



Cluster Maintenance

This chapter provides details about various cluster maintenance actions that you can perform, such as upgrade, reboot, schedule data backups, and restore data. You can also view the service and cluster status from the options available under the **Troubleshoot** menu.

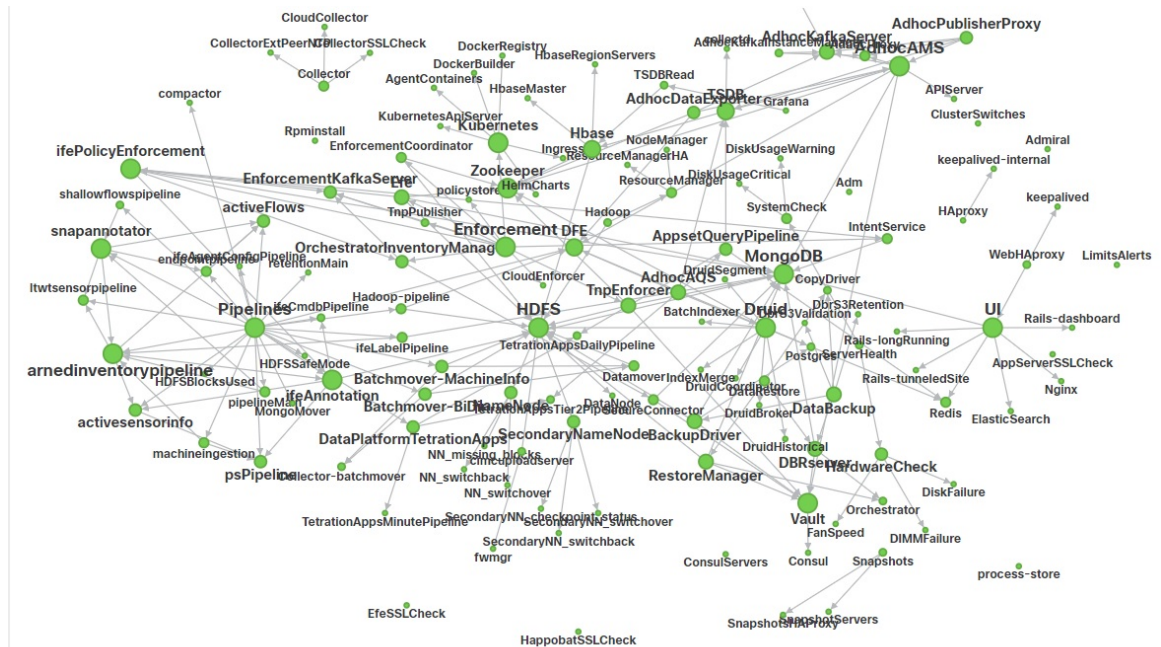
- [Service Status, on page 1](#)
- [Admiral Alerts, on page 2](#)
- [Cluster Status, on page 11](#)
- [Data Backup and Restore, on page 14](#)
- [High Availability in Secure Workload, on page 35](#)
- [VM Information, on page 43](#)
- [Upgrading a Secure Workload Cluster, on page 43](#)
- [Secure Workload Cluster Snapshots, on page 52](#)
- [Overview of Explore or Snapshot Endpoints , on page 60](#)
- [Server Maintenance, on page 72](#)
- [Disk Maintenance, on page 79](#)
- [Requirement Prechecks, on page 79](#)
- [Disk Replacement Wizard-Not Hot Swap, on page 83](#)
- [Cluster Maintenance Operations, on page 92](#)
- [Data Tap Admin: Data Taps, on page 96](#)

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 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 that is not user configurable.

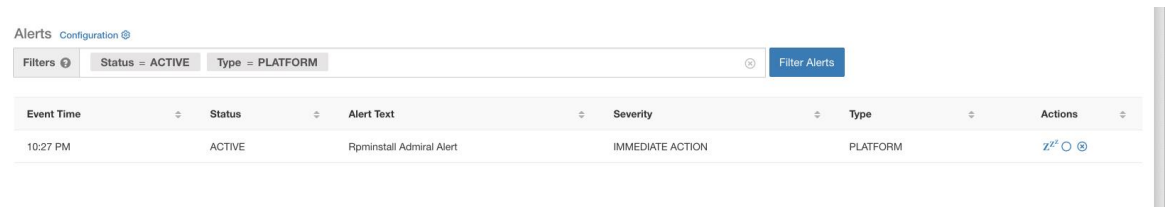
The following sections describe admiral alerts and notifications in more detail.

Lifecycle of an Admiral Alert

The admiral checks for the uptime of services on service status. It raises an alert when this uptime becomes lower than the preconfigured threshold for alerting.

As an example, Rpminstall is a service which is used to install RPMs during deploys, upgrades, patches, and so on. 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' interface. At the top, there are filters: 'Status = ACTIVE' and 'Type = PLATFORM'. Below the filters is a table with the following data:

| Event Time | Status | Alert Text | Severity | Type | Actions |
|------------|--------|--------------------------|------------------|----------|---|
| 10:27 PM | ACTIVE | Rpminstall Admiral Alert | IMMEDIATE ACTION | PLATFORM | Z O @ |

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 automatically close when its uptime goes over 80% in one hour.



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

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' interface with filters set to 'Status = CLOSED' and 'Type = PLATFORM'. Below the filters is a table with the following data:

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

There are two kinds of admiral alerts:

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

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

Alerts Configuration

Filters Status = ACTIVE Type = PLATFORM Alert Text contains Admiral Alert Filter Alerts

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

Summary Admiral Alerts

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

Alerts Configuration

Filters Status = ACTIVE Type = PLATFORM Alert Text contains Admiral Summary Filter Alerts

| Event Time | Status | Alert Text | Severity | Type | Actions |
|------------|--------|-----------------|----------|----------|---|
| 5:04 PM | ACTIVE | Admiral Summary | LOW | PLATFORM | Z Z Z |

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 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'}/e95da4521012a4789048f72a791fb58ab233bbff63e6cbcb421525d4272d469aa",

  "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

Filters
Status = ACTIVE
Type = PLATFORM
Alert Text contains Admiral Summary
Filter Alerts

| Event Time | Status | Alert Text | Severity | Type | |
|------------|--------|-----------------|----------|----------|----------------------|
| 5:04 PM | ACTIVE | Admiral Summary | LOW | PLATFORM | Force close an alert |

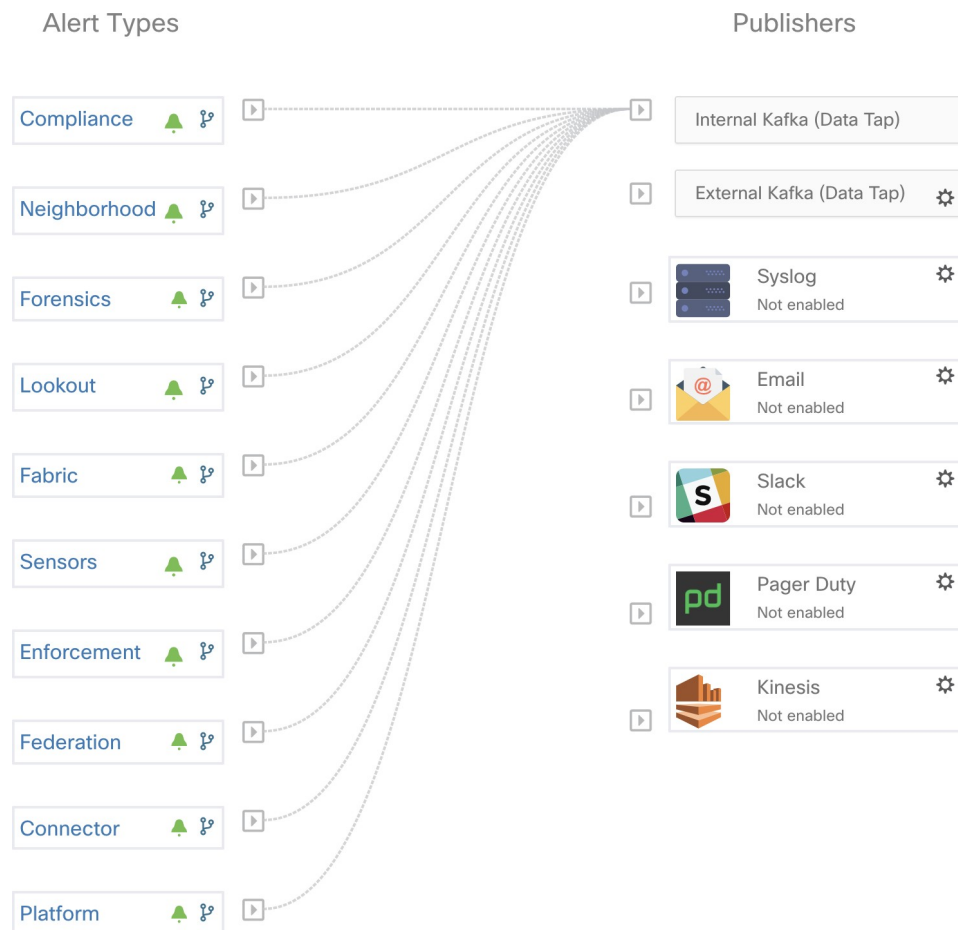
**Warning**

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

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 | CIMC Snapshots |
|--------------------------|--------------|--------|-------------|-------------|--------------------|----------------|
| <input type="checkbox"/> | Commissioned | Active | Ethernet1/1 | FCH2206V1NF | 2mo 27d 13h 3m 47s | + - |
| <input type="checkbox"/> | Commissioned | Active | Ethernet1/2 | FCH2206V1ZF | 2mo 27d 13h 2m 52s | + - |

Serial: FCH2206V1ZF

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

- collectorData mover-6
- datanode-6
- druidHistoricalBroker-4
- enforcementCoordinator-3
- orchestrator-2
- redis-1
- secondaryNameNode-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

Switch Port: Ethernet1/2

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

Commissioned Active Ethernet1/1 FCH2206V1NF 2mo 27d 13h 17m 47s + -

Serial: FCH2206V1NF

Private IP: 1.1.1.8
CIMC IP: 10.13.4.11
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

External access to CIMC UI is enabled

- 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

- adhocKafkaXL-1
- collectorData mover-5
- datanode-5
- druidHistoricalBroker-3
- elasticsearch-3
- namenode-1
- orchestrator-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

Switch Port: Ethernet1/1

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 Secure Workload 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 | <p>Select this action to power down the nodes.</p> <p>Note You cannot power down the nodes with the Inactive and Shutdown in progress status.</p> |

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

- Contact [Cisco Technical Assistance Center](#) to enable the data backup and restore options on your 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 upgrade. 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}} |

| IP Type | Default 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

Cisco Secure Workload

You do not have an active license. The evaluation period will end on Mon Nov 08 2021 23:15:57 GMT+0000. Take action now.

Cluster Configuration

Model: BRU-PROD

| | |
|--|---|
| Cluster UUID | 3b478c4d-6883-8861-c6e4-41bbca0d8d0 |
| Admin 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




- Note**
- If the **Data Backup** link under **Platform** is not available, contact [Cisco Technical Assistance Center](#) to enable the data backup and restore options.
 - If the cluster is in standby mode, you will not be able to view the **Data Backup** link.

