

Verify the Health of a Secure Workload Cluster (Tetration)

Contents

[Introduction](#)

[Background Information](#)

[When to Check the Health of the Cluster](#)

[Different Options You have to Verify the Health of the Secure Workload Cluster](#)

[Cluster Status](#)

[Service Status](#)

[Hawkeye \(Charts\)](#)

[Upgrade Prechecks](#)

Introduction

This document describes steps to verify the health of a Secure Workload cluster and highlights key aspects to review during the health check process.

Background Information

Its primary focus is on health verification; however, if you notice any issues or abnormal behavior, you must collect a snapshot and contact the Cisco Tetration Solution Support TAC team for assistance. The secure workload cluster is made up of hundreds of processes distributed across multiple virtual machines on several UCS C220 servers.

The two main tools for assessing cluster health are the Cluster Status and Service Status pages, both of which are explained in this document. Using these pages is generally the most effective way to confirm the overall health of the cluster.

When to Check the Health of the Cluster

Most of the time, it is not necessary to verify the health of your cluster. However, there are certain situations when it is a good idea:


- If you notice anything unusual or unexpected in the user interface (UI), based on your experience with how things normally work. Some common examples are listed in the Operational Display Parameters section.
- If you expect to see certain data (like flow data from software or hardware sensors) in the UI, but it is missing even though you have selected the right scope and time range.
- Before and after any scheduled maintenance, upgrades, or major changes to the cluster. It is a best practice to take a snapshot of the cluster's state both before and after these activities. If you ever need to contact TAC support, having these snapshots can help to quickly pinpoint what changed.

Different Options You have to Verify the Health of the Secure

Workload Cluster

Cluster Status

A secure workload cluster consists of either 6 servers (8RU) or 36 servers (39RU), depending on the cluster type. The Cluster Status page provides the state of the servers as well as bare metal server information.




Note: The Cluster Status page is accessible to users with Site Admin or Customer Support roles for physical clusters. Both roles are able to view and perform actions on the Cluster Status page.

From the navigation pane, choose **Troubleshoot > Cluster Status**.

The cluster status shows the status of all the servers in the Cisco Secure Workload rack. A functioning server can display a State of Commissioned and a Status of Active as shown here.

</



Caution: If you notice any node marked as inactive on the cluster status page, generate a CIMC snapshot and raise a TAC case, including the snapshot.

If the status appears as Inactive, it usually means the server is either turned off or can be down due to a hardware, cable, or connectivity issue.

When you click a server in the list, you see more details, such as

- The virtual machines (instances) running on that physical server
- The server's private IP address within the cluster
- The CIMC (management) IP address
- The current firmware versions for the BIOS, CIMC, VIC Card, LOM card, and RAID controller

Cluster Status

Model: 8RU-M6

CIMC/TOR guest password

External Access ☐ Disabled

Orchestrator State: IDLE

Displaying 6 nodes (0 selected)

Select action

Apply

	State	Status	Switch Port	Serial	Uptime	CIMC Snapshots
<input type="checkbox"/>	Commissioned	Active	Ethernet1/3	WMP272900CQ	1y 7mo 7d 18h 49m 3s	+
<div> <div>Serial: WMP272900CQ</div> <div> <div>Private IP: 192.168.1.5</div> <div>CIMC IP: 192.168.0.13</div> <div>Status: Active</div> <div>State: Commissioned</div> <div>SW Version: 3.10.1.1</div> <div>Hardware: 56 cores, 947G memory, 10 disks, 24.27T space, SSD</div> <div>Firmware: View Firmware Upgrade Logs</div> <div> <div>BIOS: C220M6.4.2.3a.01029220536</div> <div>CIMC: 4.21(3b)</div> <div>Cisco UCS VIC 1455 Slot 1: 5.2(3e)</div> <div>Cisco UCS VIC 1455 Slot 3: 5.2(3e)</div> <div>Cisco 12G SAS RAID Controller with 4GB FBWC (16 Drives) Slot MRAID: 52.20.0-4523</div> <div>Intel X550 LOM Slot L: 0x800016FD-1.826.0</div> </div> </div> <div> <div>Instances</div> <div> <div>collectorData mover-3</div> <div>datanode-3</div> <div>druidHistoricalBroker-1</div> <div>elasticsearch-1</div> <div>enforcementPolicyStore-3</div> <div>happobat-1</div> <div>hbaseMaster-1</div> <div>orchestrator-3</div> <div>redis-3</div> <div>tsdbBosunGrafana-1</div> <div>zookeeper-1</div> </div> </div> <div> <div>Disks Status</div> <div> <div>1 HEALTHY</div> <div>2 HEALTHY</div> <div>3 HEALTHY</div> <div>4 HEALTHY</div> <div>5 HEALTHY</div> <div>6 HEALTHY</div> <div>7 HEALTHY</div> <div>8 HEALTHY</div> <div>9 HEALTHY</div> <div>10 HEALTHY</div> </div> </div> <div>Switch Port: Ethernet1/3</div> </div>						

