Quantum

Unified Surveillance Platform (USP)

# User Guide

6-69079-01, Rev D

# Table of Contents

## Contents

# Overview

This document outlines how to use Unified Surveillance Platform (USP). It will guide you through deploying and configuring instances, as well as viewing the overall performance and health of the USP cluster.

# Launching the USP Management Interface

1. Navigate to the USP management interface using the Cluster IP that was configured during installation: https://<cluster_ip>/quantum/usp/dashboard/

NOTE: *Quantum recommends that you always use the latest version of Chrome to access the USP Management Application. The web application runs best at 1920*1080 resolution.*

2. Enter the username and password defined during installation.
3. Click on **Login.**

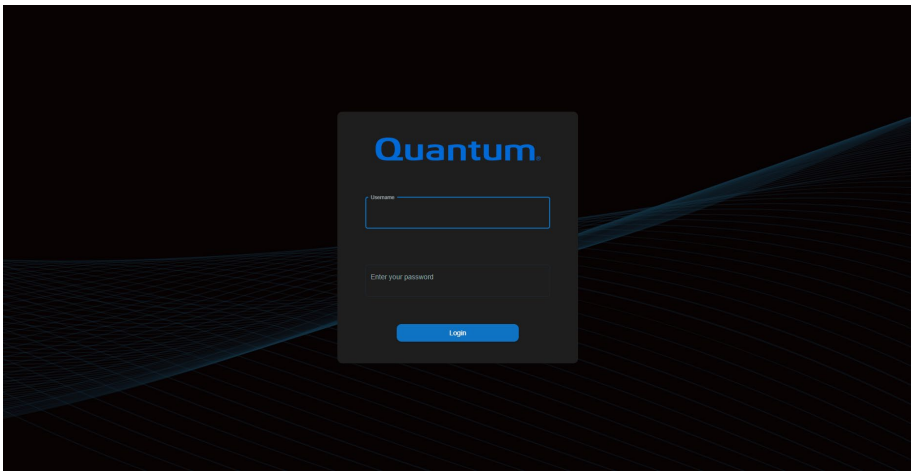


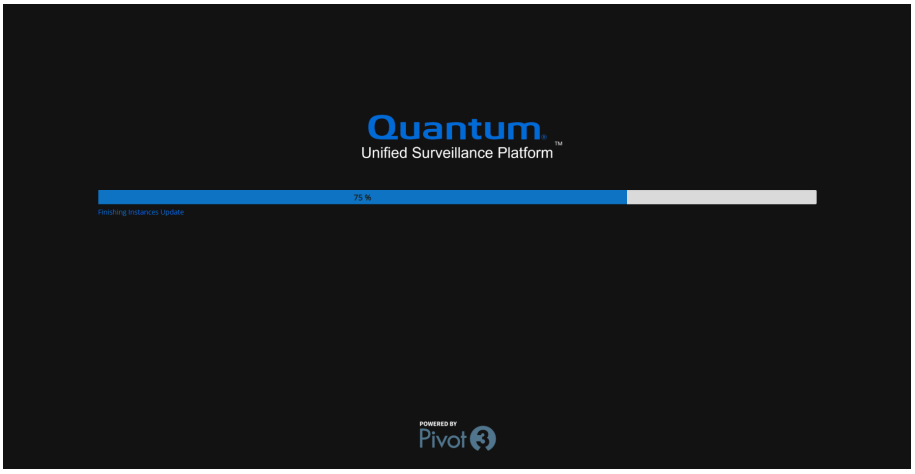

# Dashboard Overview

The following section describes the landing dashboard page. You can navigate back to this page from anywhere in the application by either clicking on the **Quantum** logo in the top-left corner or selecting **Overview** from the navigational menu .

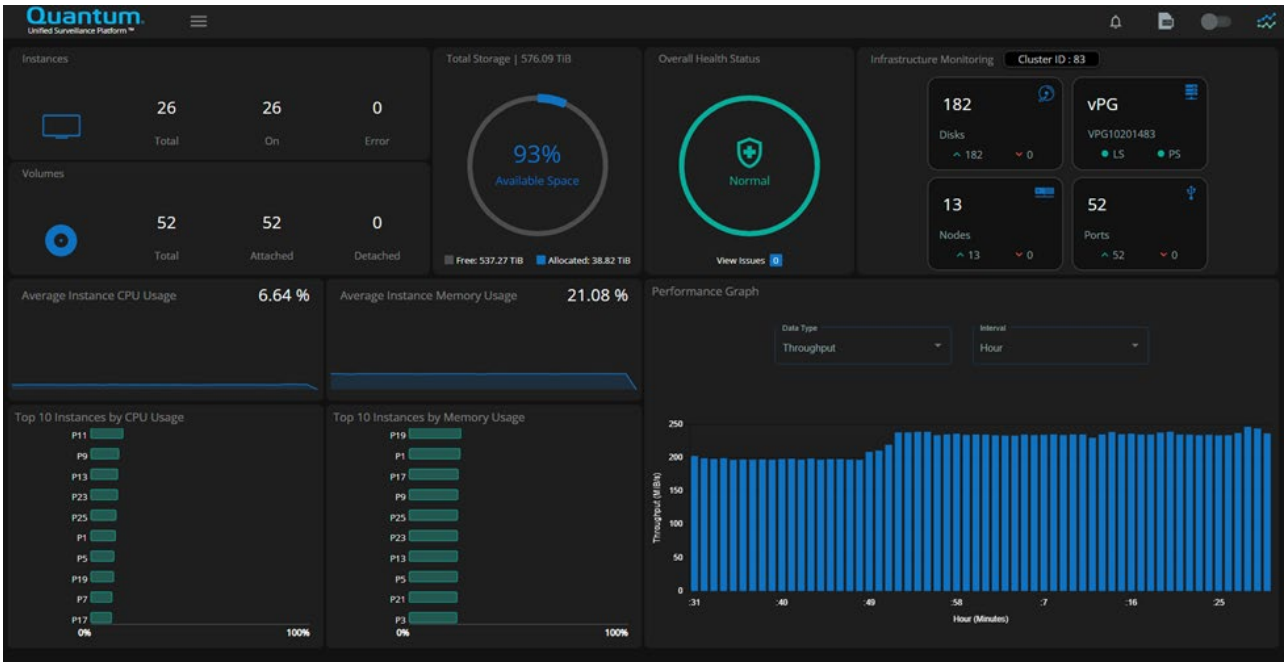


## Instances

This section of the dashboard lets you quickly see the instance count on the USP cluster. The section illustrates the number of instances that are powered on and how many are in an error state. You can click anywhere in this section on the dashboard to navigate to the Instances page.

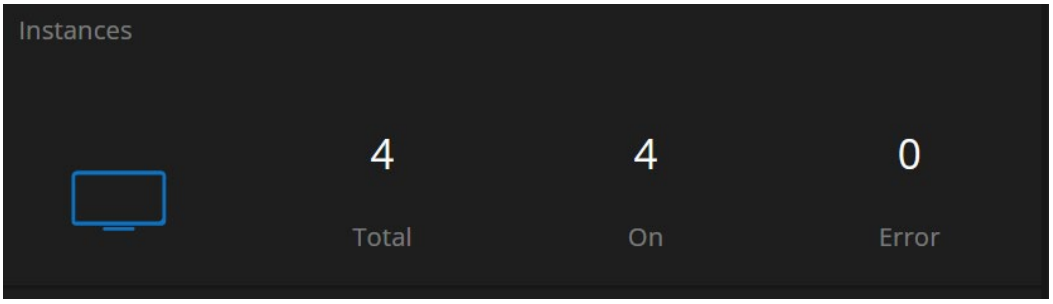


## Volumes

The Volumes section of the dashboard displays the total number of volumes on the cluster. It describes the volumes used by your instances and those that are currently detached from any instance.


## Total Storage

This dashboard section shows the available space on the system. It includes space used by your instances, any volumes attached to those instances, as well as space used to maintain the USP infrastructure.

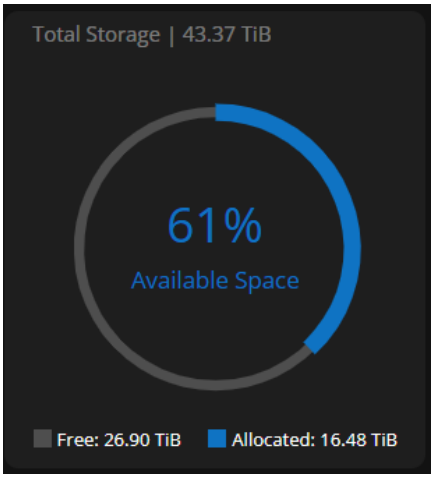


## Overall Health Status

The Overall Health Status shows an overview of the health of your system with a counter for issues reported on the system. You can click the **View Issues** text to launch a dialog that will show details about any problems on the cluster.

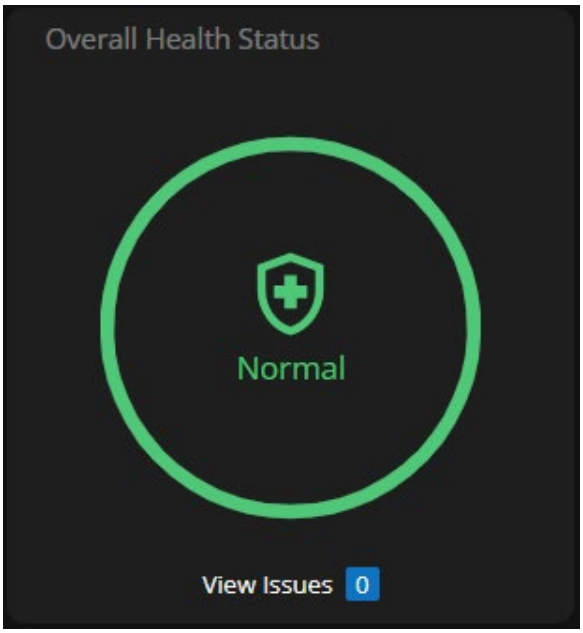


In this example, we disconnected SAN1 to demonstrate the View Issues.

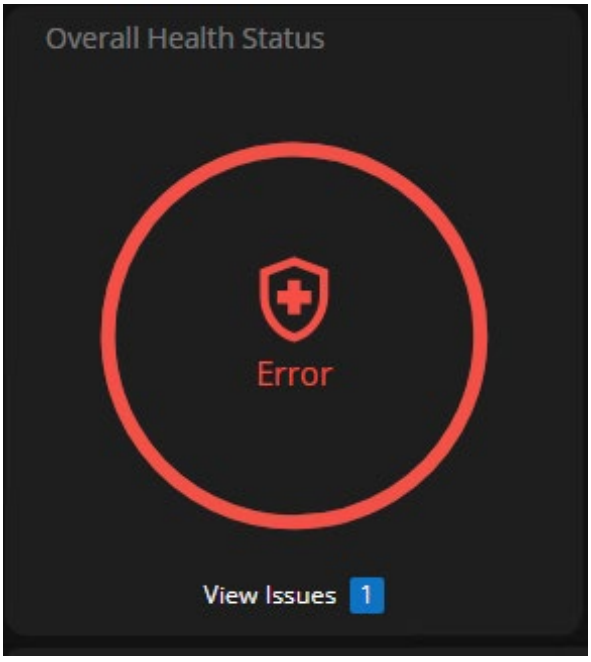


The issue is returned on the host.