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

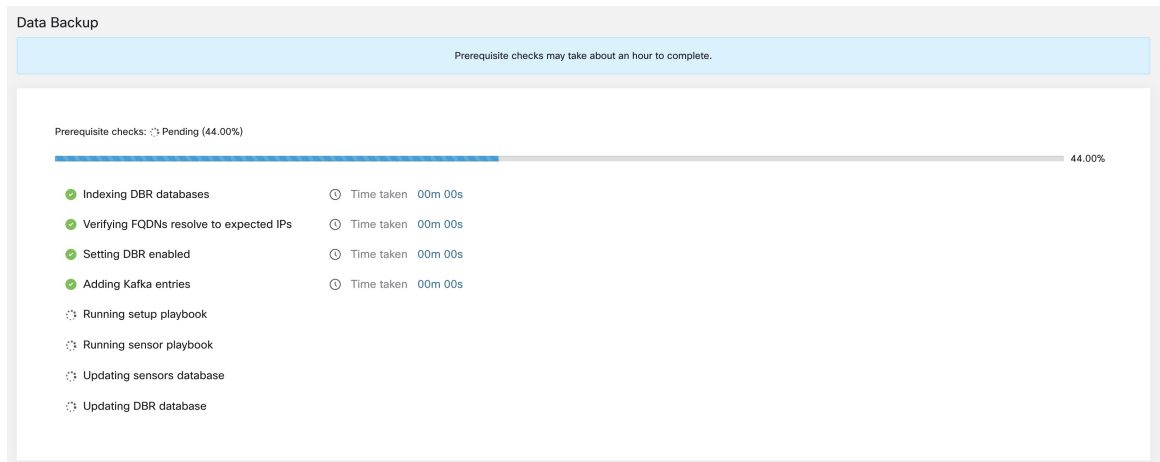
Procedure

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 Storage
 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 ⓘ

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 ⓘ
 Mbps

Est. Sensor Count ⓘ

Lean Data Mode ⓘ
☐

Retention ⓘ
 days

Est. Backup duration ⓘ
 -

Est. Max Storage ⓘ
 TB

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

CONFIGURE SCHEDULE

Configure Storage ☒ Configure Backup ☒ **Schedule Backup** ☒ Review ☐

Set starting backup point from today ☒

Continuous backup ☐

Timezone: America/Los_Angeles

Allowed start backup window ☐

Every: Day at 0:00

Enable recurring full backup ☐

Cancel Previous Next

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

Figure 22: Backup Configuration Review

CONFIGURE SCHEDULE

Configure Storage

Configure Backup

Schedule Backup

Review

| Storage | |
|------------|-------------------------|
| Name | cohesity |
| Bucket | dbr-erdos |
| Access Key | vCEASJuz5frJavfHPNSg... |

| Backup | |
|-----------------------|-----------------|
| Window | 23:15 every day |
| Duration | 4 hrs 46 min |
| Recurring Full Backup | Not scheduled |

| Bandwidth | |
|--------------|----------|
| Sensor count | 350 |
| Observed | 64 Mbps |
| Max allowed | 300 Gbps |

| Backup details | |
|---------------------------|-------|
| Required Storage / backup | 128GB |
| Allowed Storage | 189TB |
| Retention (days) | 60 |

Cancel

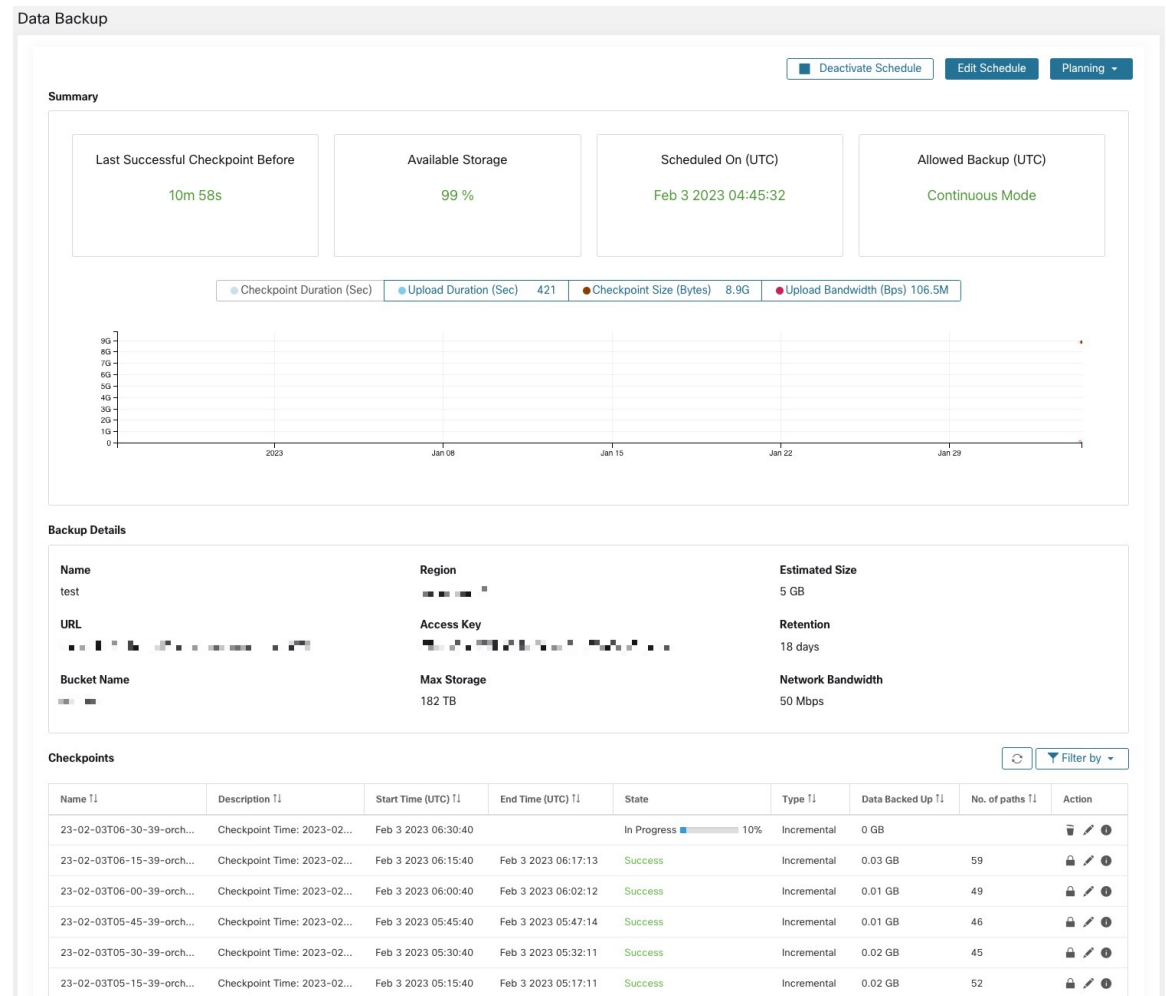
Previous

Initiate Job

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 33](#).

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

- To restore using backed up data, a cluster must be in the **DBR standby mode**. Currently, you can set a cluster to standby mode **only during initial setup**.
- After the cluster is in standby mode, choose **Platform** from the navigation pane to access the data restore option.

Secure Workload supports the following combinations:

| 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.



Note The Kafka and sensor FQDNs *must* match with the primary cluster, else the restore process fails.

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

Standby Config ☒

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.

Restore Data to a Secure Workload Cluster

Before you begin

Ensure that the cluster is deployed in standby mode. For more details, see [Deploy Cluster in Standby Mode](#).

Procedure

- Step 1** (Optional) If you have already configured the storage details, go to Step 2. To configure S3 storage, enter the following details:
- A name for the storage.
 - The URL of an S3-compliant storage endpoint.
 - An S3-compliant bucket name configured on the endpoint 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 necessary.
 - (Optional) If a CA certificate is required to authenticate the storage server, enable **Use Server CA Certificate** and enter the certificate details.
- Step 2** Click **Test** to check if the S3 storage is accessible from the Secure Workload cluster.
- The status of the tests that are performed are displayed in the table. If there are any errors connecting to the storage, read the description and troubleshoot the errors to continue to the next step.
- Step 3** Click **Next**.
- Step 4** Under **Pre-checks**, the status of the prechecks runs by Secure Workload are displayed. To manually run the prechecks, click **Perform Check**.
- The status of all the checks is displayed:
- For the checks that have an error, but do not prevent you from restoring the data, hover your cursor over the warning icon to get the details and a link to navigate to the **Service Status** page to get more details of the service.
 - If any of the checks failed, you must troubleshoot the issue to proceed with data restore. Navigate to the **Service Status** page to get more details of the service.
- Note** Ensure that the checkpoint you are restoring to is the latest with no errors.
- Step 5** Click **Start restore process**.
- Under **Restore**, all the data restore jobs that run, the configured S3 storage details, and the status of the data restore prechecks are displayed.
- Step 6** Click **Restore now**.
- Step 7** In the confirmation dialog box, check the check boxes to confirm that you agree to the fact that agent connectivity is lost and data may be lost during the data restore. Click **Confirm** to start the data restore process.
- The progress of the data restore process is displayed.

Caution At the **Pre Restore Playbook** stage, all the services within the cluster are reinitialized and there is a downtime of approximately two hours. At this stage, the Secure Workload GUI is not accessible. For more information about the phases that are involved in data restore, see [Cluster Restore Phases](#).

If the GUI is rendered inaccessible for an extended period, contact the [Cisco Technical Assistance Center](#) to troubleshoot the issue.

Note

After the **Post Restore Playbook** stage, the GUI is accessible and the status of all the jobs are updated. A confirmation message is displayed indicating that the data restore is successful.

What to do next

Update your DNS server to redirect the configured FQDNs to the cluster IP address, which ensures that the software agents communicate with the cluster after the cluster failover is complete.

Prefetch Cluster Data

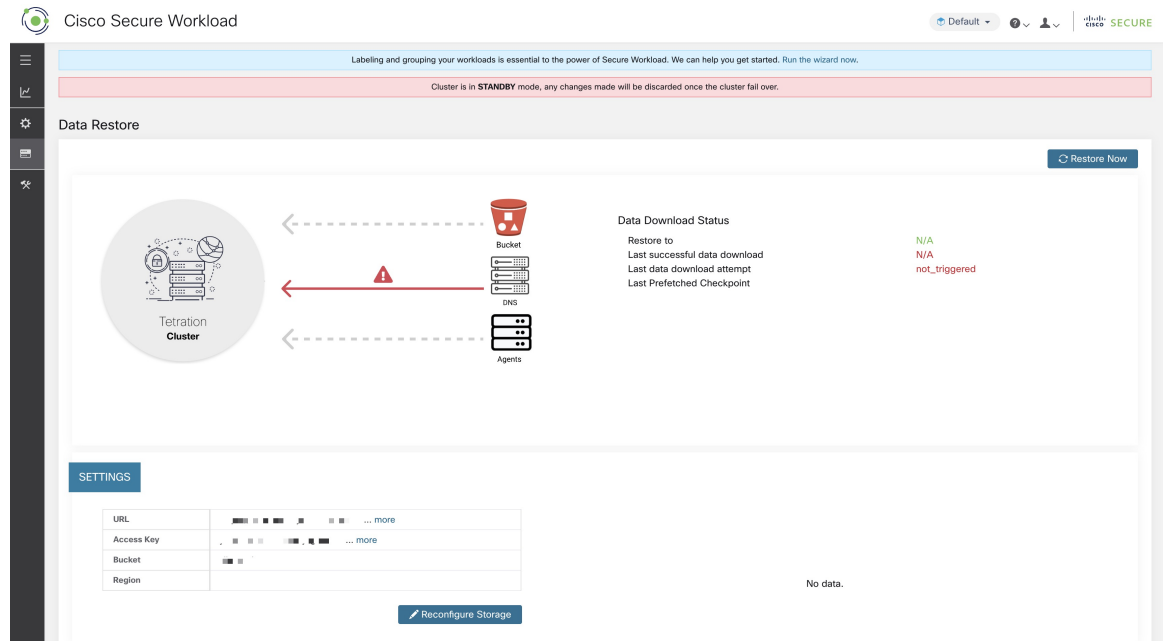
Before the cluster can be restored, it must prefetch data. The checkpoint data is prefetched from the same storage bucket that is used for backing up data. Credentials must be provided for the backup service to download from the storage. If a storage is not set up for prefetch, the **Data Restore** tab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

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 that are 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 must be accessed directly.
- The multipart 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 that is 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 that is 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 Mongoddb 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

