Data Tap	b4kafka3.tetrationanalytics.com:9092	External	Active	

Figure 86: Cert types



Java Keystore

Upon downloading the Alerts.jks.tar.gz, user you should see the following files that contain information to connect to Secure Workload MDT to receive messages:

1. kafkaBrokerIps.txt - This file contains the IP address string, that kafka client should use to connect to Secure Workload MDT.
2. topic.txt - This file contains the topic this client can read the messages from. Topics are of the format topic<root_scope_id>. This root_scope_id can be used later while setting up other properties in Java Client
3. keystore.jks - Keystore the Kafka Client should use in the connection settings shown below.
4. truststore.jks - Truststore the Kafka Client should use in the connection settings shown below.
5. passphrase.txt - This file contains the password to be used for #3 and #4.

Following the Kafka settings should be used while setting up Consumer.properties (Java client) that uses the keystore and truststore:

```
security.protocol=SSL
ssl.truststore.location=<location_of_truststore_downloaded>
ssl.truststore.password=<passphrase_mentioned_in_passphrase.txt>
ssl.keystore.location=<location_of_truststore_downloaded>
ssl.keystore.password=<passphrase_mentioned_in_passphrase.txt>
ssl.key.password=<passphrase_mentioned_in_passphrase.txt>
```

Following set of Properties should be used while setting up the Kafka Consumer in Java code:

```
Properties props = new Properties();
props.put("bootstrap.servers", brokerList);
props.put("group.id", ConsumerGroup-<root_scope_id>); // root_scope_id is same as
mentioned above
props.put("key.deserializer",
"org.apache.kafka.common.serialization.StringDeserializer");
props.put("value.deserializer",
"org.apache.kafka.common.serialization.StringDeserializer");
props.put("enable.auto.commit", "true");
props.put("auto.commit.interval.ms", "1000");
props.put("session.timeout.ms", "30000");
props.put("security.protocol", "SSL");
props.put("ssl.truststore.location", "<filepath_to_truststore.jks>");
props.put("ssl.truststore.password", passphrase);
```

```
props.put("ssl.keystore.location", <filepath_to_keystore.jks>);
props.put("ssl.keystore.password", passphrase);
props.put("ssl.key.password", passphrase);
props.put("zookeeper.session.timeout.ms", "500");
props.put("zookeeper.sync.time.ms", "250");
props.put("auto.offset.reset", "earliest");
```

Certificate

If end user wants to use Certificates, they can use Go clients using Sarama Kafka library to connect to Secure Workload MDT. Upon downloading Alerts.cert.tar.gz, user should see the following files:

1. kafkaBrokerIps.txt - This file contains the IP address string that Kafka Client should use to connect to Secure Workload MDT
2. topic - This file contains the topic this client can read the messages from. Topics are of the format topic<root_scope_id>. This root_scope_id can be used later while setting up other properties in Java Client.
3. KafkaConsumerCA.cert - This file contain the KafkaConsumer certificate.
4. KafkaConsumerPrivateKey.key - This file contains the Private Key for the Kafka Consumer.
5. KafkaCA.cert - This file should be used in the root CA certs listing in the Go client.

See the following example of Go Client to connect to Secure Workload MDT. (Attach the Sample Go Code) [Sample Go Client to consume alerts from MDT](#)