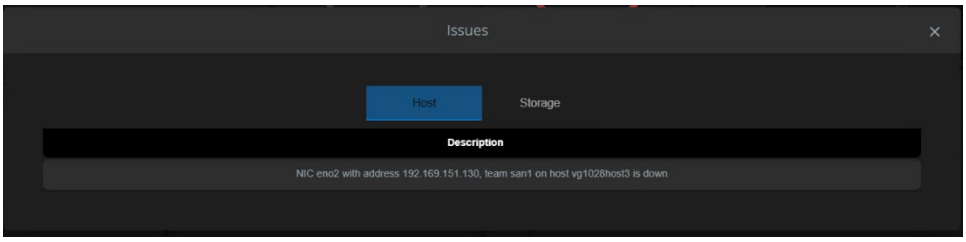


Infrastructure monitoring displays the port down on the dashboard.

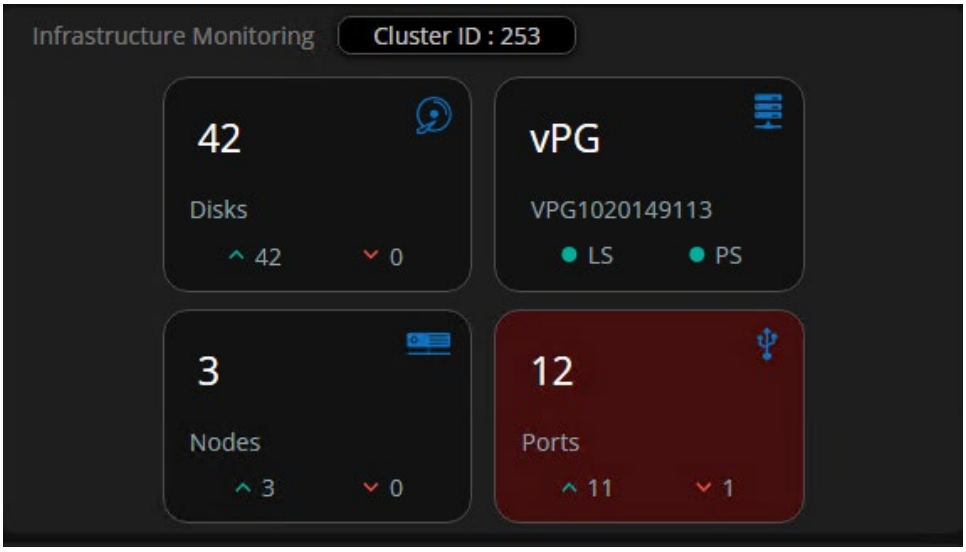


If you open infrastructure monitoring and select the host with the red highlight, you can drill down to the ports and see the port that is disconnected.

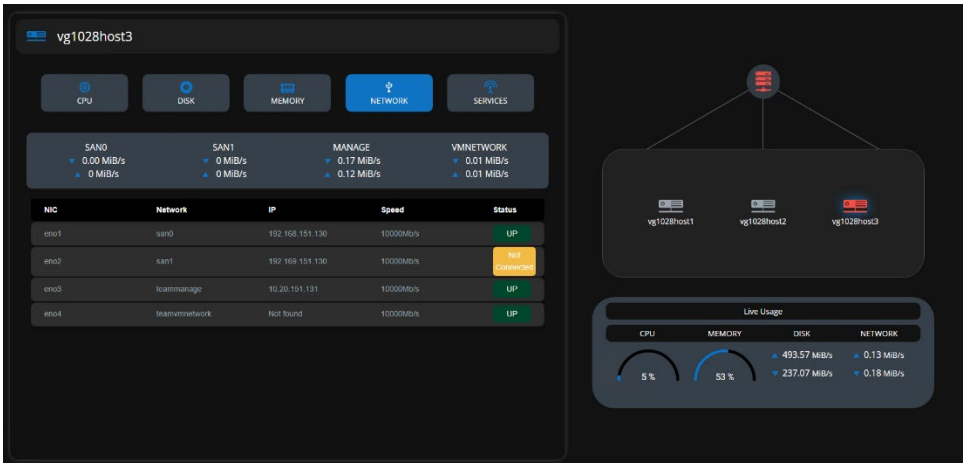


A Warning is also entered in the NICs log.

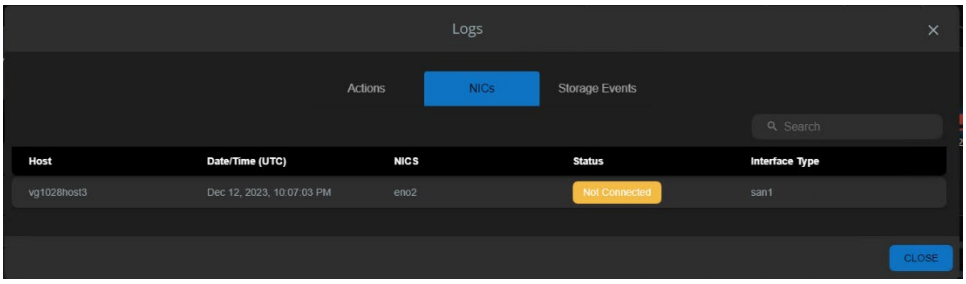


## Infrastructure Monitoring

Infrastructure Monitoring shows the overall health of your vPG, broken down into sections of hardware with issues.

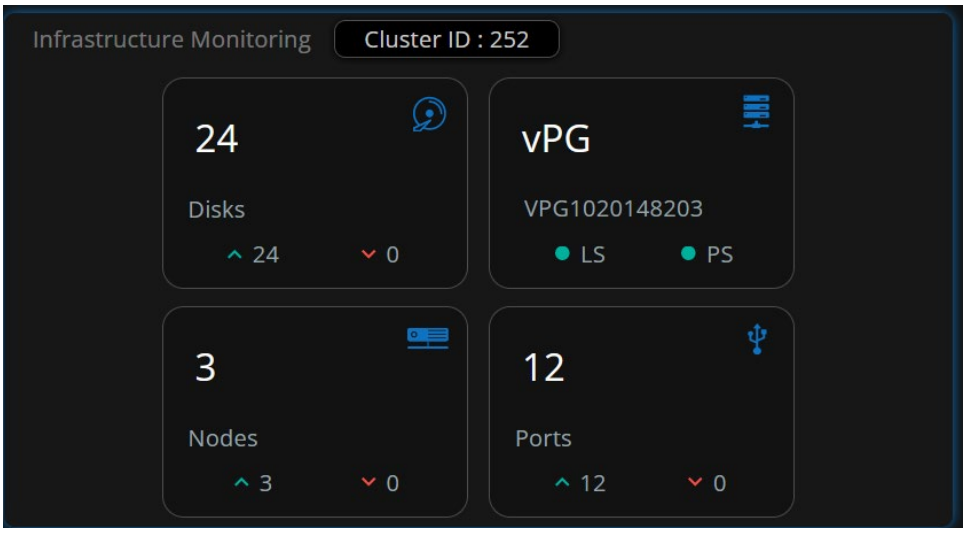


### *Disks*

This section provides a count of the disks on the system. If there are any disks next to the red down arrow, those disks have failed and should be replaced.

### *vPG*

This section shows your system's backend storage health. Both Logical (volume) and Physical (node) states are displayed, as reported by the storage system.

### *Nodes*

This section shows the cluster health from the OpenStack perspective. If any nodes are not responding to heartbeats, they will be listed next to the red down arrow.

### *Nics*

This section shows the health of the Nics (Network Interface Cards) across the cluster.

### *Cluster ID*

The Cluster ID is set during installation and must be unique on the network. Make sure to gather the existing Cluster ID from previously installed clusters and use a unique value for new cluster installed on the same network.

## Instance Performance Overview

The top portion of this graph shows the average CPU and memory usage for all instances on the cluster. The total disk and network access values are also shown for all instances. The bottom sections show the top 5 instances in each category (CPU, memory, disk access, and network access).

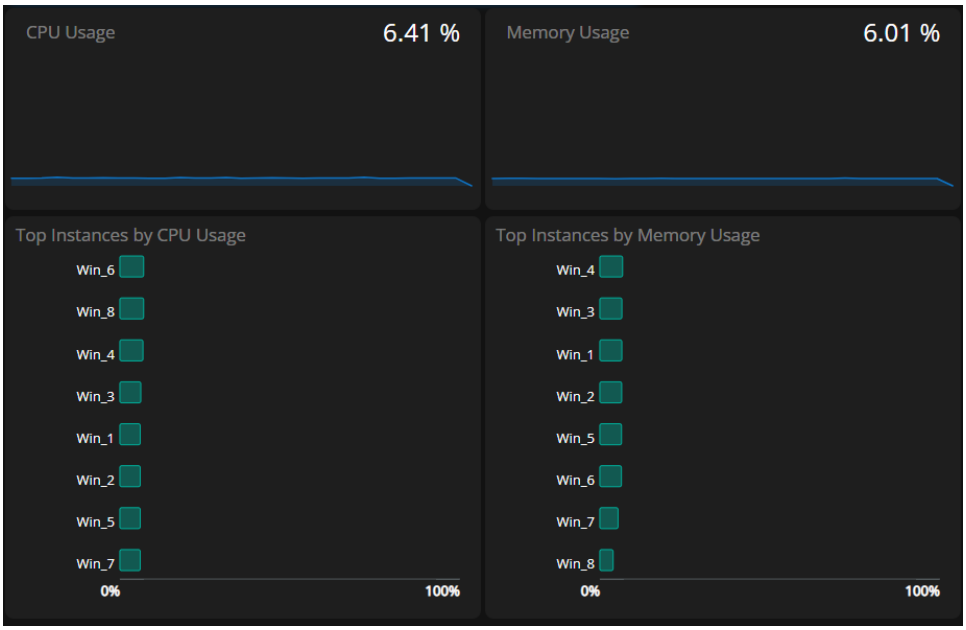


## Detailed Instance Performance

This section describes how to view the detailed performance statistics for one or more instances.

1. Right-click on the instance and select **Add to Deck**.

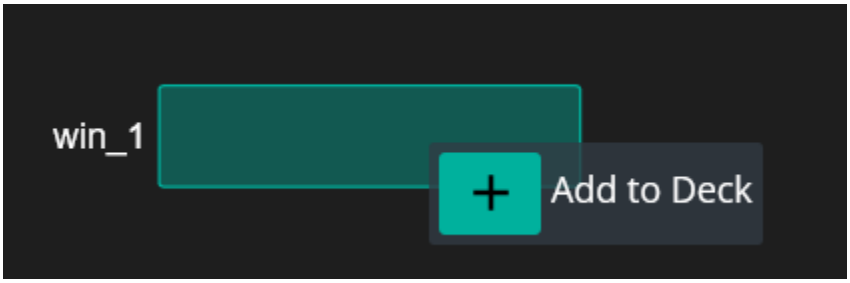


2. After the instances have been added to the deck, they will turn red.

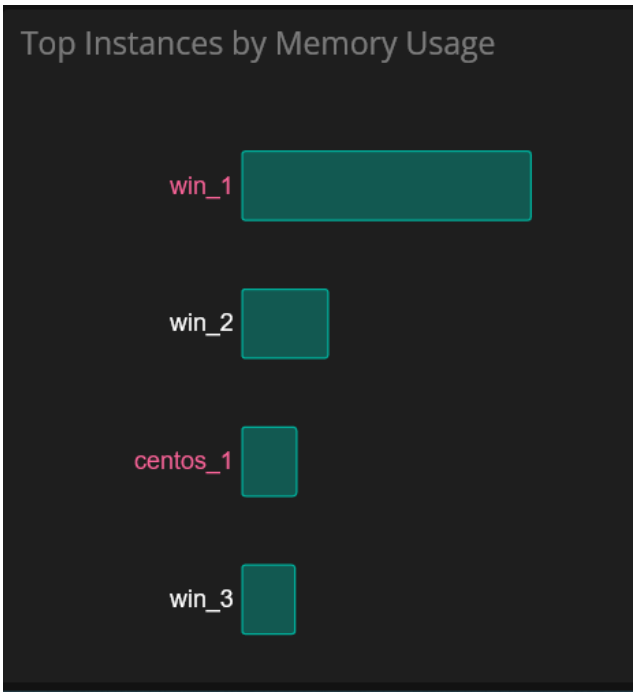


3. In the top-right corner of the USP Management Application, click on the slider to display the advanced performance chart.