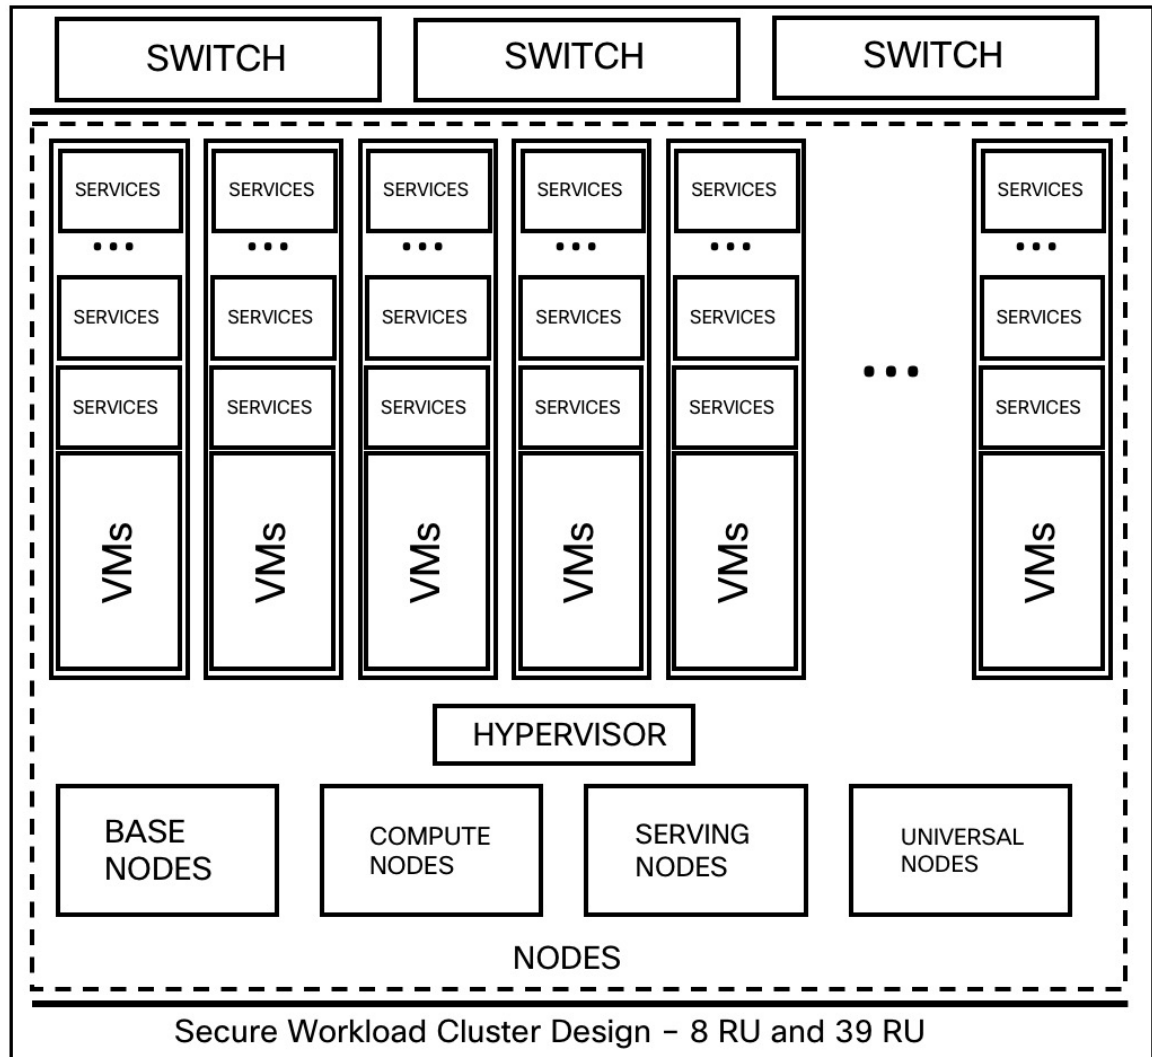


Table 10: Secure Workload Cluster Components

| 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

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

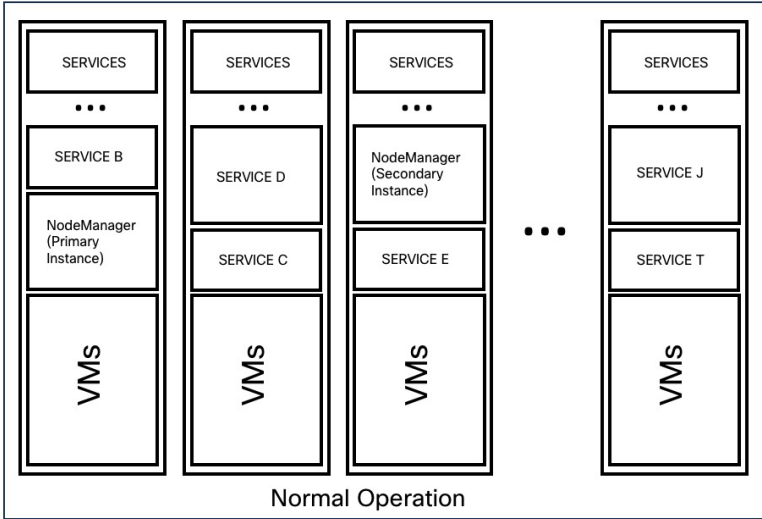


Figure 28: Failure Scenario of a Service

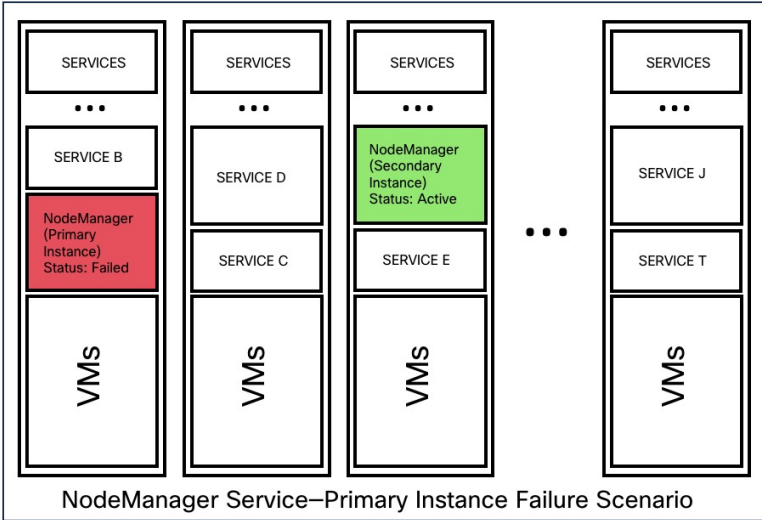


Table 11: Services Failure Impact and Recovery

| | |
|----------|---|
| 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

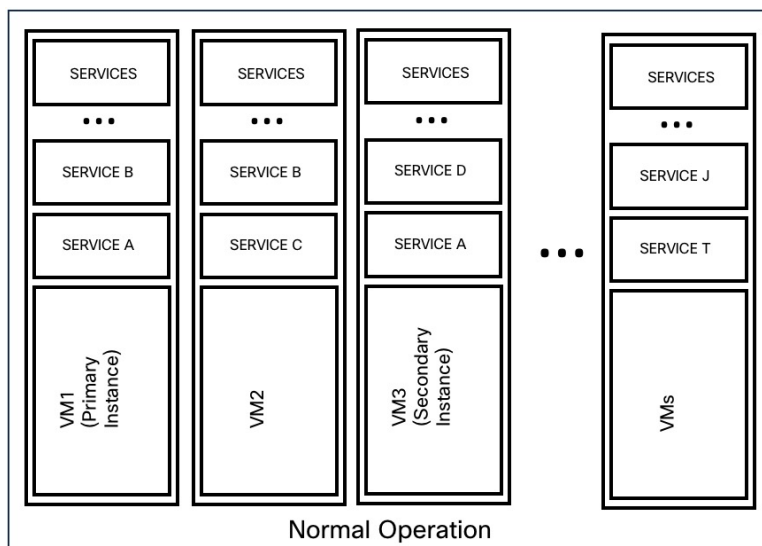
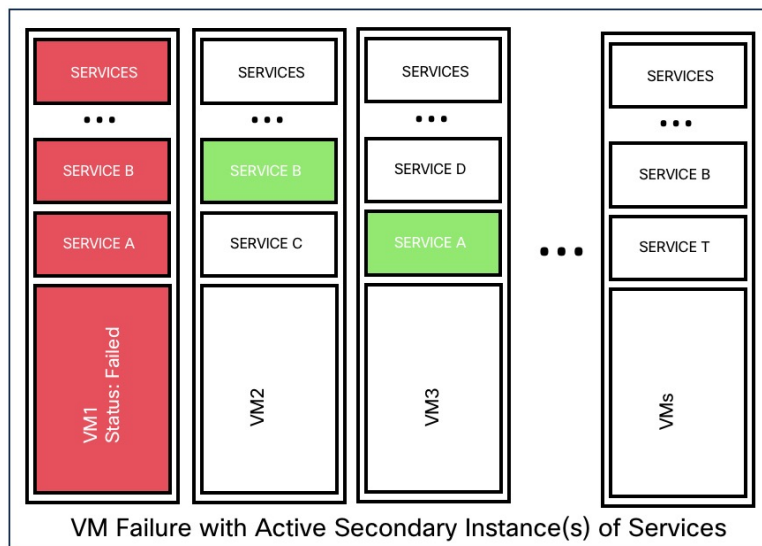


Figure 30: Failure Scenario of a VM



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

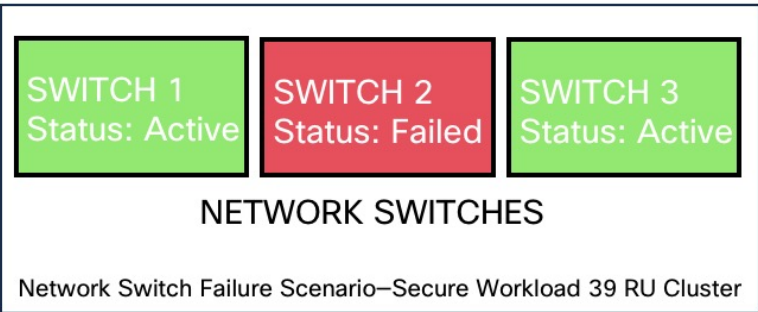


Table 16: Number of Switch Failures Tolerated

| Form Factor | 8 RU | 39 RU |
|--|--|--|
| Number of switch failures that are tolerated for high availability | 1 | 1 |
| | Note If two or more switches fail, it is likely to have an impact on the entire functionality of the cluster. | 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 a Secure Workload Cluster

Secure Workload supports two types of upgrade—Full upgrade and patch upgrade. The following sections describe the full upgrade process. During the full upgrade, all VMs in the cluster are shut down, new VMs are deployed, and the services are reprovisioned. All the data within the cluster are persisted during this upgrade, except for downtime during the upgrade.

Cluster Upgrade Options

Supported upgrade types for a Secure Workload cluster:

- **Full Upgrade:** To initiate full upgrade, from the navigation pane, choose **Platform > Upgrade/Reboot/Shutdown**. In the **Upgrade** tab, select **Upgrade**. During the full upgrade process, the VMs are powered off, and VMs are upgraded and redeployed. There is a cluster downtime during which the Secure Workload UI is inaccessible.
- **Patch Upgrade:** Patch upgrade minimizes the cluster downtime. The services that must be patched are updated and do not result in VM restarts. The downtime is usually in the order of a few minutes. To initiate patch upgrade, select **Patch Upgrade** and click **Send Patch Upgrade Link**.

An email with a link is sent to the registered email address to initiate the upgrade.

Figure 35: Email with the Upgrade Link

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, the 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, covering the following:
 - Drive failure
 - Drive predicted failure.
 - Drive missing
 - StorCLI failures
 - MCE log failures
- Checks to ensure that bare metals are in the commissioned state, nothing fewer than 36 servers for 39RU and six for 8RU.



Note

If there are any failures, an upgrade link is not sent to the registered email address and a 500 error is displayed with information such as HW failure, or missing host and check orchestrator logs for more information. In this scenario, use `explore to tail -100` on `/local/logs/tetration/orchestrator/orchestrator.log` in the host `orchestrator.service.consul`. The log provides detailed information about which of the three checks caused the failure. It usually requires fixing the hardware and recommissioning the node. Restart the upgrade process.







Upload RPMs

After you click the upgrade link that is received on your email, the **Secure Workload Setup** page is displayed. The Setup UI is used to deploy or upgrade the cluster. The landing page displays the list of RPMs that are currently deployed in the cluster. You can upload the RPMs to upgrade the cluster.

Figure 36: RPM Upload

Tetration Setup Diagnostics » RPM Upload » Site Config » Site Config Check » Run norfleet

RPM Upload

| | | |
|---|---|---------------|
| ✓ |  tetration_os_rpminstall_k9 | 3.5.0.7.devel |
| ✓ |  tetration_os_UcsFirmware_k9 | 3.5.0.7.devel |
| ✓ |  tetration_os_adhoc_k9 | 3.5.0.7.devel |
| ✓ |  tetration_os_mother_rpm_k9 | 3.5.0.7.devel |
| ✓ |  tetration_os_enforcement_k9 | 3.5.0.7.devel |
| ✓ |  tetration_os_base_rpm_k9 | 3.5.0.7.devel |

Select RPM file

No file selected.

Upload the RPMs in the order that is shown on setup UI. Not uploading RPMs in the correct order result in upload failures and the **Continue** button remains disabled. 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, ensure that you also upgrade the tetration_os_ova_k9 RPM and that you do not upload the tetration_os_base_rpm_k9 RPM.