Service Status

The Service Status page is located in the left navigation pane under **Troubleshoot > Service Status**.

The Service Statuspage displays the health of all services that are used in your CiscoSecure Workloadcluster along with their dependencies.

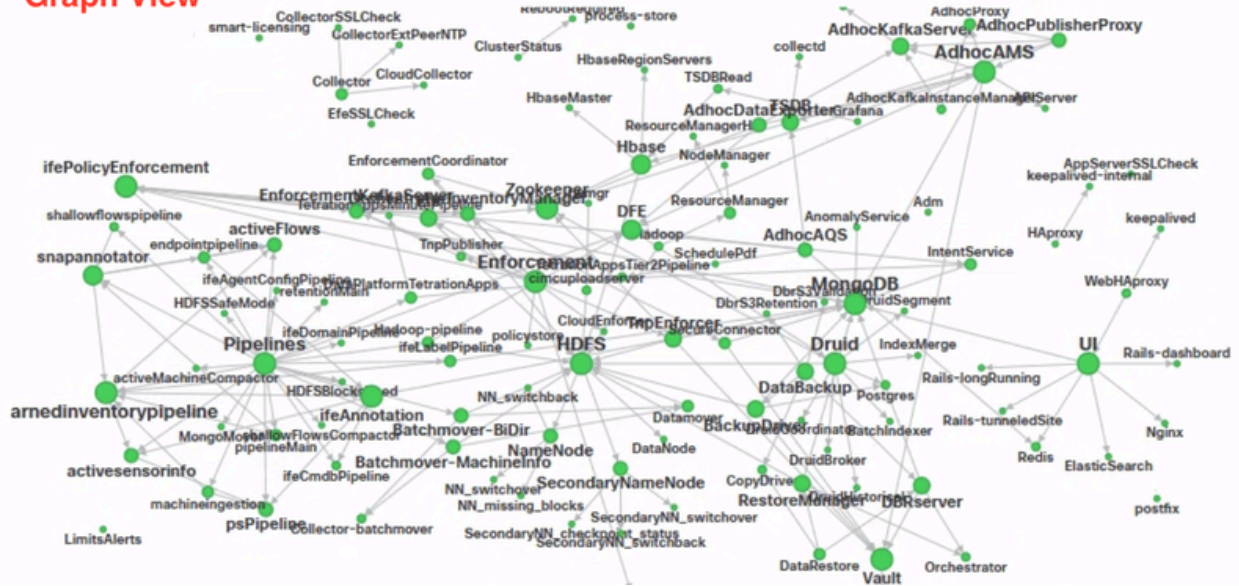
The graph view shows the health of the service, each node in the graph shows the health of the service, and an edge represents dependency on other services. Unhealthy services are marked either red when the service is unavailable, and orange when the service is degraded but available. A green color or sky blue color indicates that the service is healthy. For more debug information on these nodes, use tree view which has theExpand Allbutton 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.