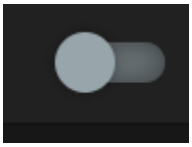

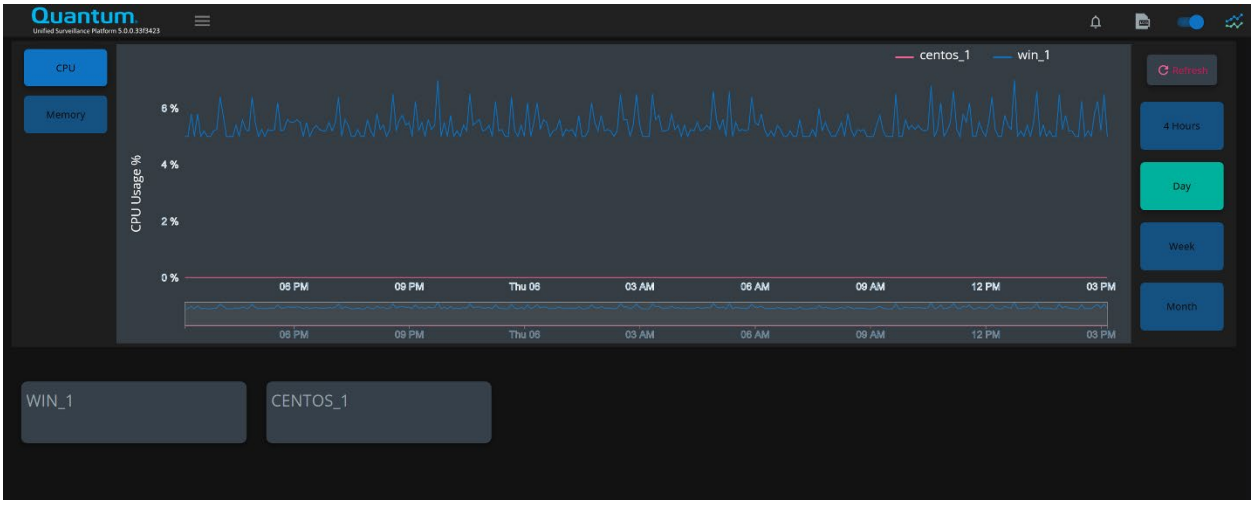


4. To remove instances from the deck, first exit the performance chart, right-click on the instance, and then remove it. You can remove a single instance or remove all instances at once.

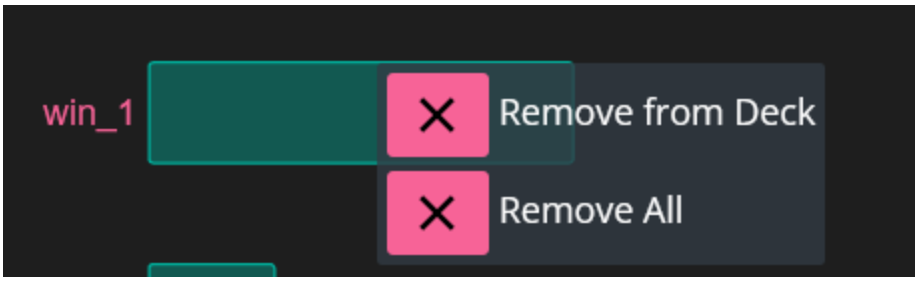


## Infrastructure Monitoring

To view the infrastructure monitoring page, click **Infrastructure Monitoring** displayed on the dashboard side menu.

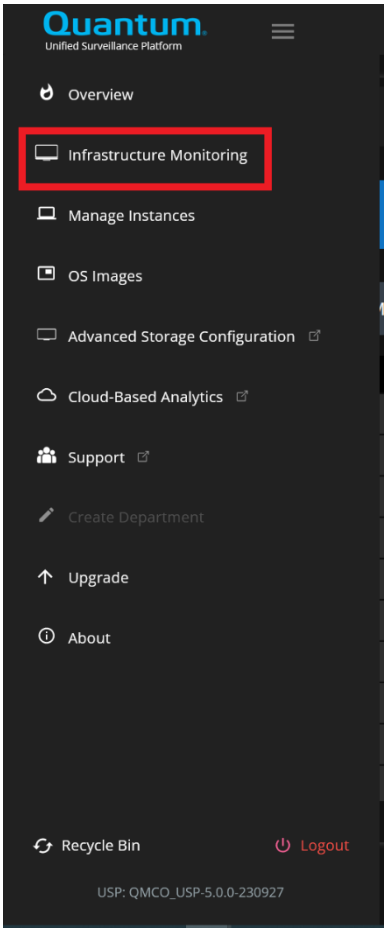


Alternatively, you can launch the Infrastructure Monitoring page by clicking anywhere in the **Infrastructure Monitoring** pane on the dashboard.

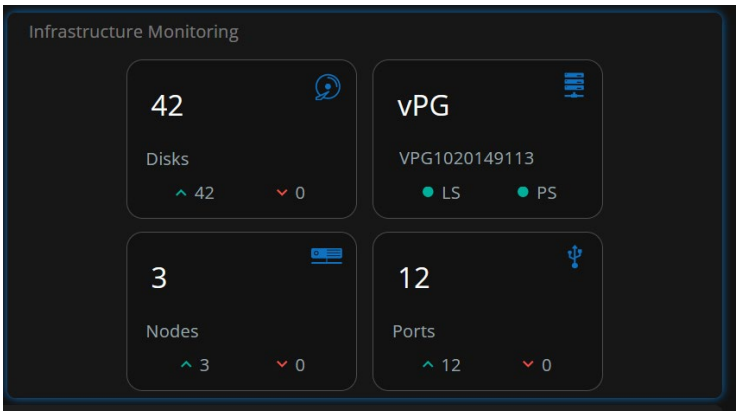


## Hosts View

By default, the Hosts view is shown first. It lists the CPU, disk, memory, services, and network usage for each host in the cluster. This view also shows the status of each host. On the right side of the screen, the view shows the specific hardware information for CPU, disk, memory, and status.

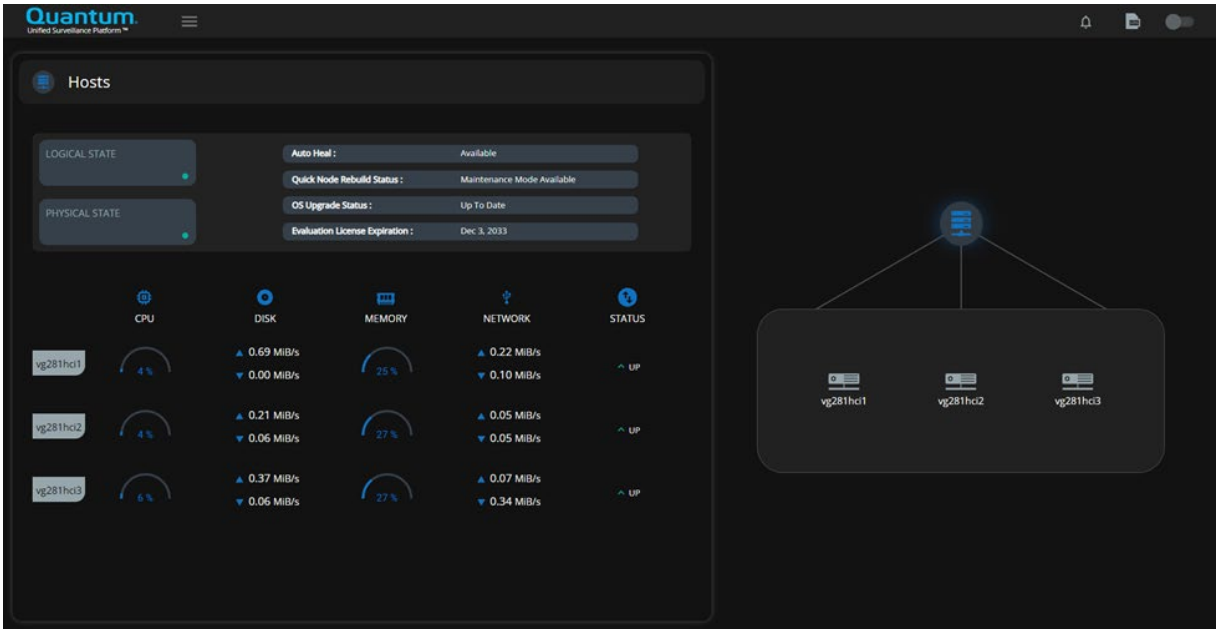


On the right side of the display, select one of the hosts to see its details.

## CPU View


## Acuity Storage Disks View

To view the Acuity Storage details, click on the **Storage** icon on the right side of the screen. This shows physical state of the disks in the server and the health of each disk that is passed through to the Acuity Storage VMs.