To view the logs for each upload, click 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 norfleet

RPM Upload

| | | |
|---|-----------------------------|---------------|
| ✓ | tetration_os_rpminstall_k9 | 3.5.0.7.devel |
| ✓ | tetration_os_UcsFirmware_k9 | 3.5.0.7.devel |
| ✓ | tetration_os_adhoc_k9 | 3.5.0.7.devel |
| ✓ | tetration_os_mother_rpm_k9 | 3.5.0.7.devel |
| ✓ | tetration_os_enforcement_k9 | 3.5.0.7.devel |
| ✓ | tetration_os_base_rpm_k9 | 3.5.0.7.devel |

Select RPM file

tetration_os_enforcement_k9-3.5.0.8.devel.rpm

verifying RPM...

RPM downloaded

RPM install failed

Figure 38: Upload RPMs

Cisco Secure Workload Setup Diagnostics » Software Upgrade » Site Config » Site Config Check » Run 1

Software Upgrade

To upgrade Secure Workload, perform these steps:

1. Download the RPMs from CCO.
2. Upload "tetration_os_rpminstall_k9" package and click Install.
3. Upload the dependent RPMs and click Install.

| Package Name | Currently Installed Version | Staged Version | Status |
|-----------------------------|-----------------------------|----------------|--------|
| tetration_os_rpminstall_k9 | 3.9.0.14.devel | | |
| tetration_os_UcsFirmware_k9 | 3.9.0.14.devel | | |
| tetration_os_nxos_k9 | 3.9.0.14.devel | | |
| tetration_os_adhoc_k9 | 3.9.0.14.devel | | |
| tetration_os_mother_rpm_k9 | 3.9.0.14.devel | | |
| tetration_os_enforcement_k9 | 3.9.0.14.devel | | |
| tetration_os_base_rpm_k9 | 3.9.0.14.devel | | |

Select RPM file

No file selected.

There are no RPMs staged for install.

For detailed instructions, see the [Cisco Secure Workload Upgrade Guide](#).

Site Information

The next step in upgrading the cluster 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**. Configure TAN Appliance to export data to syslog. For more details, see [External syslog tunneling moving to TAN](#).
 - Secure Workload supports secure SMTP communication with mail servers that support SSL or TLS communication using 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.
-

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

Preupgrade Checks

Before upgrading the cluster, a few checks are performed on the cluster to ensure things are in order. The following preupgrade checks are performed:

- **RPM version checks:** Checks to ensure all the RPMs are uploaded and the version is correct. It does not check if the order was correct, just checks if it was uploaded. Note that order checks are done as a part of the upload itself.
- **Site Linter:** Performs Site Info linting
- **Switch Config:** Configures the Leafs or Spine switches
- **Site Checker:** Performs DNS, NTP, and SMTP server checks. Sends an email with a token, the email is sent to the primary site admin account. If any of the services - DNS, NTP or SMTP is not configured, this step fails.
- **Token Validation:** Enter the token that is sent in the email and continue the upgrade process.

Upgrade Secure Workload Cluster



Caution

- We recommend that you do not select the **Ignore Stop Failures** option. This is a recovery option for upgrade failures when certain services do not shut down. Using this option shuts down the VMs that can create failures when the services become active.
- Use this option under supervision.

Figure 39: Upgrading the Cluster

| Serial | Base/Model ID | Instance Type | Instance Index | Private IP | Public IP | Uptime | Status | Deploy Progress | |
|-------------|---------------|------------------|----------------|------------|-----------|----------|---------|-----------------|--------------------------|
| FD42111V2R0 | 1.1.1.5 | baseRegionServer | 2 | 1.1.1.29 | | 12 hours | Stopped | 100% | View Log |
| FD42113V0W0 | 1.1.1.7 | adhoc | 2 | 1.1.1.83 | | 12 hours | Stopped | 100% | View Log |
| FD42112V1X0 | 1.1.1.9 | adhoc | 1 | 1.1.1.82 | | 12 hours | Stopped | 100% | View Log |
| FD42113V0W0 | 1.1.1.7 | haproxy | 2 | 1.1.1.81 | | 12 hours | Stopped | 100% | View Log |
| FD42111V2M7 | 1.1.1.6 | haproxy | 1 | 1.1.1.80 | | 12 hours | Stopped | 100% | View Log |

Before you begin

Complete the preupgrade checks and enter the token that is received in the *verify token email*.

Procedure

Step 1

Click **Continue** to start the upgrade.

Step 2

(Optional) Click the cluster name to view the site information.

The Secure Workload RPMs and the versions are displayed. The upgrade bar displays the upgrade progress. Blue color indicates activities that are in progress, green color indicates the activities that are complete, and red color indicates the activities that failed.

Four buttons are available:

- Refresh: Refreshes the page.
- Details: Click **Details** to view the steps that have been completed during this upgrade. Click the arrow next to it to display the logs.
- Reset: This has an option to Reset Orchestrator State. This option cancels the upgrade and take you back to the start. DO NOT use this unless the upgrade had failed and a few minutes have passed after the upgrade had failed to let all the processes reach completion before restarting the upgrade.
- Restart: When an upgrade fails, click **Restart** to restart the cluster and initiate a new upgrade. This can help resolve any pending cleanup operations or issues that may be blocking the upgrade processes.

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

- Serial: Bare metal serial that hosts this VM
- Baremetal IP: the Internal IP assigned to the bare metal
- Instance Type: the type of VM
- Instance Index: Index of the VM - there are multiple VMs of the same type for high availability.
- Private IP: the Internal IP assigned to this VM
- Public IP: the routable IP assigned to this VM - not all VMs have this.
- Uptime: Uptime of the VM
- Status: Can be Stopped, Deployed, Failed, Not Started or In Progress.
- Deploy Progress: Deploy Percentage.
- View Log: Button to view the deploy status of the VM

Cluster Upgrade Logs

There are two types of logs:

Procedure

-
- Step 1** VM deployment logs: Click **View Log** to view VM deployment logs.
 - Step 2** Orchestration Logs: Click the arrow next to the **Details** button to view orchestration logs.

Figure 40: Orchestration Logs

Running playbooks on the instances ...

The screenshot shows a web interface for viewing orchestration logs. At the top, there's a header "Running playbooks on the instances ...". Below it, there are three buttons: "Refresh", "Details" (which is currently selected and has a dropdown arrow), and "Reset". A dropdown menu is open from the "Details" button, listing various log categories: Orchestrator, Orchestrator-Upgrade, Orchestrator-consul, Orchestrator-scheduler, Orchestrator-server, Playbooks-Orch-bare_metal, Playbooks-Orch-bigbang, Playbooks-Orch-consul_server, Playbooks-Orch-get_upgrade_logs, Playbooks-Orch-orchestrator_during_instance_deploy, Playbooks-Orch-orchestrator_postinstall_setup, Playbooks-Orch-orchestrator_setup, Playbooks-Orch-pre_orchestrator_setup, Playbooks-Orch-switch_config, SiteInfoChecker, and VM Manager. In the background, a table is visible with columns "Instance", "Serial", and "Instance Type". The "Instance" column contains links to logs, and the "Instance Type" column contains the type of instance.

| Instance | Serial | Instance Type |
|-----------------------------|--------|-------------------|
| FCH2111V... | | hbaseRegionServer |
| FCH2111V... | | adhocKafkaXL |
| FCH2113V... | | happobat |
| FCH2111V... | | happobat |
| FCH2111V... | | zookeeper |
| FCH2112V... | | zookeeper |
| FCH2111V... | | zookeeper |
| FCH2112V... | | datanode |

Each of the links point to the logs.

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

Run Preupgrade Checks

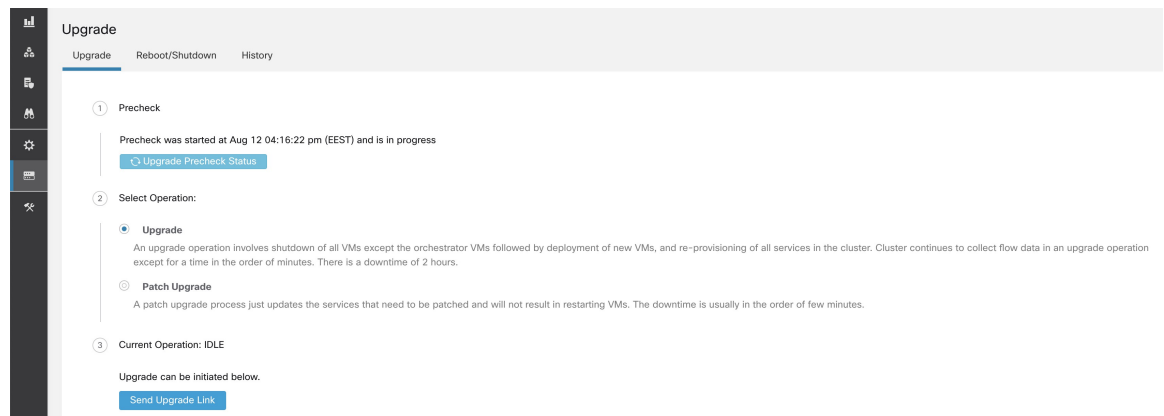
Occasionally, there may be hardware failures or the cluster may not be ready to be upgraded after scheduling an upgrade and while initiating an upgrade. These errors must be fixed before proceeding with upgrades. Instead of waiting for an upgrade window, you can initiate preupgrade checks which can be run any number of times and anytime, except when upgrade, patch upgrade, or reboot is initiated.

To run preupgrade checks:

1. In the **Upgrade** tab, click **Start Upgrade Precheck**.

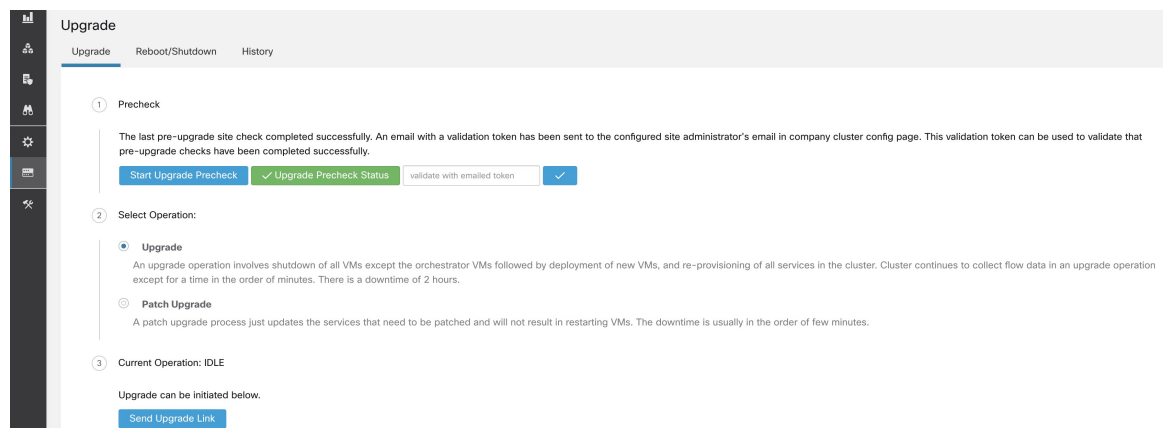
This initiates the preupgrade checks and is transitioned to a running state.

Figure 41: Running Preupgrade Checks



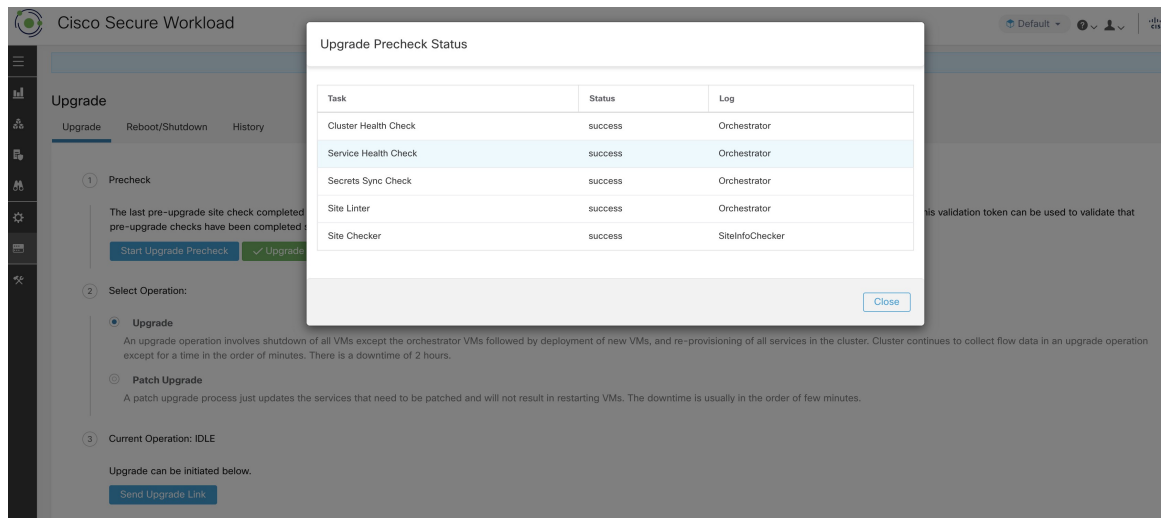
2. After all the checks that are run by orchestrators pass, an email with a token is sent to the registered email ID. Enter the token to complete the preupgrade checks.