Service Status Graph




Choose a service

Graph View

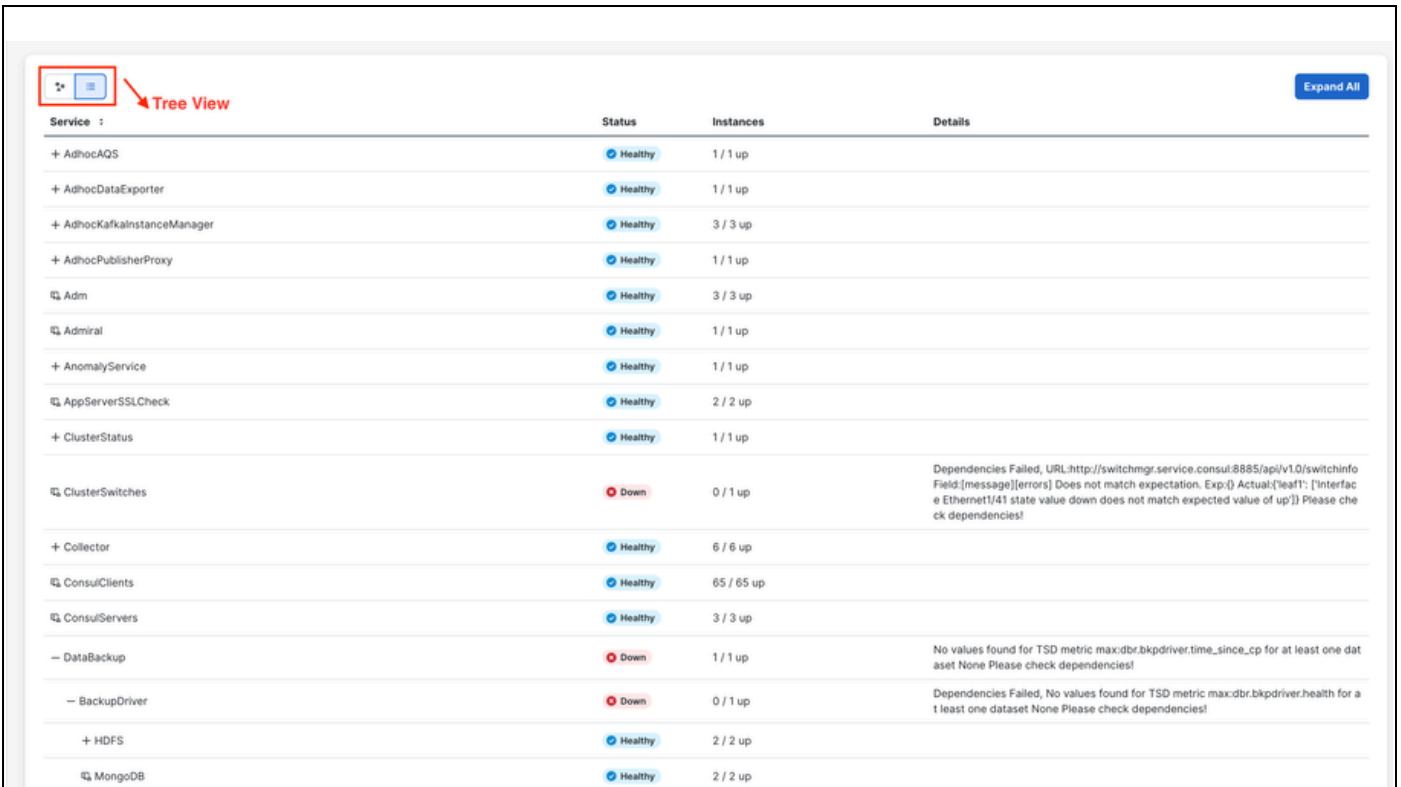




Note: Starting with patch version 3.10.2.11, the service status page appears in sky blue. A green color or sky blue color indicates that the service is healthy.



Caution: If you see that any of the services are unhealthy and showing red in color, contact the Technical Assistance Center (TAC) for support in resolving these issues. Quick engagement with TAC can help restore full functionality.