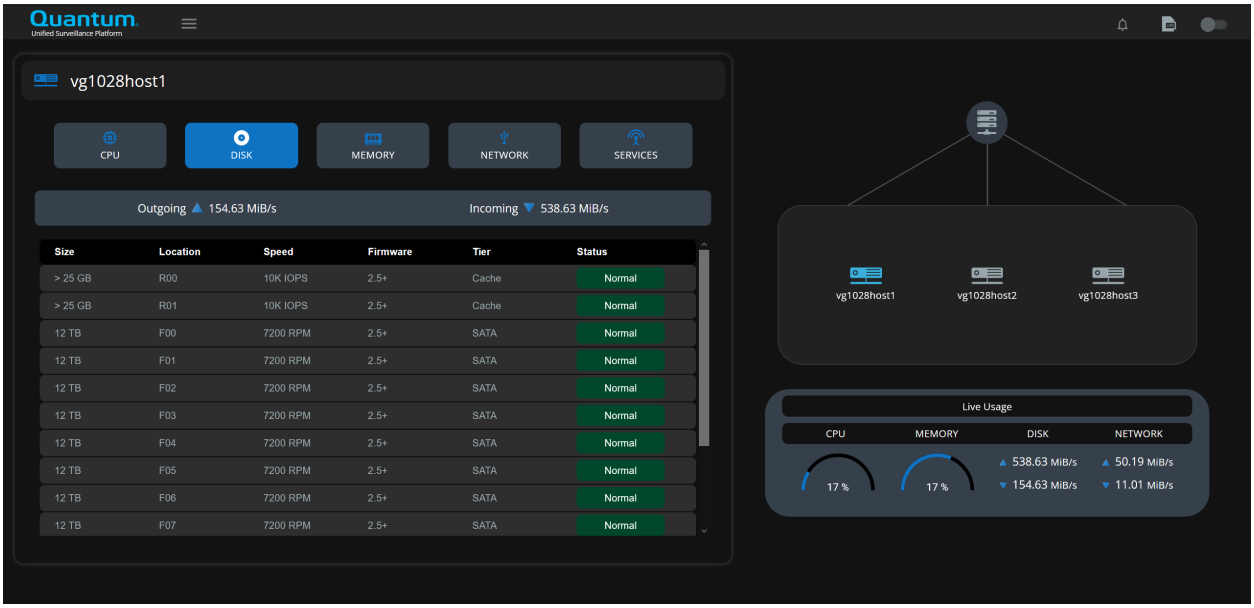

## Memory View


## Network View

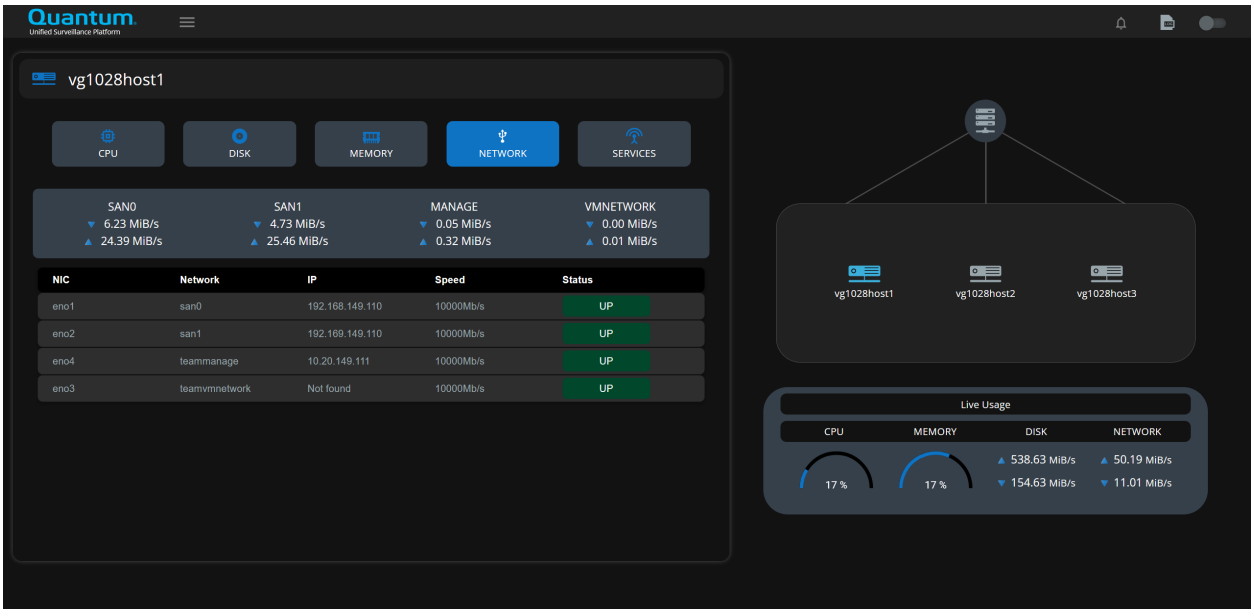

## Services View


# Manage Instances

## Upload OS Image

To create an instance, you first must upload an instance OS image.

1. To upload an image, select **OS Images** from the navigation menu.

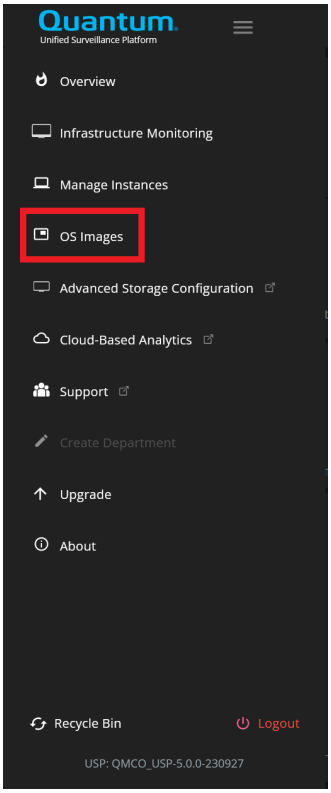

2. Press the plus button in the top right corner.

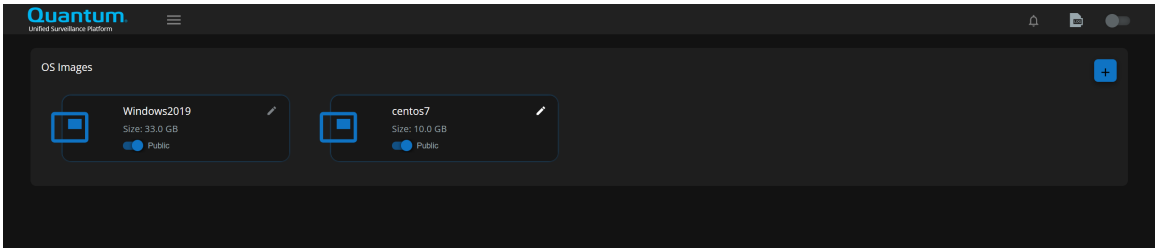


3. Enter the Name, OS Type, OS Distro, and browse to the raw image on your local file system. Press **CONFIGURE** to upload the image.

NOTE: *The amount of time for this operation varies and will depend on the local network configuration.*

CRITICAL: *Do not refresh the web browser during this operation. Doing so can cause the operation to fail and it will need to be restarted.*

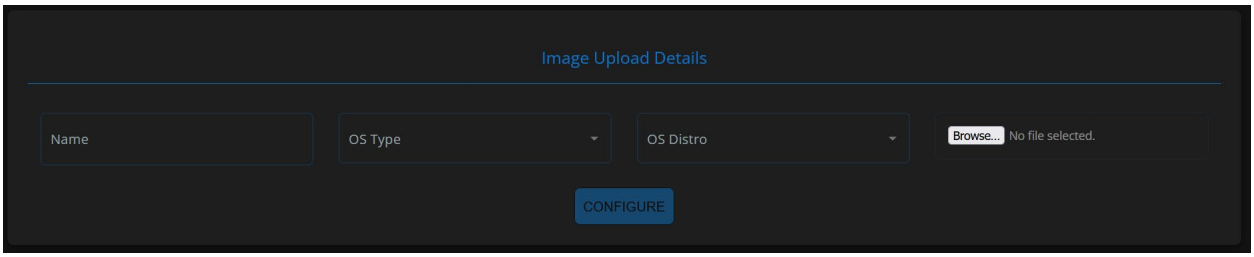


## Create Instance

This section describes how to create an Instance.

1. Navigate to the Instances view by selecting **Manage Instances** from the left-hand navigation menu.

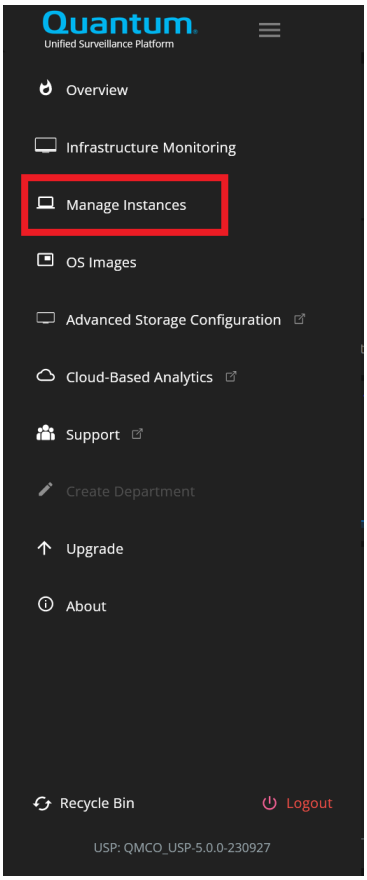


2. Click on the 'Plus' icon in the top-right corner of the **Manage Instances** pane.

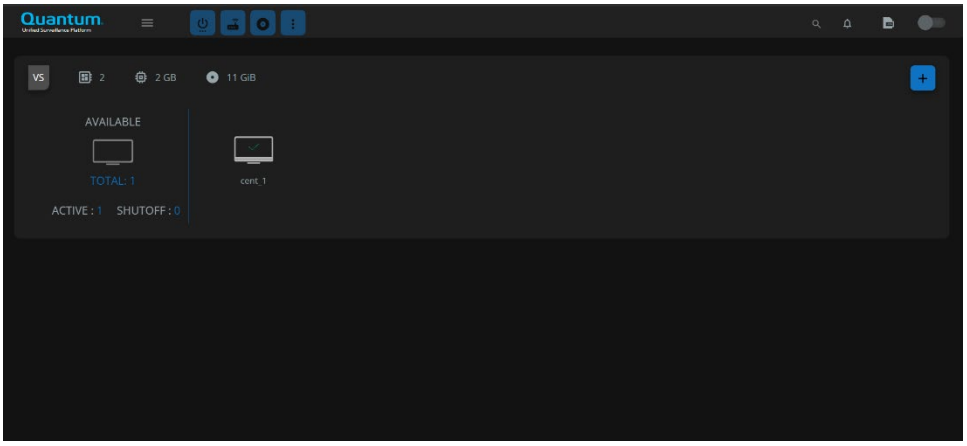

3. Enter the Instance details and press **NEXT.**

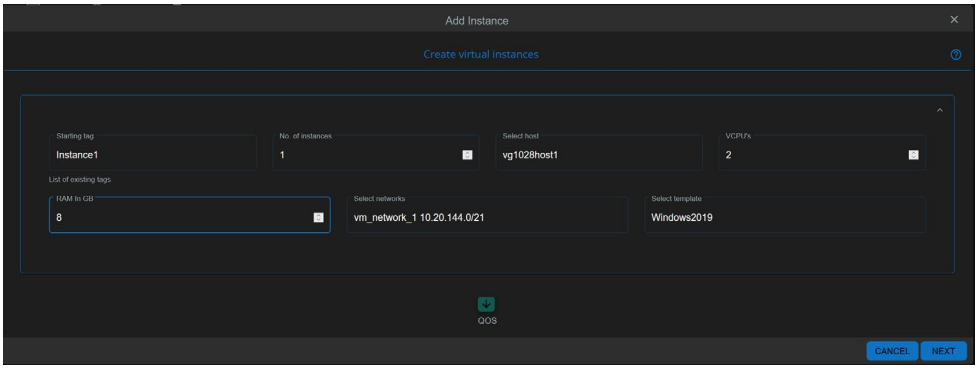


CRITICAL: *This release only supports creating a single instance at a time. Do not attempt to create more than one instance at a time.*

NOTE: *Changing the QoS settings is not supported in this release.*

4. Confirm the instance name and enter the IP information.
5. Select a GPU to be attached to the instance if needed.

NOTE: *GPUs will not be available until the steps in section [Configure Cluster for GPU Passthrough]() are completed.*

6. Check the **Auto Evacuate** checkbox if you want the instance to be migrated to a new host in case of failure.
7. Enter the Acuity storage details for the boot volume.
8. Select the Tier (SSD or HDD) for the OS boot disk if your cluster is configured with multi-tier support.