Figure 42: Enter Token for Preupgrade Checks



You can verify the status of the checks. If there are any failures during preupgrade checks, you can view the failed checks and the respective check transition to the failed state.

Figure 43: Status of Preupgrade Checks



Data Backup and Restore (DBR)

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

Secure Workload Cluster 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 44: Snapshot Tool Runner

logs

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

yarn log state:

application status (RUNNING, FAILED, KILLED, UNASSIGNED) used to get information for, default: all

alerts

start date:

number of days to fetch alert history for, default: 10

comments

These comments will appear in the snapshots for:

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

Figure 45: CIMC Technical Support Runner

CIMC serial number

select a node's serial number to collect cimc logs for

Create a Snapshot

Select **Create Snapshot** with the default options, the Snapshot tool collects:

- Logs
- State of Hadoop or 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 is displayed in the Snapshot file list page and the comments section will reflect that the CIMC Technical Support Bundle collection has been triggered. When 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 46: 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 47: Terse List of Alert State Changes

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

- Reproduction of grafana dashboards.

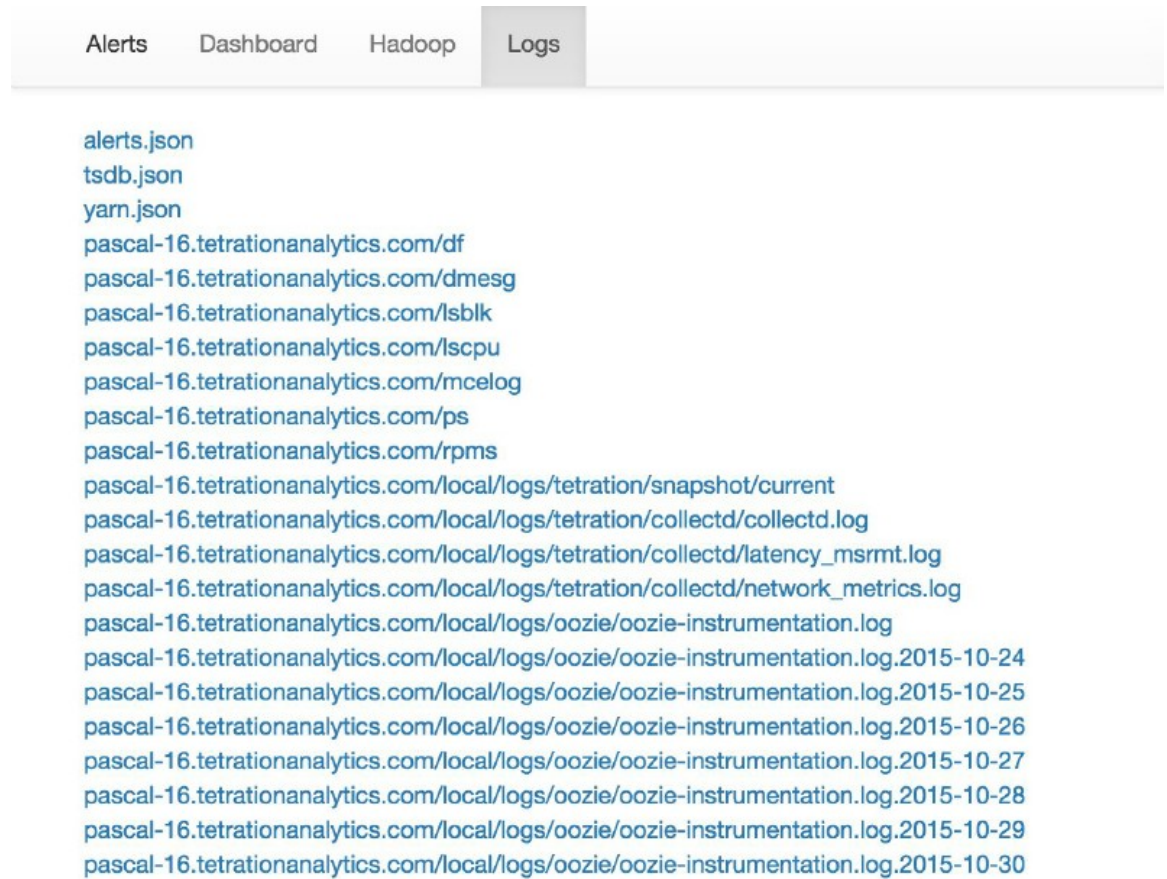
Figure 48: Reproduction of Grafana Dashboards

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

Figure 49: Reproduction of 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 50: List of 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 51: Snapshot Service for Debugging and Maintenance

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 52: Using the Snapshot Service for Debugging and Maintenance

POS

Status: 200

```
total 52
drwxr-xr-x 2 root    users    4096 Nov  3 20:08 BDPipeline
drwxr-xr-x 2 root    users    4096 Nov  3 20:22 activeflowpipeline
drwxr-xr-x 5 tetter   tetter    4096 Jun  9 22:53 adm
drwxr-xr-x 2 collectd collectd 4096 Oct 17 04:29 collectd
drwxr-xr-x 4 druid    users    4096 Aug  7 22:08 druid
drwxr-xr-x 3 root     root     4096 Oct 12 18:08 mongo_indexer
drwxr-xr-x 2 collectd collectd 4096 Sep 15 20:49 netmond
drwxr-xr-x 2 root     root     4096 Nov  3 15:02 policy_server
drwxr-xr-x 2 root     users    4096 Oct 19 17:20 repl
drwxr-xr-x 2 tetter   tetter    4096 Nov  3 21:47 retentionPipeline
drwxr-xr-x 2 root     users    4096 Oct 14 23:28 snapshot
drwxr-xr-x 2 root     users    4096 Nov  3 22:04 ubercp_bidir
drwxr-xr-x 2 root     users    4096 Nov  3 22:03 ubercp_machineinfo
```