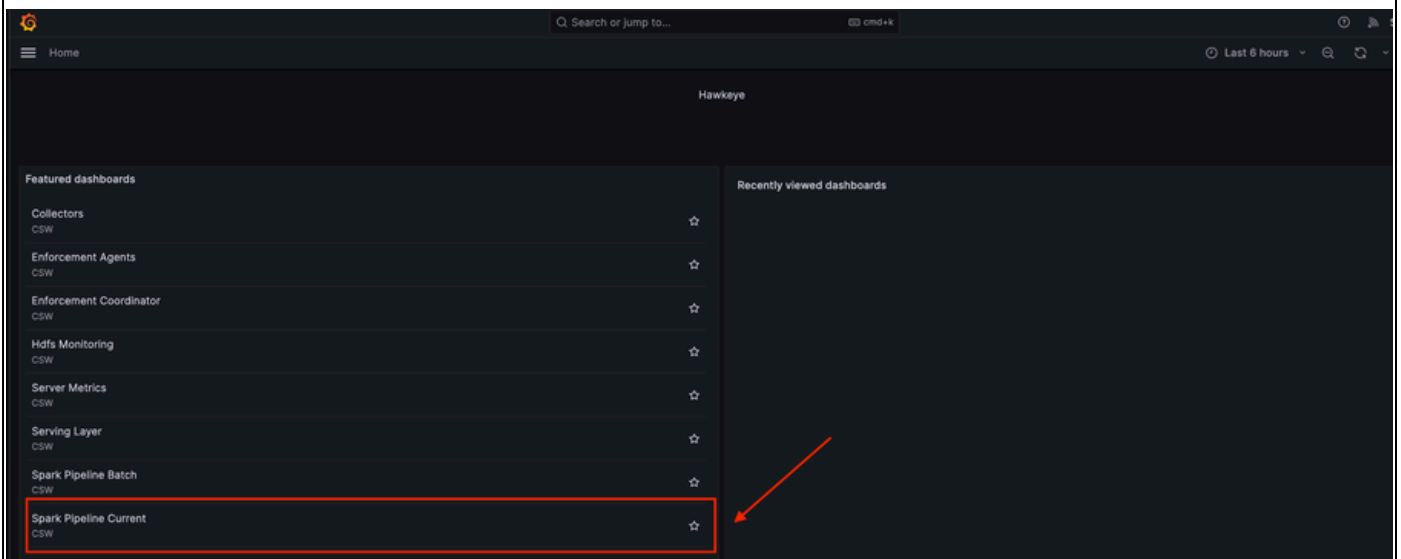
Service	Status	Instances	Details
+ AdhocAQS	Healthy	1 / 1 up	
+ AdhocDataExporter	Healthy	1 / 1 up	
+ AdhocKafkaInstanceManager	Healthy	3 / 3 up	
+ AdhocPublisherProxy	Healthy	1 / 1 up	
Adm	Healthy	3 / 3 up	
Admiral	Healthy	1 / 1 up	
+ AnomalyService	Healthy	1 / 1 up	
AppServerSSLCheck	Healthy	2 / 2 up	
+ ClusterStatus	Healthy	1 / 1 up	
ClusterSwitches	Down	0 / 1 up	Dependencies Failed, URL:http://switchmgr.service.consul:8885/api/v1.0/switchinfo Field:[message][errors] Does not match expectation. Exp:() Actual:{"leaf1": ["Interface Ethernet1/41 state value down does not match expected value of up"]} Please check dependencies!
+ Collector	Healthy	6 / 6 up	
ConsulClients	Healthy	65 / 65 up	
ConsulServers	Healthy	3 / 3 up	
- DataBackup	Down	1 / 1 up	No values found for TSD metric max:dbr.bkpdriever.time_since_cp for at least one dataset None Please check dependencies!
- BackupDriver	Down	0 / 1 up	Dependencies Failed, No values found for TSD metric max:dbr.bkpdriever.health for at least one dataset None Please check dependencies!
+ HDFS	Healthy	2 / 2 up	
MongoDB	Healthy	2 / 2 up	

Hawkeye (Charts)

Hawkeye dashboards offer visibility into the health of the secure workload cluster, as well as metrics and insights to assist with troubleshooting

The Hawkeye (Charts) page is located in the left navigation pane under **Troubleshoot > Hawkeye (Charts)**.