NOTE: *It is recommended to choose EC 3p for the Instance OS boot volume to get the best protection to performance ratio.*

9. Verify the setting and then click on **CREATE** to create the instance.

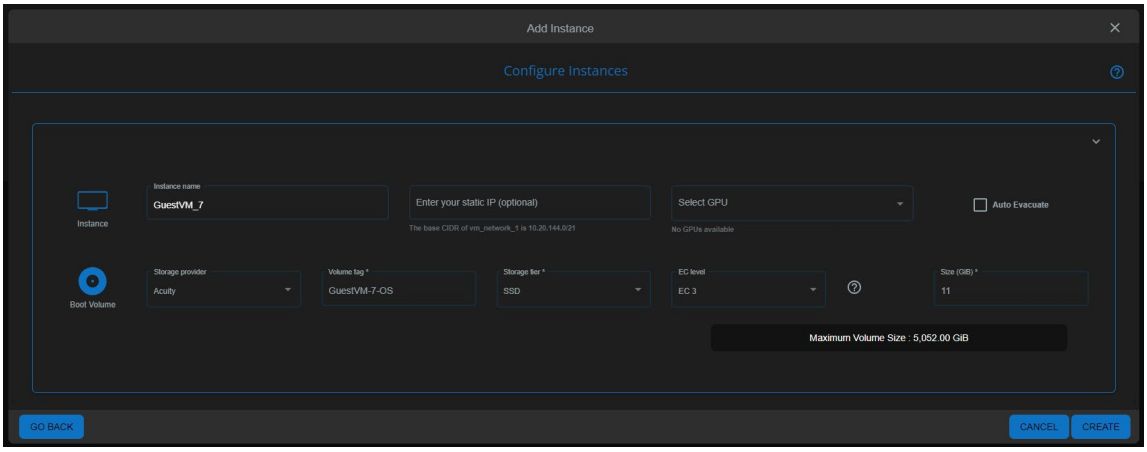

An example for multi-tier Instance creation, placing the OS volume on the SSD tier is shown below.

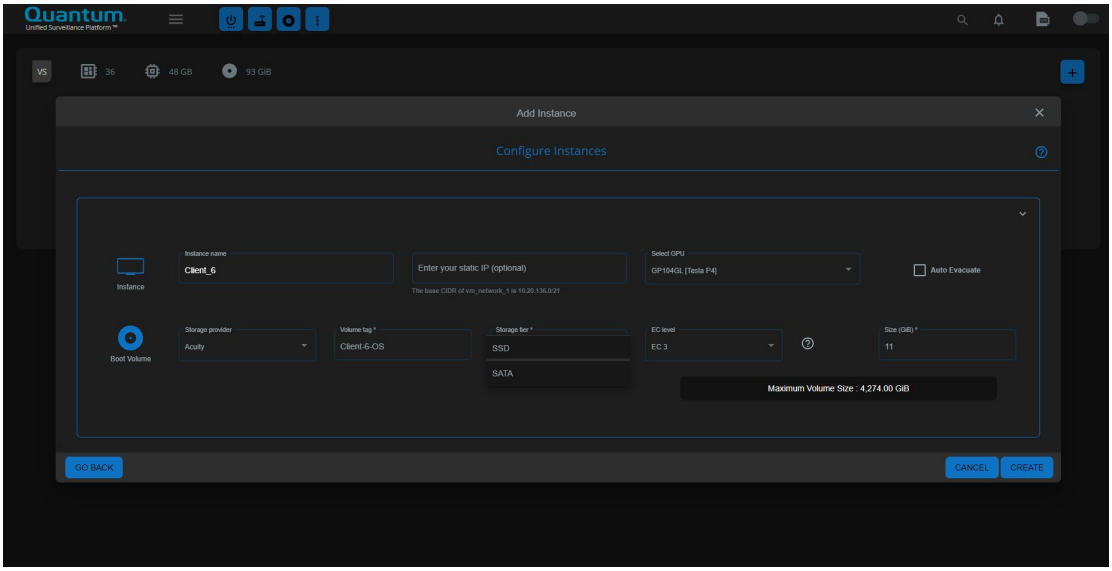


## View Instance Details

To view the instance details, hover your mouse over the instance icon in the **Manage Instances** pane.

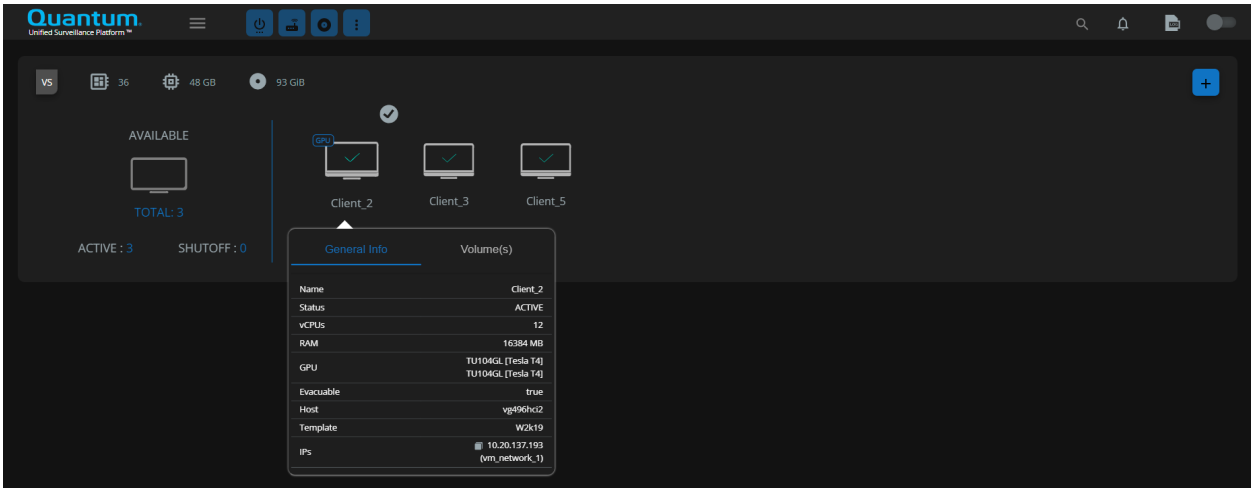


## VNC Console

To view the VNC console of an instance, double-click the instance icon in the **Manage Instances** page. This launches a separate browser tab where you can view the VNC console. Make sure to enable popups for the USP Management Application.

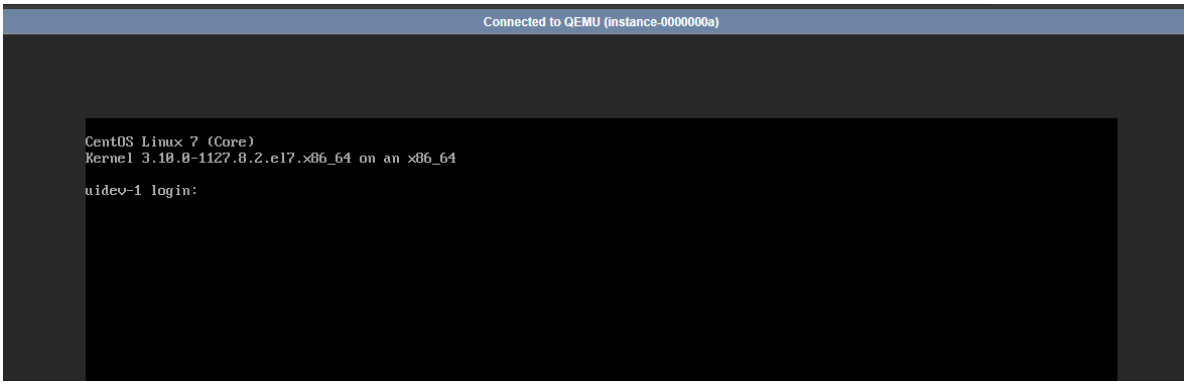


## Instance Power Options

Select an instance in the **Manage Instances** section, and select the **Power Settings** icon in the top of the web application to perform power options on an instance.

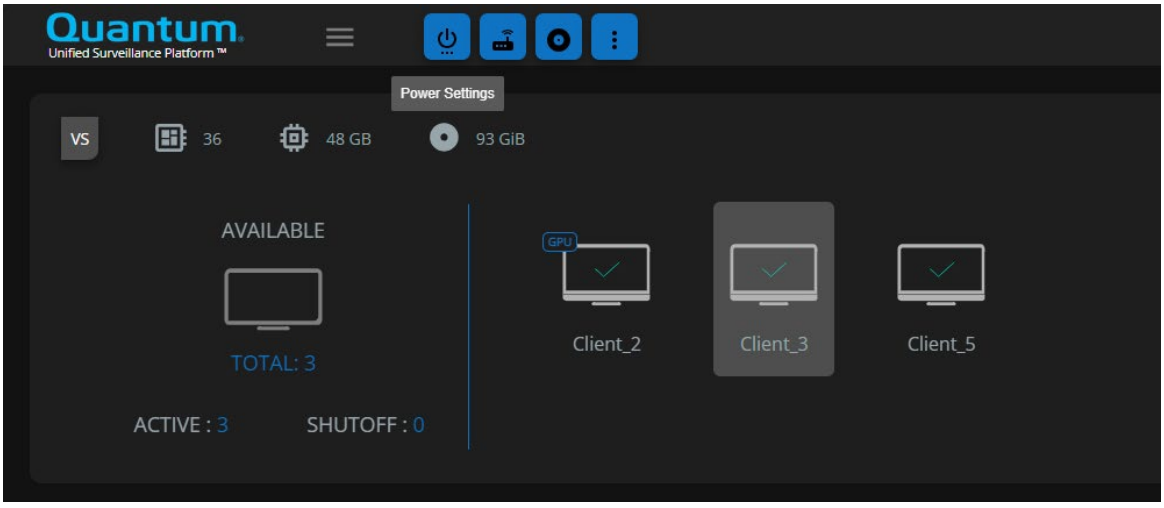


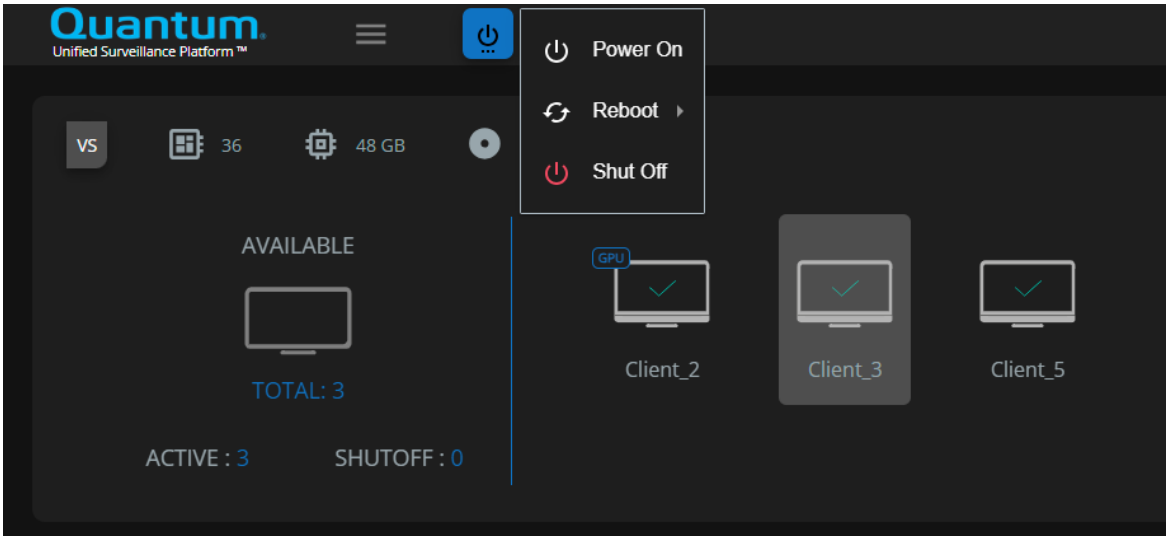

## Live Migration

To move an instance to a different host, navigate to the **Manage Instances** page, click on the **More Settings** icon at the top, and then select **More > Migrate.** From this dialog, you can choose where to move your instance.