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 that are signed by Secure Workload. It runs gpg -d on the POSTed data. If it can be verified against a signature, it runs 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

```

Overview of Explore or Snapshot Endpoints

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

post Commands

Table 18: post Commands

| Endpoint | Description |
|------------------------------|--|
| bm_shutdown_or_reboot | <ul style="list-style-type: none"> Gracefully shut down 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?</code> <code>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 | Wrapper command for <i>cat</i> Unix command |
| cimc_password_random | <ul style="list-style-type: none"> Randomizes the CIMC password. Should be run on the orchestrator hosts only |
| cleancmdlogs | 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 needs to be rebooted to activate the BIOS and other component firmware. • Run as: cluster_fw_upgrade • This endpoint kicks 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 endpoint. |
| 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. • <i>Use with caution because the cluster is brought down.</i> • Run the endpoint as <code>cluster_powerdown?args=-start.</code> |
| 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 | Wrapper command for <i>df</i> Unix command |
| dig | Wrapper command for <i>dig</i> Unix command |
| dmesg | Wrapper command for <i>dmesg</i> Unix command |
| dmidecode | Wrapper command for <i>dmidecode</i> Unix command |

| Endpoint | Description |
|--|--|
| druid_coordinator_v1 | Displays the druid stats. |
| du | Wrapper command for <i>du</i> Unix command |
| dusorted | Wrapper command for <i>dusorted</i> Unix command |
| externalize_change_tunnel | <ul style="list-style-type: none"> Changes the collector IP that will be used to tunnel the CIMC UI Run as: 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 status of externalizing the CIMC UI 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 |
| fsck | <ul style="list-style-type: none"> Wrapper command for <i>fsck</i> Unix command Should be run on Baremetal host only |
| get_cimc_techsupport | <ul style="list-style-type: none"> Input Internal IP address of bare metal. Retrieves the CIMC technical support bundle. After it is completed, it will be available for download from the snapshots page in the UI. This can be run from any host on the cluster and requires the bare metal internal IP address as an argument. Example: get_cimc_techsupport?args=1.1.0.9 |

| Endpoint | Description |
|-------------------------|--|
| syslog_endpoints | <ul style="list-style-type: none"> Controls the syslog configurations for one or more of the UCS servers. Run the command with <i>-h</i> to get a full list of parameters. |
| grep | Wrapper command for <i>grep</i> Unix command |
| hadoopbalancer | <ul style="list-style-type: none"> Distributes HDFS data uniformly across all nodes Must be run on hosts that have hdfs. For example, launcher host |
| hadoopdu | <ul style="list-style-type: none"> Prints the directory utilization of hdfs Should be run on hosts that have hdfs. For example, launcher host |
| hadoopfsck | <ul style="list-style-type: none"> Runs hadoop fsck and reports the state of the provided hdfs file system It also takes “-delete” as an argument to clear corrupt 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 that is run as: <code>hadoopfsck?args=/raw</code> To delete corrupt files, run as: <code>hadoopfsck?args=/raw -delete</code> |
| hadoopls | <ul style="list-style-type: none"> Lists the Hadoop File System Should be run on hosts that have hdfs for example launcher host. |

| 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 an inconsistency, run as: <code>hbasehbk?args=-details</code>• To repair a corrupted HBase, run as: <code>hbasehbk?args=-repair</code>• Output written to <code>/local/logs/etcd/rapid-cmls/rapid-hbasebk_log.txt</code>• <i>Repair with caution</i> |
| hdfs_safe_state_recover | <ul style="list-style-type: none">• Removes HDFS from safe state• Required if HDFS is in READ_ONLY_STATE due to 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 | Wrapper command for <i>initctl</i> Unix command |
| head | Wrapper command for <i>head</i> Unix 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 | Wrapper command for <i>ip</i> Unix 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 |
| ipmisensorlist | <ul style="list-style-type: none">• Prints the IPMI sensor information• Should be run on the Baremetal hosts only |

| Endpoint | Description |
|------------------------|--|
| jstack | Prints Java stack traces of Java threads for a given Java process or core file |
| ls | Wrapper command for <i>ls</i> Unix command |
| lshw | Wrapper command for <i>lshw</i> Unix command |
| lsuf | Wrapper command for <i>lsuf</i> Unix command |
| lvdisplay | Wrapper command for <i>lvdisplay</i> Unix command |
| lvs | Wrapper command for <i>lvs</i> Unix command |
| lvscan | Wrapper command for <i>lvscan</i> Unix command |
| makecurrent | <ul style="list-style-type: none"> Resets or fast forwards 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 enforcementpolicy hosts |
| mongo_stats | <ul style="list-style-type: none"> Displays the mongo stats Should be run on either the mongodb or the enforcementpolicy hosts |
| mongodump | <ul style="list-style-type: none"> Dumps the collections from the database Should be run on either the mongodb or the enforcementpolicy hosts Run as: <code>mongodump?args=<collection>[-db DB]</code> |
| monit | Wrapper command for <i>monit</i> Unix command |
| namenode_jmx | Displays the primary name node jmx metrics |
| ndisc6 | Wrapper command for <i>ndisc6</i> Unix command |
| netstat | Wrapper command for <i>netstat</i> Unix command |
| ntpq | Wrapper command for <i>ntpq</i> Unix command |

| Endpoint | Description |
|----------------------|--|
| 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 primary orchestrator and trigger a switchover Should be run on the orchestrator.service.consul host only USE WITH CAUTION |
| ping | Wrapper command for <i>ping</i> Unix command |
| ping6 | Wrapper command for <i>ping6</i> Unix command |
| ps | Wrapper command for <i>ps</i> Unix command |
| pv | Wrapper command for <i>pv</i> Unix command |
| pvs | Wrapper command for <i>pvs</i> Unix command |
| pvddisplay | Wrapper command for <i>pvddisplay</i> Unix command |
| rdisc6 | Wrapper command for <i>rdisc6</i> Unix 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 |
| 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> |

| Endpoint | Description |
|------------------------|--|
| 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 | Wrapper command for <i>service</i> Unix command |
| smartctl | <ul style="list-style-type: none"> • Run the smartctl executable • Should only be run on a bare metal node |
| storcli | Wrapper command for <i>storcli</i> Unix command |
| sudocat | Wrapper for <i>cat</i> command that works only under /var/log or /local/logs |
| sudogrep | Wrapper for <i>grep</i> command that works only under /var/log or /local/logs |
| sudohead | Wrapper for ‘head’ command that works only under /var/log or /local/logs |
| sudols | Wrapper for ‘ls’ command that works only under /var/log or /local/logs |
| sudotail | Wrapper for ‘tail’ command that works only under /var/log or /local/logs |
| sudozgrep | Wrapper for ‘zgrep’ command that works only under /var/log or /local/logs |
| sudozcat | Wrapper for ‘zcat’ command that works only under /var/log or /local/logs |
| svrestart | Restarts the entered service. Run the command as svrestart?args=<servicename> |
| svstatus | Prints the status of the entered service, run as svstatus?args=<servicename> |
| switchinfo | Get information about the cluster switches. |

| Endpoint | Description |
|---------------------------------|---|
| switch_namenode | <ul style="list-style-type: none"> Manually fail over name node from primary or secondary Should be run on the orchestrator.service.consul host only Run while recommission or decommission of name node 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 name node hosts Run endpoint as switch_secondarynamenode?args=--start |
| 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 decommission or decommission of resource manager hosts Run endpoint as switch_yarn?args=--start |
| tail | Wrapper command for <i>tail</i> Unix 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 launcher host hosts |

| Endpoint | Description |
|--------------------------|---|
| 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 | Prints the haproxy stats and status for external haproxy |
| uptime | Wrapper command for <i>uptime</i> Unix command |
| userapps_kill | <ul style="list-style-type: none"> • Kills all the running user application • Should be run on the launcherhost hosts only |
| vgdisplay | Wrapper command for <i>vgdisplay</i> Unix command |
| vgs | Wrapper command for <i>vgs</i> Unix 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 shut down the VM • Should be run on the Baremetal hosts only • Run endpoint as vmshutdown?args=<vmname> |

| Endpoint | Description |
|-----------------|---|
| 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 shut down 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 |
| 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 | Wrapper command for <i>zcat</i> Unix command |
| zgrep | Wrapper command for <i>zgrep</i> Unix command |

Server Maintenance

Server maintenance involves replacement of any faulty server components such as hard disk, memory, or replacing the server.



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.

To perform all the steps involved in server maintenance, from the navigation pane, choose **Troubleshoot > Cluster Status**. 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 the Cisco Secure Workload rack.

Figure 53: Server Maintenance

Model: 8RU-PROD

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

Orchestrator State: IDLE

Displaying 6 nodes (0 selected)

| <input type="checkbox"/> | State T1 | Status T1 | Switch Port ↑ | Serial T1 | Uptime T1 | |
|--------------------------|--------------|-----------|---------------|-------------|---------------------|---|
| <input type="checkbox"/> | Commissioned | Active | Ethernet1/1 | FCH2206V1NF | 2mo 27d 18h 25m 47s | <div> <div> + Commission + Decommission + Reimage + Firmware upgrade + Power off + Reboot </div> <div> <div>Switch Port: Ethernet1/2</div> </div> </div> |
| <input type="checkbox"/> | Commissioned | Active | Ethernet1/2 | FCH2206V1ZF | 2mo 27d 18h 24m 52s | |

Serial: FCH2206V1ZF

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.0
- 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

- collectorDatamover-6
- datanode-6
- druidHistoricalBroker-4
- enforcementCoordinator-3
- orchestrator-2
- redis-1
- secondaryNameNode-1

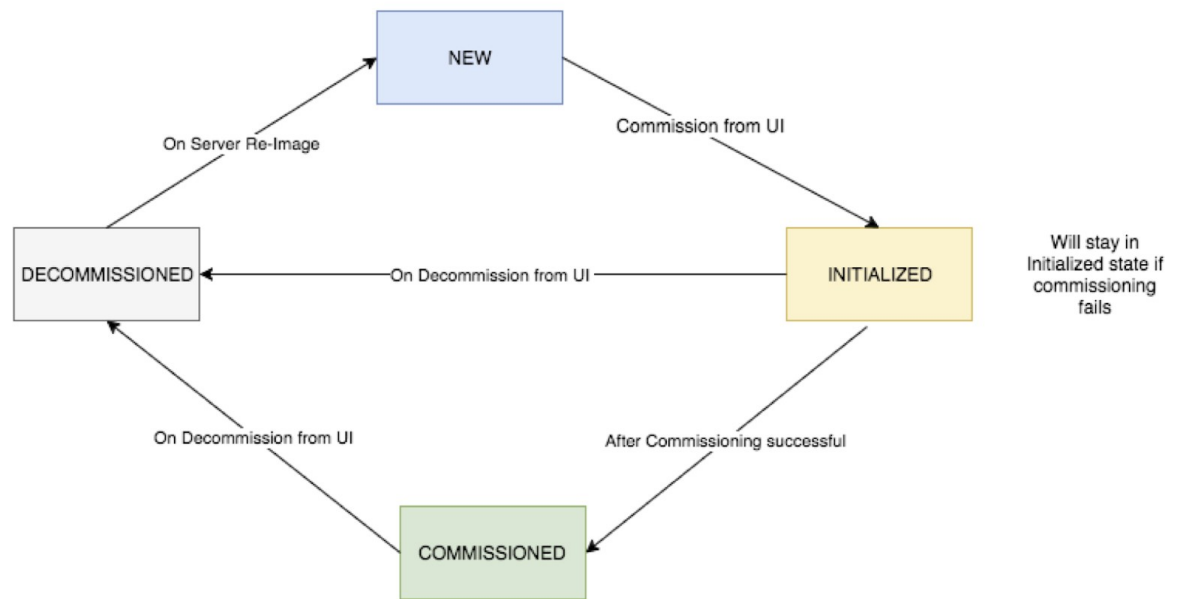
Disks Status

- 252:1 HEALTHY
- 252:2 HEALTHY
- 252:3 HEALTHY
- 252:4 HEALTHY
- 252:5 HEALTHY
- 252:6 HEALTHY
- 252:7 HEALTHY
- 252:8 HEALTHY

| | | | | | | |
|--------------------------|--------------|--------|-------------|-------------|---------------------|-----|
| <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 | + ↓ |

Figure 54: Server State Transition Diagram

Server State Transition Diagram



Steps involved in server or component replacement

- **Determine the server that requires maintenance:** This can be done using the server *Serial* number or the *switch port* to which 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 boxes find out the VMs or instances present on the server and check if any special actions must be carried out for those VMs. The next section lists out Actions for VMs during server maintenance.
- **Decommission the server:** when 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, you can

still perform all the server maintenance steps. Decommission steps can be performed even if the server is powered off.

Figure 55: Decommission the Server

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

Select action ▾
Apply
Clear

| <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:

- 1. 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 the same as that of the replaced server. The CIMC IP for each server is available on the **Cluster Status** page
- 2. 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.
- 3. 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 software version is different from the version of the cluster, then reimage the server.

Figure 56: Replacing Server

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

Select action Apply Clear

| <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

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 [A](#)
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
- happobat-2
- hbaseRegionServer-2
- orchestrator-3
- resourceManager-2
- zookeeper-1

Switch Port: Ethernet1/3

4. **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 provisions the VMs on the server. Commissioning of a server takes about 45 mins. The server will be marked *Commissioned* after commissioning completes.

Figure 57: Commission the Server

Displaying 6 nodes (0 selected)

Select action Apply Clear

| <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 |

Serial: FCH2048V2VY

Private IP: 1.1.1.4
CIMC IP: 172.26.230.178
Status: Active
State: Initialized
SW Version: 2.3.1.24.devel
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
- hbaseRegionServer-2
- orchestrator-3
- resourceManager-2
- zookeeper-1

Switch Port: Ethernet1/3

| | | | | | |
|--------------------------|--------------|--------|-------------|-------------|----------------|
| <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.

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


Figure 58: Maintenance Explorer

If you try to decommission a server with a primary orchestrator, the following error is displayed.

Figure 59: Decommission a Server with Primary Orchestrator Error

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

- **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.
- **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.
- **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.
- **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 60: Server Maintenance: Datanode

Troubleshooting server maintenance

- **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 61: Server Maintenance Log

POST : orchestrator-3 cat?lang=/local/logs/181866p/orchestrator/orchestrator.log

+ Add HTTP Header

Body

POST/PUT body to send

Status: 200