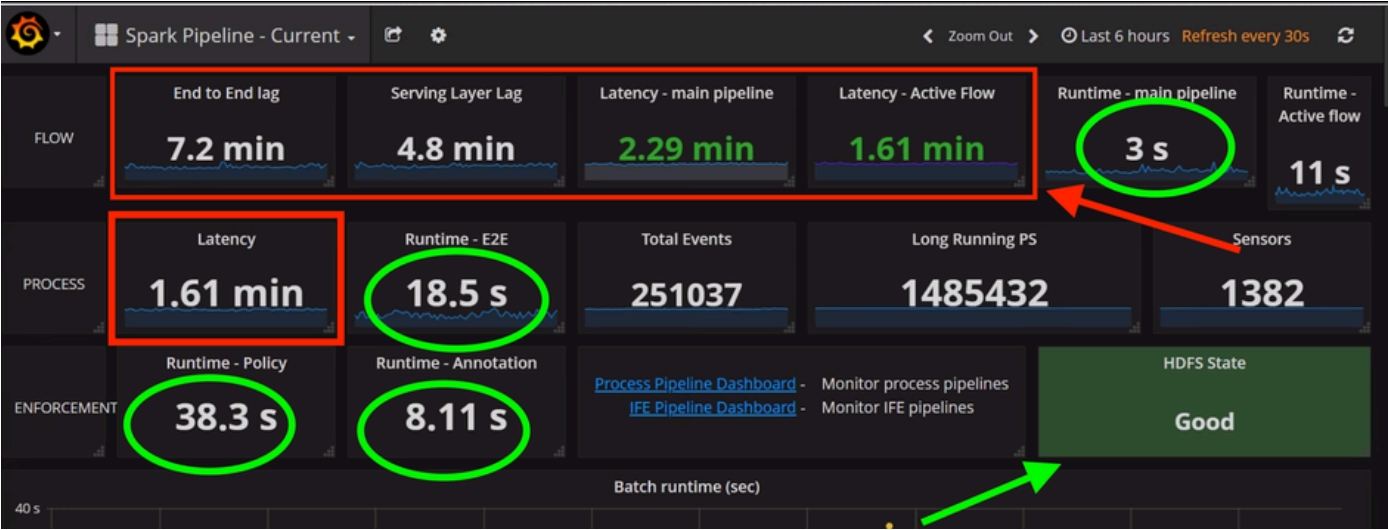
When you click Hawkeye (Charts), a new browser tab automatically opens, displaying the Hawkeye dashboard as shown here.




From the Hawkeye dashboard, click the **Spark Pipeline Current** tab to monitor the secure workload cluster's health.

On the Spark Pipeline Current page, verify that the End-to-End Lag, Serving Layer Lag, Main Pipeline Latency, and Active Flow Latency values are all under 10 minutes.

Also, confirm that the runtime values are less than 1 minute and are presented in seconds and the HDFS state is Good, as illustrated next.





Caution: If you observe any latency values, including end-to-end lag or service layer lag, exceeding 6 hours without showing a gradual decrease, please reach out to the Technical Assistance Center (TAC).

Upgrade Prechecks

Prior to and after maintenance tasks, use the upgrade precheck to run cluster health checks; this process ensures that services, configurations, and hardware components are all in proper working order

1. Navigate to **Upgrade Precheck**.

Navigate to the **Tetration UI** and follow these steps:

- Click **Platform**.
- Select **Upgrade/Reboot/Shutdown**.
- Click **Start Upgrade Precheck**.

Wait a few minutes for the output of the upgrade prechecks. If everything is successful as shown in this image, then you can proceed with the next actions of the cluster maintenance activities.

Secure Workload

Your license usage is out of compliance

Upgrade/Reboot/Shutdown

Upgrade

Reboot/Shutdown

1 Precheck


The last pre-upgrade site check was successful. The configured site admin's email address was validated that pre-upgrade checks passed.

Start Upgrade Precheck

Upgrade Precheck Status

Task	Status	Log
Cluster Health Check	success	Orchestrator
Service Health Check	success	Orchestrator
Secrets Sync Check	success	Orchestrator
Site Linter	success	Orchestrator
Site Checker	success	SiteInfoChecker

Close



Caution: If any upgrade prechecks are unsuccessful, please contact the Technical Assistance

Center (TAC) for assistance.