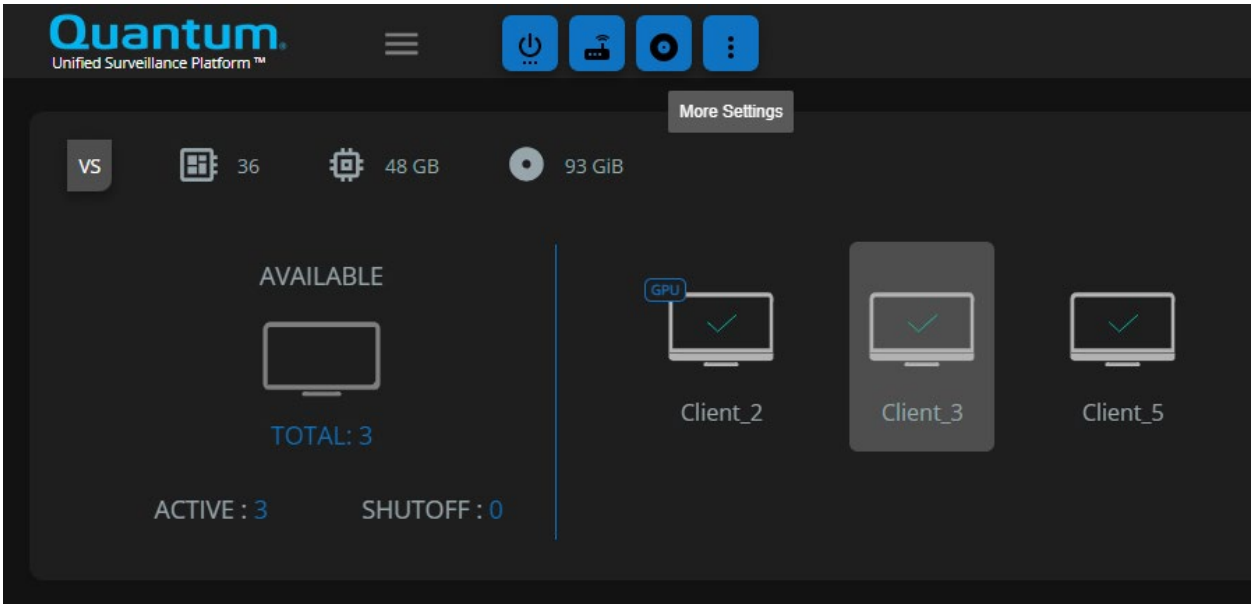


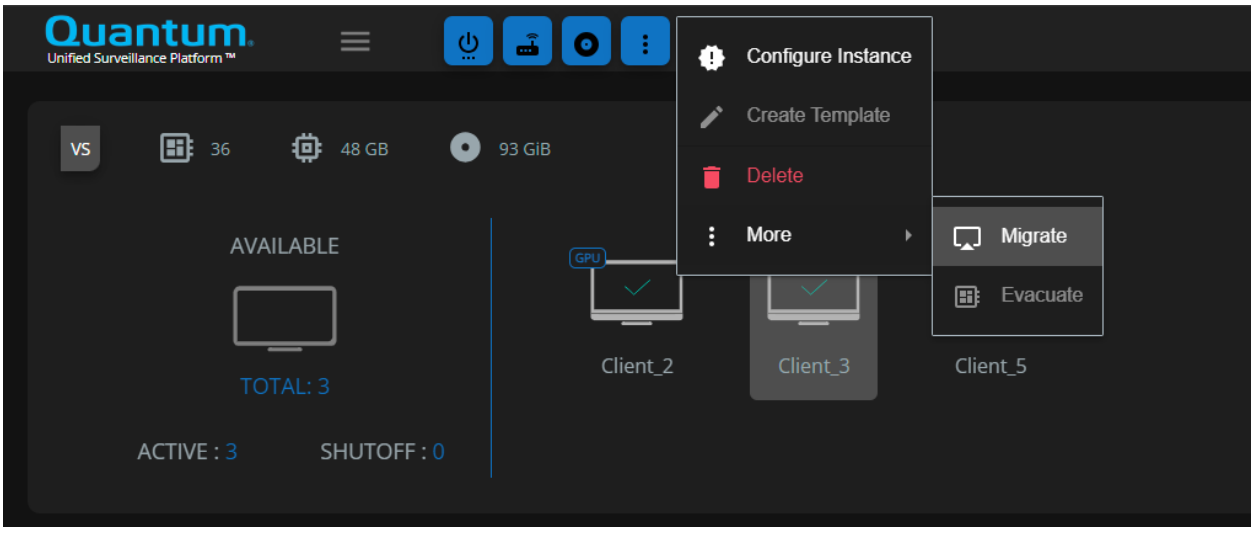


NOTE: *Certain virtual machine features, such as live migration and evacuation, require similar CPU architectures across servers.*

## Instance Volume Update

1. Select an instance in the **Manage Instances** section and select the **Disk** icon in the top of the web application to perform Volume actions on an instance. Then Click on Attach/Detach Volumes.

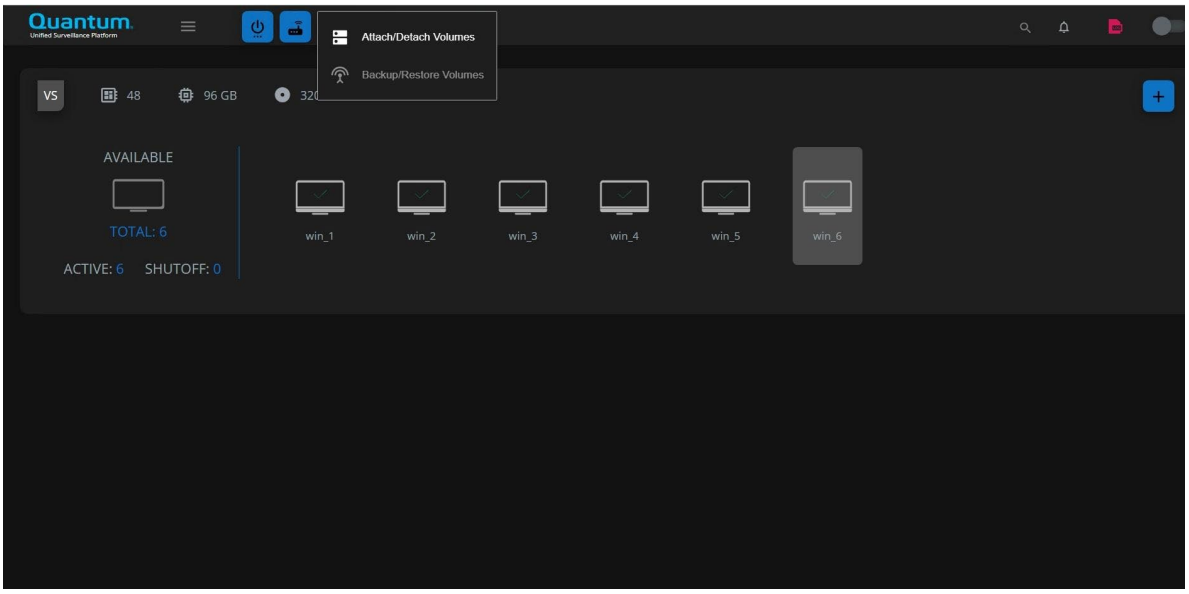


2. Click on the Disk you want to edit. Enter the Details and click on “Apply Changes.”

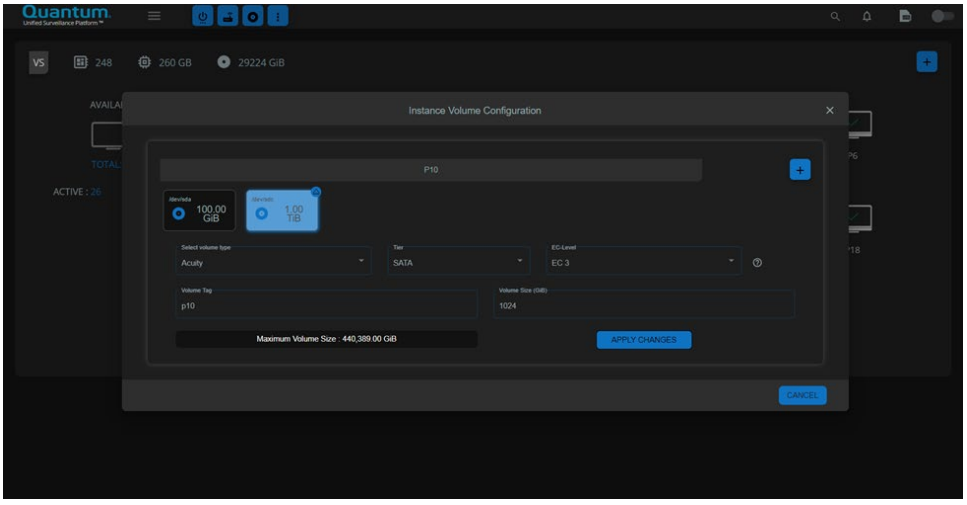

3. Click on “Confirm.”

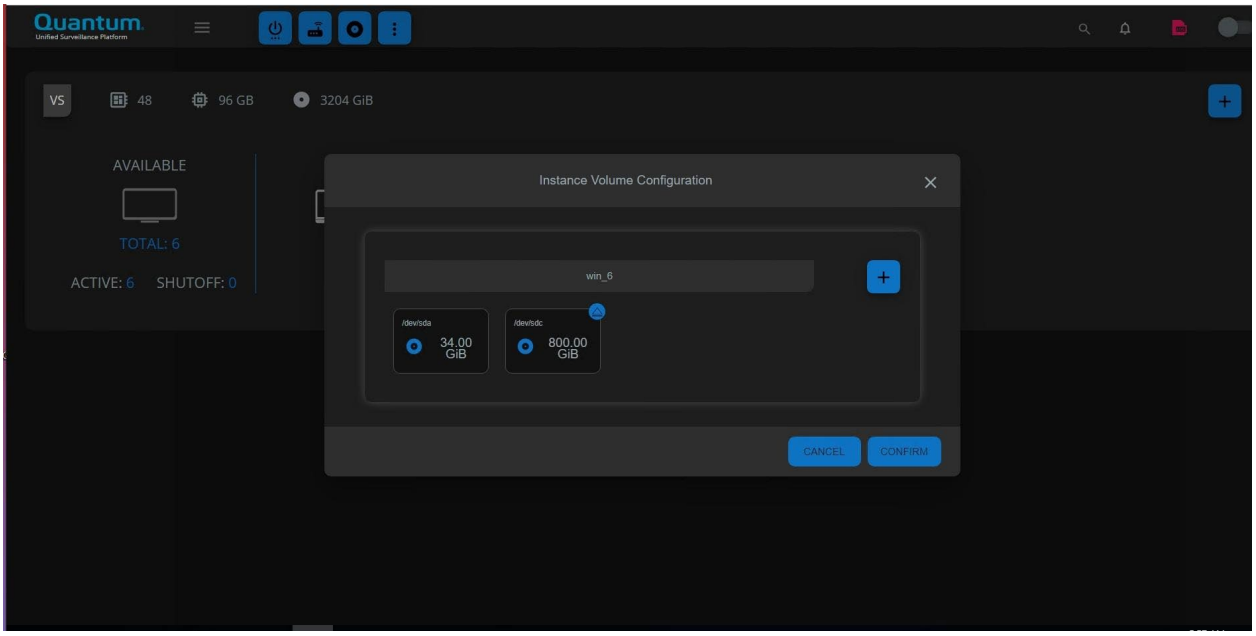


## Instance Volume Addition

1. Select an instance in the **Manage Instances** section and select the **Disk** icon in the top of the web application to perform Volume actions on an instance. Then click on Attach/Detach Volumes.