```

2017-04-07_17:27:17.07953 "inst_deployed": {
2017-04-07_17:27:17.07954   "archestrator-2",
2017-04-07_17:27:17.07954   "servernode-1",
2017-04-07_17:27:17.07954   "datanode-5",
2017-04-07_17:27:17.07955   "baseRegionServer-1",
2017-04-07_17:27:17.07955   "collectorDatabase-2",
2017-04-07_17:27:17.07955   "druidHistoricalBroker-2",
2017-04-07_17:27:17.07955   "mongodb-2",
2017-04-07_17:27:17.07955   "enforcementCoordinator-1",
2017-04-07_17:27:17.07955   "hadoop-1"
2017-04-07_17:27:17.07955 },
2017-04-07_17:27:17.07955 "network": {
2017-04-07_17:27:17.07955   "private_ip": "1.1.1.0"
2017-04-07_17:27:17.07955 },
2017-04-07_17:27:17.07955 "state": "Commissioned",
  
```

• Decommission

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



Note

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

- 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 that this server does not join back the cluster in decommissioned state. It must be reimaged and added back to the cluster.

• Reimage

- This step installs the Secure Workload base OS or Hypervisor OS on the server.
- It also formats the hard drives and installs few Secure Workload libraries on the server.
- Reimage runs a script called **mjolnir** to initiate the server imaging. mjolnir 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 information about the reimage, like `/var/log/nginx` logs during pxe boot up, `/var/log/messages` to check for DHCP IP and pxe boot configs.

- 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 in case the server is replaced and set the cimc password to the default password
- 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.

- **Commission**

- Commissioning schedules of the VMs on the server and runs playbooks in the VMs to install Secure Workload software.
- Commissioning requires about 45 minutes to complete.
- The workflow is similar to deploy or upgrade.
- The logs indicate any failures during commissioning.
- The server on the cluster status page will be initialized during commissioning and marked commissioned only after you complete the steps.

Bare Metal 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 the user to reboot (upgrade) with a bad hardware, after which the regular RMA process for the failed bare metal can be performed.

The user is expected to use a post to explore the endpoint with serial of the bare metal to be excluded:

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

The orchestrator performs a few checks to determine if the exclusion is feasible. In which case, it will setup a few consul keys and return a success message indicating which bare metal and VMs will be excluded in the next reboot/upgrade workflow. If the bare metals include certain VMs, they cannot be excluded as described in the Limitation section below, the explore endpoint replies back with the message indicating why the exclusion is not possible. After successful post on the explore endpoint, the user can initiate reboot/upgrade through the main GUI 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 or reboot again with exclude BMs, users are expected to post to the bmexclude explore endpoint again.

Limitations

The following VMs cannot be excluded:

- namenode

- secondaryNamenode
- mongodb
- mongodArbiter

Disk Maintenance

Disk Maintenance involves replacement of any faulty hard disks from one or more servers. The orchestrator monitors the health of the disks as reported by bmmgr on every server in the cluster. If there are any faulty disks, a banner displays the error on the **Cluster Status** page. From the navigation pane, choose **Troubleshoot > Cluster Status**.

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

Figure 62: Faulty Disk Banner

The screenshot shows the Cisco Tetration interface for the CLUSTER STATUS page. At the top, there is a navigation bar with the Cisco Tetration logo and the text 'CLUSTER STATUS'. Below this, 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 displays '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 lists 6 nodes, all with a 'Commissioned' state and 'Active' status. The table columns are: State, Status, Switch Port, Serial, Uptime, and CIMC Snapshots. The nodes are listed with their respective switch ports (Ethernet1/1 to Ethernet1/6) and serial numbers (FCH2148V1EU to FCH2148V1NE). The uptime for each node is shown as '16d 11h 22m 40s' or similar. The CIMC Snapshots column contains icons for viewing 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 | |

Requirement Prechecks

Before performing the decommission or commission of disks, various checks are performed at the backend. All checks must pass before you can proceed with the decommission or commission of the disks.

Failed checks are reported on the **Disk Replacement Wizard** with the failure details and corrective action, which must be taken care before you proceed to the next step, for example, only one datanode can be decommissioned at a time. Namenode and secondaryNamenode cannot be decommissioned together; also, check if Namenode is healthy before commissioning the disk.

Figure 63: Disk Replacement Prechecks

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 | druid-HistoricalBroker-4 |
| FCH2148V1N9 | 252:7 | UNHEALTHY | datanode-6 |

Prechecks

Start Prechecks

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

Decommission

Start Decommission

You can select any set of failed disks to be decommissioned together and start the decommission prechecks. Changing the set of failed disk requires a rerun of the prechecks. Check the prechecks again before you start decommission or commissioning of the disks. Ensure that there are no new precheck failures between the last precheck run and the start of the decommission task.

Figure 64: Unhealthy Disks for Decommission

The screenshot shows the Cisco Tetration interface for 'CLUSTER STATUS - DISK REPLACEMENT'. The workflow progress bar indicates four steps: Prerequisites (1, green), Decommission Drives (2, blue), Replace Drives (3, grey), and Commission Drives (4, grey). The 'Decommissioning Unhealthy Drives' section contains the following instructions:

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.

The 'Select Disks' section shows a dropdown menu with the text 'Select unhealthy disks for decommission'. Below the dropdown, a table lists the selected disks:

| Serial | Enclosure:Slot | Status | Affected VMs |
|-------------|----------------|-----------|-------------------------|
| FCH2148V1EP | 252:3 | UNHEALTHY | druidHistoricalBroker-4 |

The 'Prechecks' section has a 'Start Prechecks' button and a note: 'Prechecks should be run successfully to proceed with decommission.'

The 'Decommission' section has a 'Start Decommission' button.

If a precheck fails, a detailed message is displayed. Click on the failure message and a suggested action will be shown in a pop-over when the pointer hovers over the cross button.

Figure 65: Suggested Action for Failed Precheck

The screenshot shows the Cisco Tetration interface for 'CLUSTER STATUS - DISK REPLACEMENT'. The 'Selected 1 disk' section shows a table with the following data:

| Serial | Enclosure:Slot | Status | Affected VMs |
|-------------|----------------|-----------|---|
| FCH2148V1NG | 252:1 | UNHEALTHY | resourceManager-2, orchestrator-3, hbaseRegionServer-2, enforcementPolicyStore-3, datanode-1, appServer-1, redis-2, zookeeper-1, collectorDataMover-3 |

The 'Prechecks' section has a 'Start Prechecks' button and a message: 'Prechecks failed at May 6 11:24:52 am (PDT). Please find details below.'

The 'Decommission' section has a 'Start Decommission' button.

A pop-over titled 'Action Required' is displayed, showing the message: 'Please check if any disk is missing from the list of disks to be decommissioned.'

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

Figure 66: Select UNHEALTHY Disks to Decommission

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

- ✓ FCH2148V1EP | 252:3 | druidHistoricalBroker-4
- ✓ FCH2148V1N9 | 252:1 | druidCoordinator-1, orchestrator-2, enforcementPolicyStore-1, enforcementCoordinator-3, redis-1, sec...
- ✓ FCH2148V1N9 | 252:7 | datanode-6

| Host | Device | Status | Component |
|-------------|--------|-----------|-------------------------|
| FCH2148V1EP | 252:3 | UNHEALTHY | druidHistoricalBroker-4 |

Prechecks

Start Prechecks

Prechecks should be run successfully to proceed with decommission.

Decommission

Start Decommission

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

Figure 67: Suggested Action in Popover for Failed Precheck

The screenshot shows the Cisco Tetration interface for 'CLUSTER STATUS - DISK REPLACEMENT'. At the top, there's a dropdown menu labeled 'Select unhealthy disks for decommission'. Below this, a table shows 'Selected 1 disk' with the following details:

| Serial | Enclosure:Slot | Status | Affected VMs |
|-------------|----------------|-----------|--|
| FCH2148V1NG | 252:1 | UNHEALTHY | resourceManager-2, orchestrator-3, hbaseRegionServer-2, enforcementPolicyStore-3, datanode-1, appServer-1, redis-2, zookeeper-1, collectorData mover-3 |

Below the table, there's a 'Prechecks' section with a 'Start Prechecks' button. A message indicates that prechecks failed at May 6 11:24:52 am (PDT). A yellow bar highlights the failed precheck: 'check_disk_ready_for_decomm'. A popover titled 'Action Required' is shown, stating: 'Please check if any disk is missing from the list of disks to be decommissioned.' Below this, there's a 'Decommission' section with a 'Start Decommission' button. At the bottom right, there are 'Previous' and 'Next' navigation buttons.

Disk Replacement Wizard-Not Hot Swap

Before you Begin

Before you start the replacement process for unhealthy disks, ensure that the new disks are available.

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

Figure 68: Disk Replacement Wizard

Drive Replacement Process

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

Before you begin

1. Keep the **replacement disks** with following configuration in hand.
 - 2 disks of type 1.454 TB SSD INTEL SSDSC2BB016T7K
 - 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](#)



Note Physically, the drives and hardware support hot swap.

Disk Status Transitions

In any cluster, for non-RAID, there are six states for the hard disks—**HEALTHY**, **UNHEALTHY**, **UNUSED**, **REPLACED**, **NEW**, and **INITIALIZED**. After you deploy or upgrade the cluster, the status of every disk in the cluster is **HEALTHY**. The status of one or more disks can change to **UNHEALTHY** based on various error detections.



Note Non-hot swappable drives are available only for M4 and M5 clusters.

No action is taken unless the status of a disk changes to **UNHEALTHY**. Before you start commissioning the disks, deploy all VMs that were removed as part of the decommissioning process.

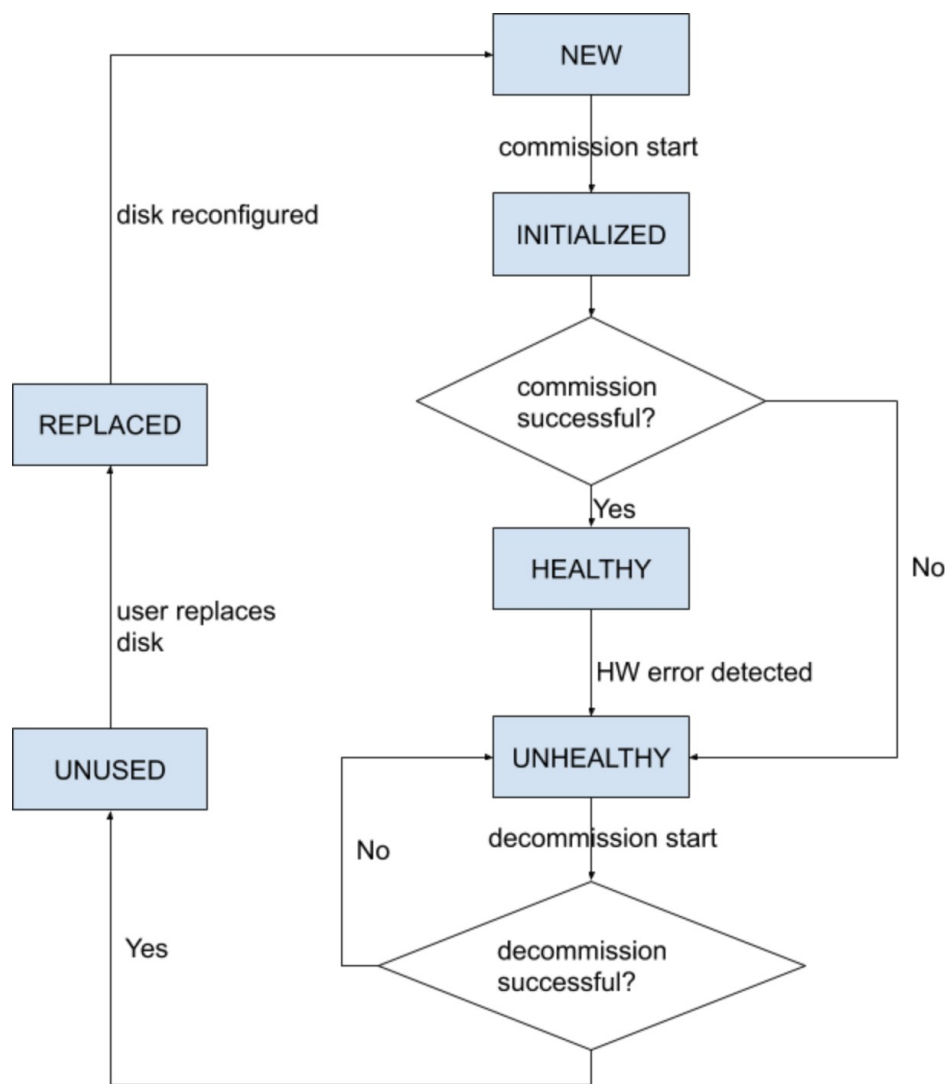
After you have successfully commissioned the disks without any errors, the status of the disks changes to **HEALTHY**. In case the disk commission is unsuccessful, the status displays as **UNHEALTHY**. For the

disks that are **UNHEALTHY**, start the process of disk decommission. If the decommissioning process is successful, the status of the disk changes to **UNUSED**, and if the disks fail during decommission, repeat the process until the status of the disks changes to **UNUSED**.

Remove the **UNHEALTHY** disks from the cluster and replace them with new disks, the status changes to **REPLACED**. Reconfigure the replacement disks and scan the hardware for any anomalies. In case there are no anomalies detected, the status of the disks changes to **NEW**, else, you may need to troubleshoot the issue; status transition can take upto three minutes.

To understand how disk status transitions are handled, see the flow diagram below:

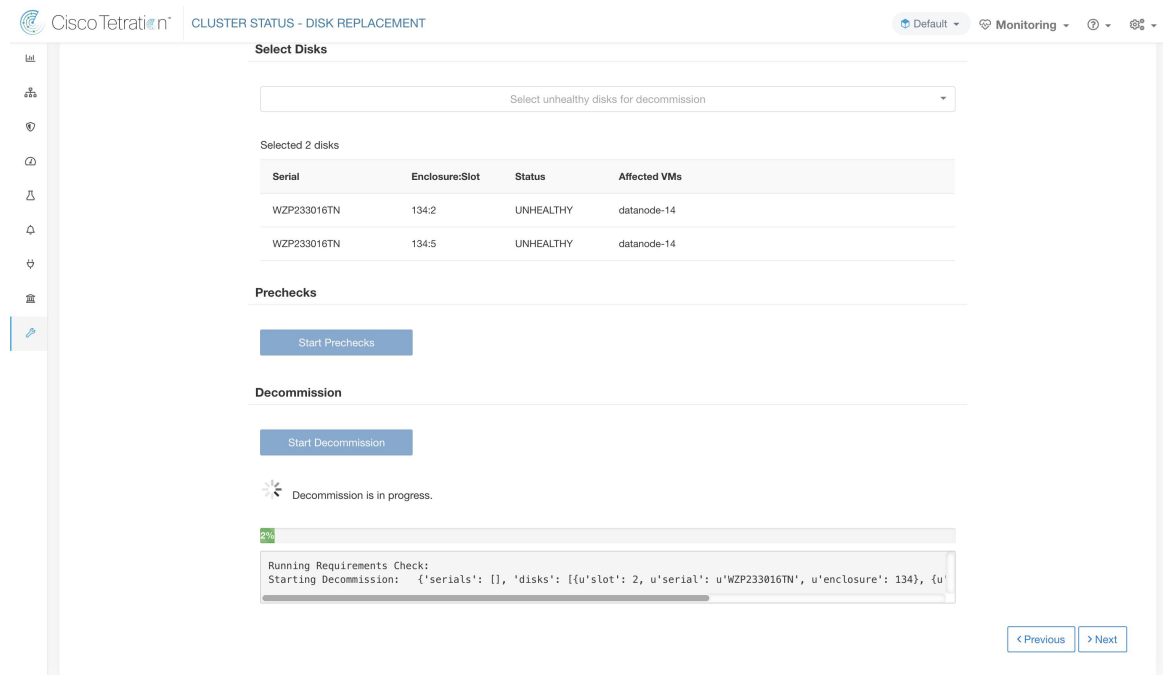
Figure 69: Disk Status Transitions



Decommission Disk

After the prechecks passes, you can proceed to decommission the disk. The progress of decommission will be shown on the disk replacement wizard. When the progress of decommission reaches 100%, all the decommissioned disks' status changes to UNUSED.

Figure 70: Monitor Disk Decommission Progress



Replace Disk

After decommissioning the disks, remove the disks and replace with new disks. To help with this process, we have added disk and server locator LED access on the replace page. Ensure you switch off the server and disks locator LEDs.

Figure 71: Reconfigure Newly Added Disks (Not Hot Swappable)

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 | | |

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 on the UI and the backend. On the UI, you will have a replace button active for disk with the lowest slot number with the status UNUSED.

Commission Disk

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

Cisco Tetration CLUSTER STATUS - DISK REPLACEMENT

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.

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 were successful at May 4 11:21:14 pm (PDT).

Commission

Start Commission

Previous

Progress of commission is monitored on the disk commission page. At the end successful commission, the status of 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](#)

Failure Recovery During Disk Commission

After you deploy the VMs and there is a failure, you can recover using the **Resume Commission** button. To continue commission of disk, click the **Resume Commission** button to restart 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, Run
```

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

Disk Failure 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 the 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 is the only path available in cases of nonresumable failures. If the post-deploy failure is caused due to the newly failed disks, the second path will again be only way forward, although we have resume button available.