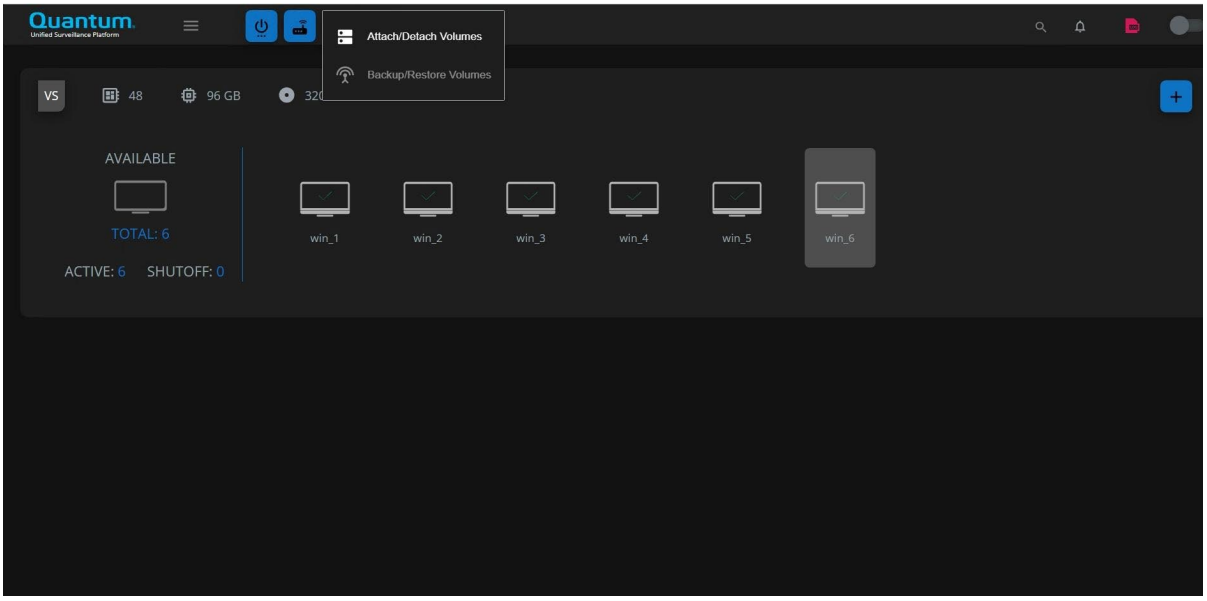

2. Click on the **+** Button. Either select an already **Existing Volume** or fill in all the details to create a **New Volume.** Click on “Add Volume.”

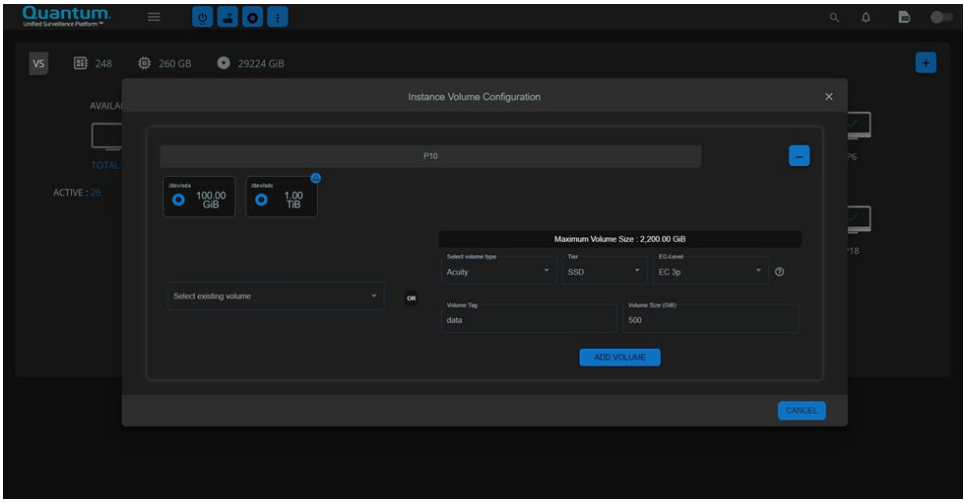


3. The **New Disk** will appear in the list. Click on **Confirm** to add the volume to your **VM.**

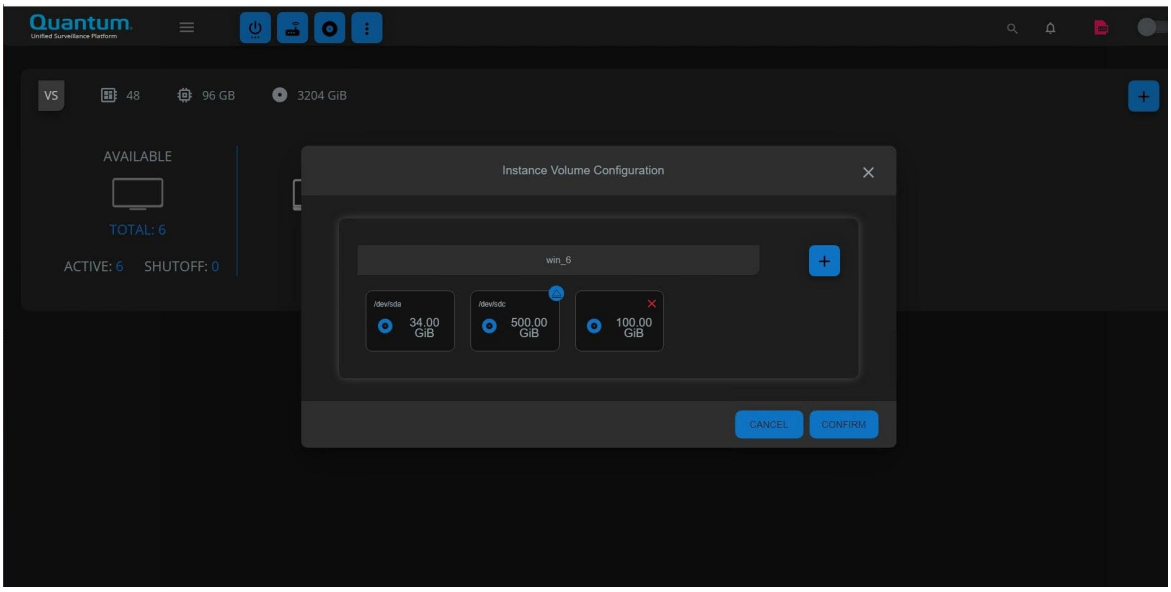

## Instance CPU, RAM & GPU UPDATE

1. Select an instance in the Manage Instances section and select the “Three Dots” icon in the top of the web application to perform CPU/RAM/GPU change actions on an instance. Then Click on “Configure Instances.”

NOTE: *GPUs will not be available until the steps in section [Configure Cluster for GPU Passthrough](#) are completed.*

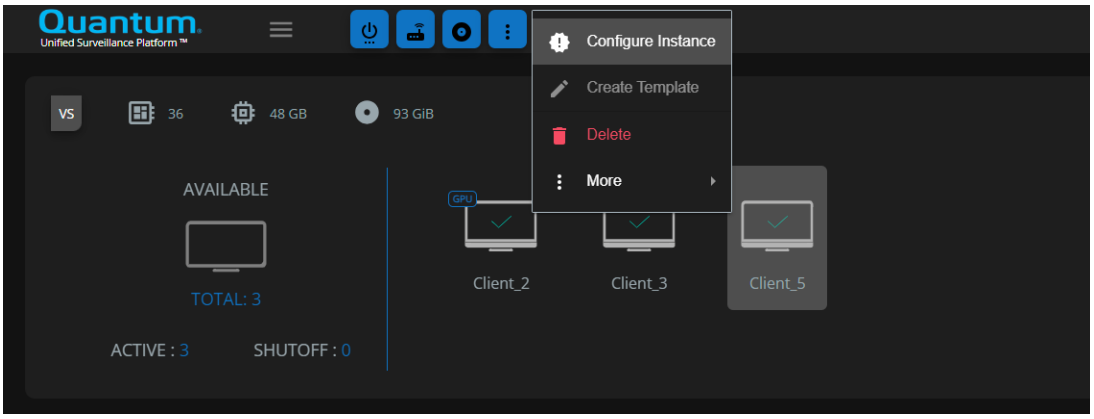


2. You will see the values for vCPU, RAM, along with the GPU attach/detach option. Adjust them to your requirements. Then Click on “Apply.”

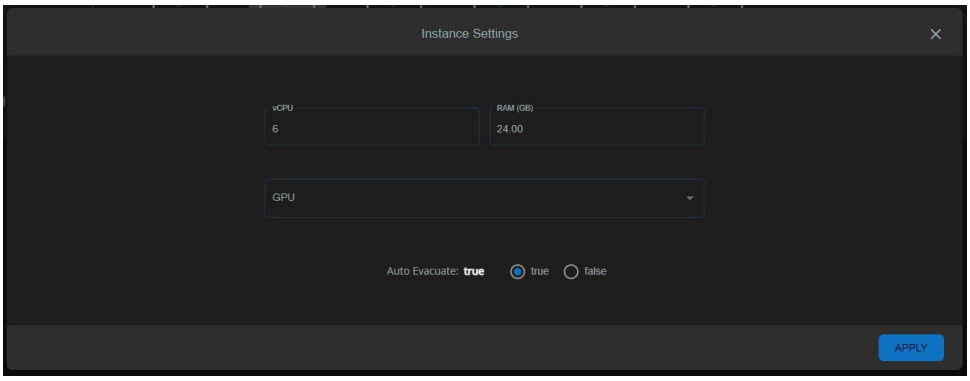

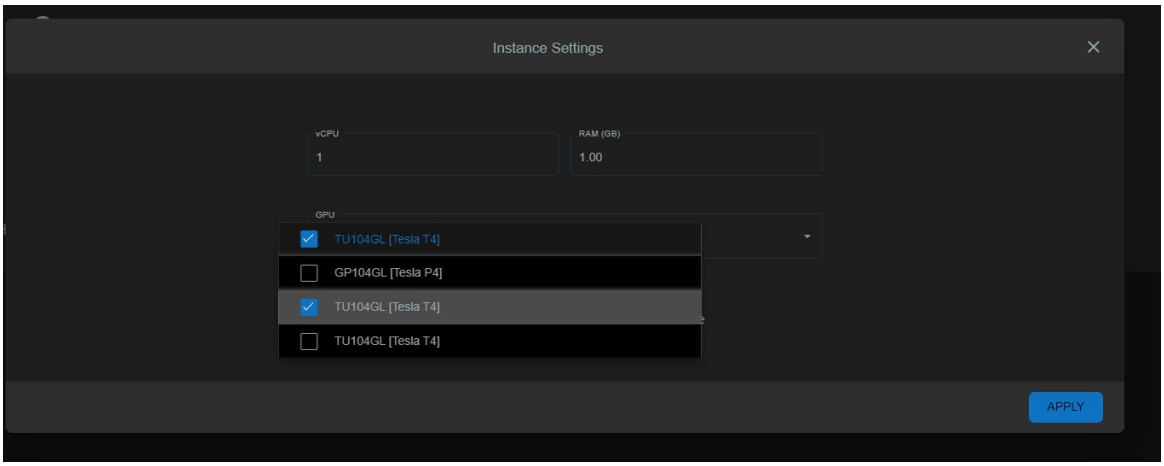


3. Click on “Apply” again. Then, wait for the VM to reboot.

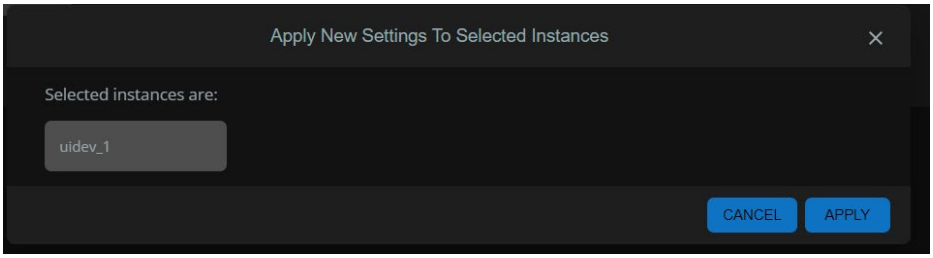


## Security & Firewall

The VM ports can be managed through the management interface, and any rules can be applied.