Known Issues and Troubleshooting

- Disk containing server root volumes cannot be replaced using this procedure. Such disk failures must be corrected using server maintenance process.

- Disk commissioning can happen only when all servers are active and in commissioned state. See *Special Handling* section that describes how to proceed in the cases where a combination of disk and server replacement is needed.
- SSD disks are too expensive and have a very low failure rate and therefore we do not want to lose valuable capacity for redundant data storage.
- On M6 clusters originally deployed with 3.8 software, when a server is commissioned with 3.9 software, RAID configuration will be applied on the HDD drives. This will cause a cluster to contain some nodes with RAID and some with the non-RAID disk configuration from 3.8. It is likely that your Secure Workload 39RU hardware originally shipped with 3.9 already installed, but some early M6s were shipped with 3.8 deployed.
- You can convert a cluster to RAID if server decommission and commission is performed progressively across all servers after upgrading to 3.9 software.
- M6 8RU clusters are all-SSD nodes and RAID is not configured on SSD drives and therefore 8RUs do not get RAID.
- Drive configuration on older generations (M4/M5) prevents us from supporting RAID on those generations of Secure Workload hardware.

Disk and Server Replacements

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*

Cisco Tetration[™] CLUSTER STATUS - DISK REPLACEMENT

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 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.

Cisco Tetration[™] CLUSTER STATUS

Model: 39RU-M5

Orchestrator State: IDLE

There are 3 unhealthy disks in the appliance. You can replace them. Please check [here](#)

Displaying 1 nodes (1 selected)

| State | Status | Switch Port | Serial | Uptime | CIMC Snapshots |
|-------|--------|--------------|-------------|--------------|----------------|
| New | Active | Ethernet1/12 | WZP232913LX | 6d 2h 2m 35s | 0 |

Commission aborted: Disks ['WZP233016TN']-[134:4] Status(UNHEALTHY)', ['WZP233016TN']-[134:2] Status(UNHEALTHY)] status is not ['NEW']. Please complete replace task in disk wizard

TetrationOS Software, Version 3.5.2.69949.ravi.pra.mrgm.build
Privacy and Terms of Use
TAC Support: <http://www.cisco.com/tac>
© 2015-2020 Cisco Systems, Inc. All rights reserved.

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 74: 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.

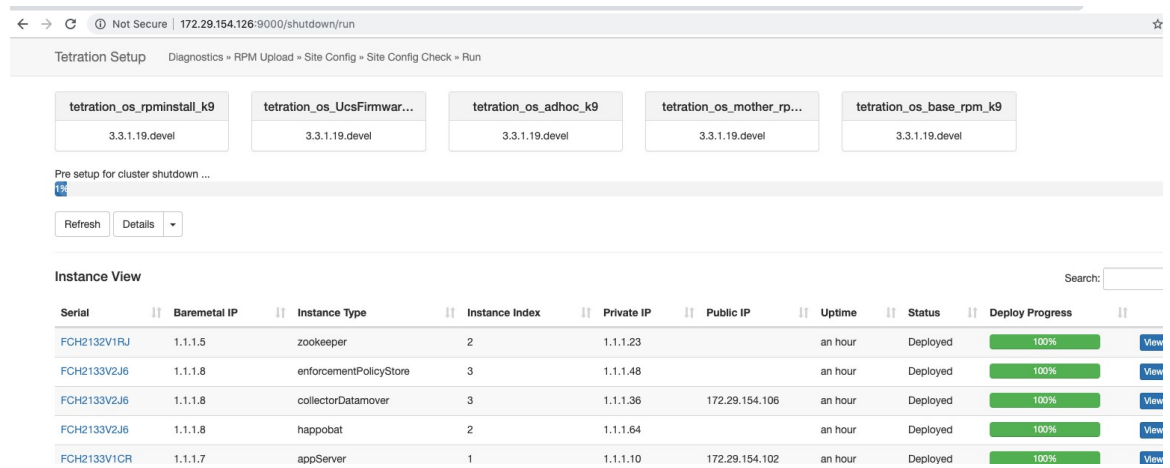
Shutdown will not be triggered until you actually click the above link.

- 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.

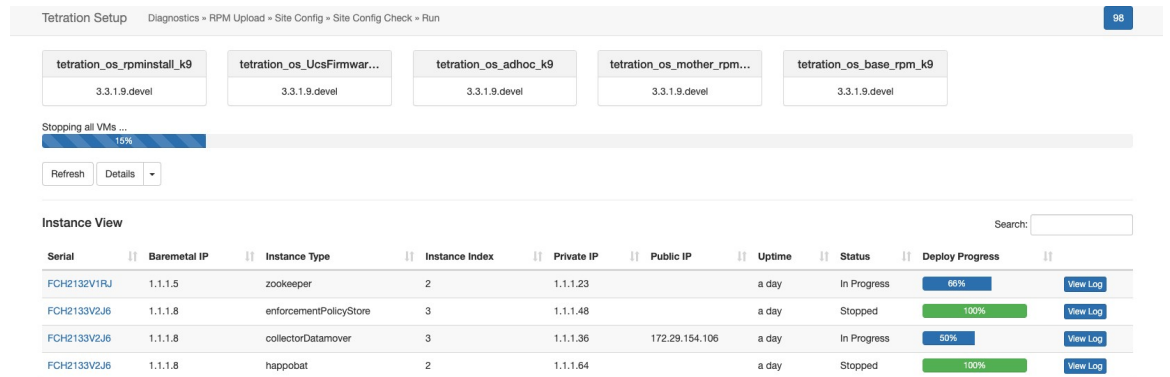
Figure 75: Cluster Shutdown Progress



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 76: Stopping VMs

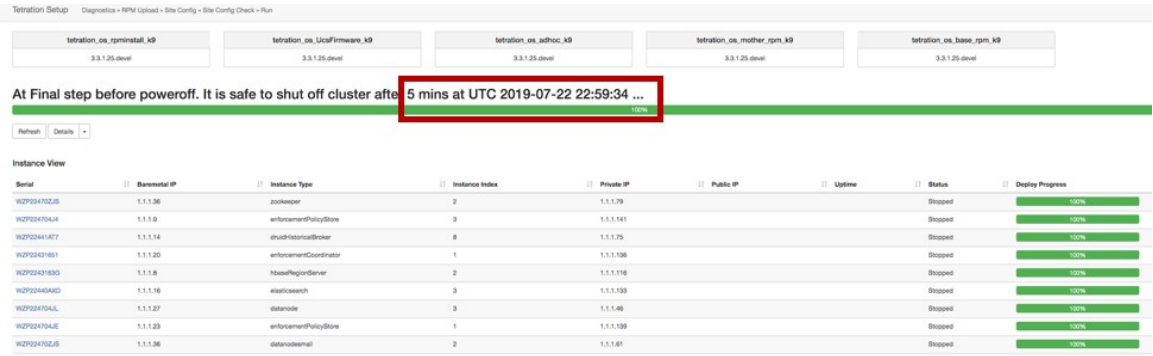


When the cluster is ready to be shut down, the progress bar will go to a 100% and indicate the time after which it is safe to power off the cluster. See the highlighted in the following screenshot.



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

Figure 77: 100 Percent Shutdown



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

Data Taps



Note Secure Workload supports writing to Kafka Brokers 0.9.x, 0.10.x, 1.0.x and 1.1.x for data taps.

To send alerts from the Secure Workload cluster, you must use a configured data tap. Data Tap Admin users can configure and activate new or existing data taps. You can view data taps of your **Tenant**.

Figure 78: Available Data Taps

Data Tap Admin - Data Taps

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

[+ New Data Tap](#)

To manage data taps, in the navigation pane, choose **Manage > Data Tap Admin**.

Recommended Kafka Configuration

While configuring the Kafka cluster, we recommend using the ports from 9092, 9093, or 9094 because Secure Workload opens these ports 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 view the available data taps and configure them by navigating to **Manage > Data Tap Admin > Data Taps**. The data taps are configured per **Tenant**.

Figure 79: All 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 | |
| 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

+ New Data Tap

Data Tap Admins can click the to add a new data tap.

Figure 80: Adding New Data Tap

New Data Tap

Name

Name of Data Tap

Description

Description of the Data Tap

Kafka Broker

IP/Hostname(s). Ex: kafka1.ci

Enter Topic Name here

Topic

default -- Kafka Topic for

Cancel

Test Settings







Note Changing data tap values require settings to be validated.

Deactivating a Data Tap

To temporarily prevent outgoing messages from 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 be reactivated at any time.

Figure 81: Deactivating a Data Tap

Data Tap Admin - Data Taps

| 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 |    |

Note: A red arrow points to the deactivate icon in the Actions column of the first row, with the text "Click here to deactivate".

Deleting a Data Tap

Deleting a data tap deletes 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 secure in terms of authentication, encryption, and authorization. To send and receive messages from MDTs, clients must be authenticated, and data that is sent over the wire is encrypted, and only authorized users can read or write messages from or to Secure Workload MDT. Secure Workload provides client certificates to be downloaded from the GUI. Secure Workload uses Apache Kafka 1.1.0 as the messages broker, and we recommend that the clients use secure clients compatible with the same version.

MDTs are automatically created after the creation of root scope. Every root scope has an Alerts MDT created. To retrieve alerts from the Secure Workload cluster, you must use the Alerts MDT. Only Data Tap Admin users can download the certificates. You can view MDTs of your **root scope**.

Figure 82: List of Configured Data Taps

Data Tap Admin - 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 alerts are sent to MDT by default, but can be changed to other Data Taps.

There are two choices for downloading the certificates:

- Java KeyStore: JKS format works well with Java Client.
- Certificate: Regular certificates are easier to use with Go Clients.


Figure 83: Download Certificates

Data Tap Admin - Data Taps

| Name ↑ | Topic ↑ | Description ↑ | Kafka Broker ↑ | Type ↑ | Status ↑ | Actions ↑ |
|------------|--|------------------------------|--------------------------------------|----------|----------|---|
| Alerts | topic-610881bf497d4f7bd287a224 | DataTap Managed by Tetration | 172.21.156.186:443 | Internal | Active |     |
| 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 |    |

Note: A tooltip "Download Client Certificate" is shown over the download icon in the Actions column of the Alerts row.

Figure 84: Certificate Types



Java Keystore

After downloading *alerts.jks.tar.gz*, you should see the following files that contain information to connect to Secure Workload MDT to receive messages:

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

The following 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>
```

While setting up the Kafka Consumer in Java code, use the following properties:

```
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 you want to use certificates, use Go clients using the Sarama Kafka library to connect to Secure Workload MDT. After downloading *alerts.cert.tar.gz*, you should see the following files:

- *kafkaBrokerIps.txt*: This file contains the IP address string that the Kafka Client uses to connect to Secure Workload MDT
- *topic*: This file contains the topic that this client can read the messages from. Topics are of the format *topic<root_scope_id>*. Use this *root_scope_id* while setting up other properties in Java Client.
- *KafkaConsumerCA.cert*: This file contains the Kafka Consumer certificate.
- *KafkaConsumerPrivateKey.key*: This file contains the private key for the Kafka Consumer.
- *KafkaCA.cert*: This file should be used in the root CA certs listing in the Go client.

To view an example of a Go client connecting to Secure Workload MDT, see [Sample Go Client to consume alerts from MDT](#).