By default, there are six rules already present for every virtual machine: three egress rules and three ingress rules. The six rules keep all traffic open for the virtual machine. These rules can be deleted, or new rules can be added, as per the requirements.

To manage firewall settings, follow these steps:

1. Select the VM and click on network settings, under Network Settings select **Manage Firewall.**

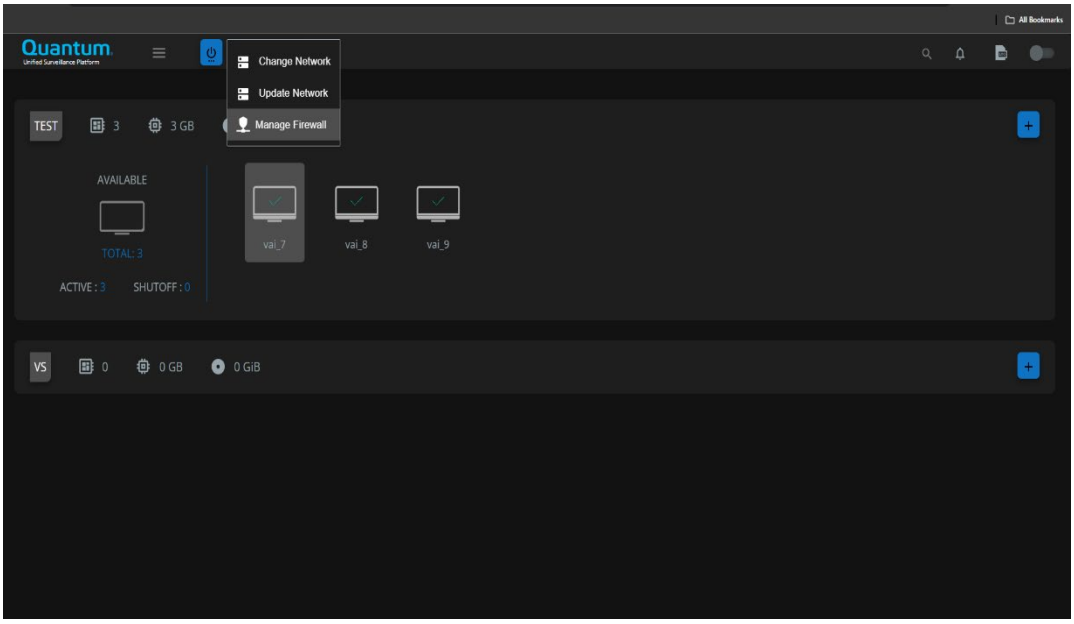


2. The dialog box shows the current rules that are available for the virtual machine. There are two types of rules that can be added either Ingress or Egress.

## *Adding Rules*

1. By default, an Ingress rule will be created. To view Egress rules or to create Egress rules, first switch the slider to Egress (can be seen on top right of the pop up)

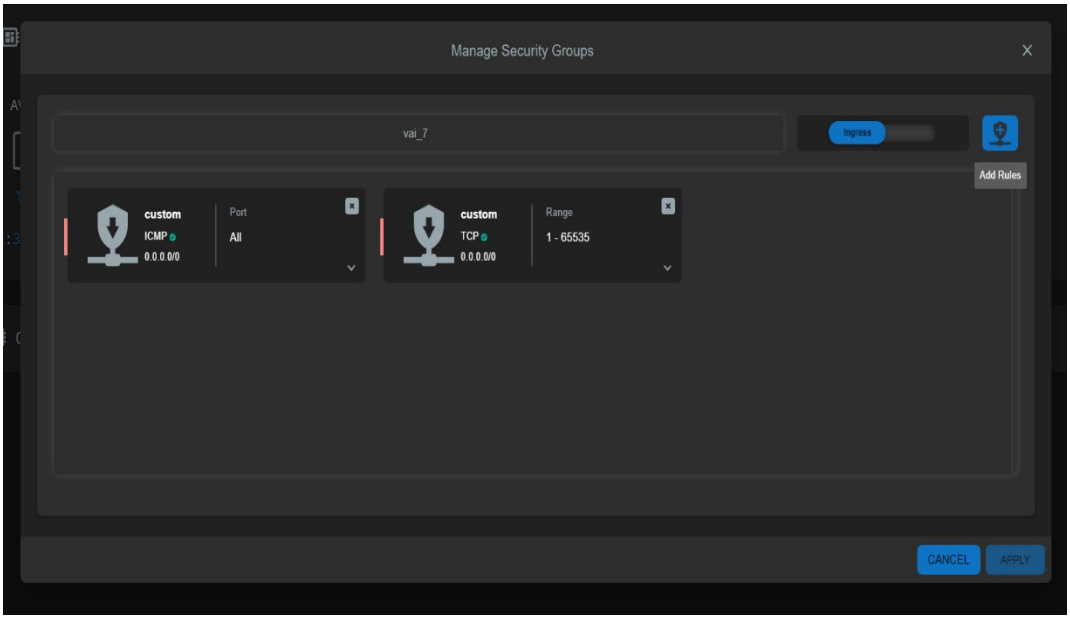

2. When you click Add Rules, a pop-up displays two options – **Pick Common Rules** and **Create Custom Rules**
    - **Common Rules -** A pre-configured list of rules. Currently there is only one rule in common rule: **SSH**.
    - **Custom Rule** – Any rule other than common rules can created by the user, based on the requirement.

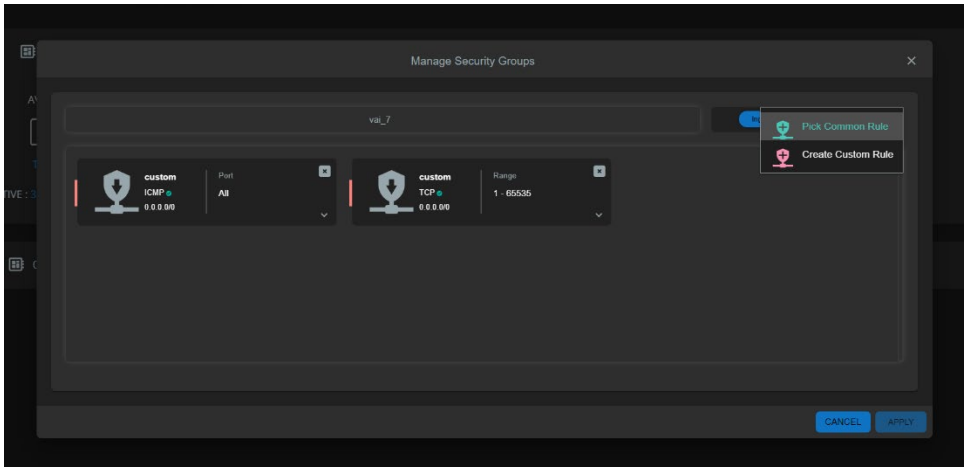


3. If you click Pick Common Rule, the following interface displays with the common rule as SSH. To apply the SSH rule, select **SSH** and click on **APPLY** button.

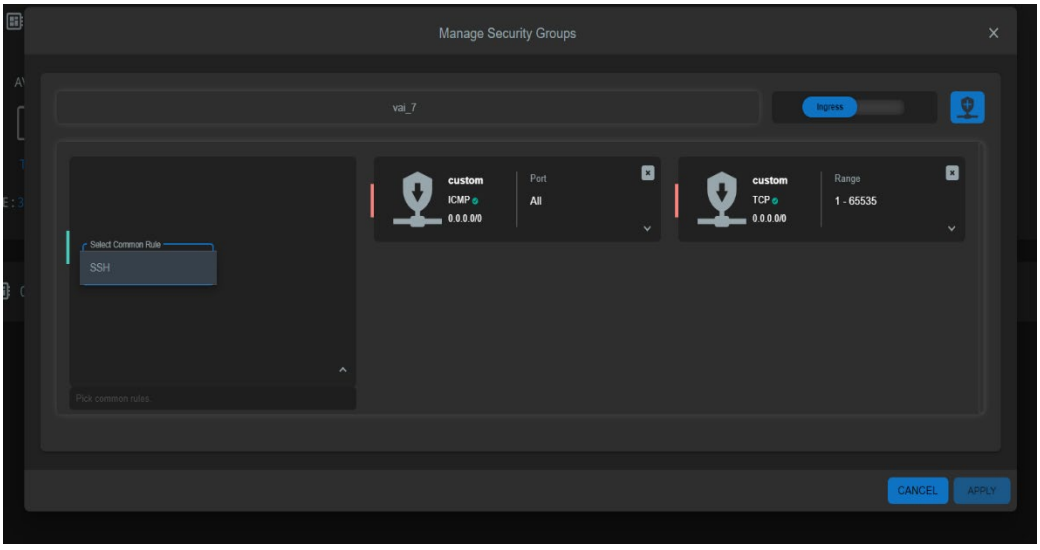

4. After clicking the Apply button, an interface displays to confirm the added rule. Click on **CONFIRM** button to add the rule.

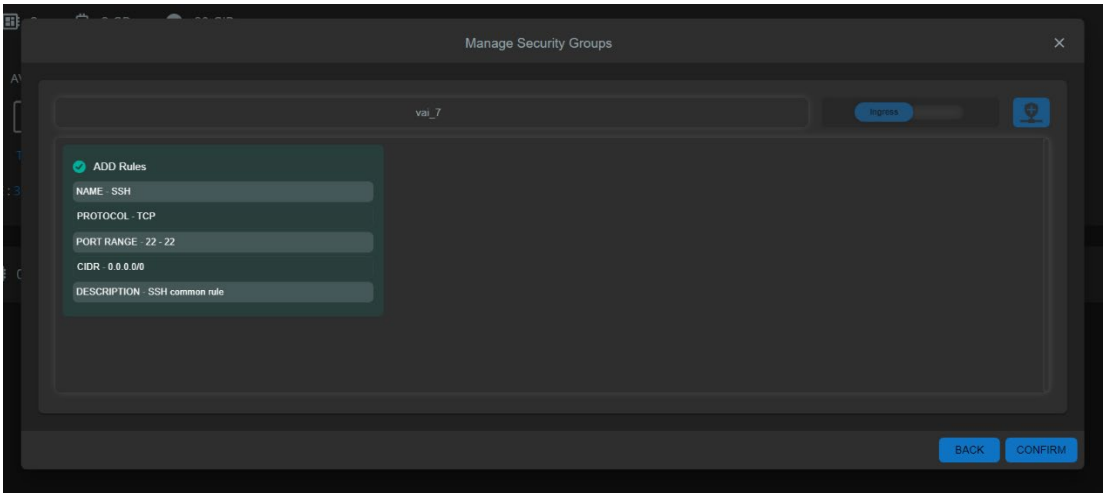


5. Similarly, instead of common rule, a custom rule can be added. Click on **Create Custom Rule.**

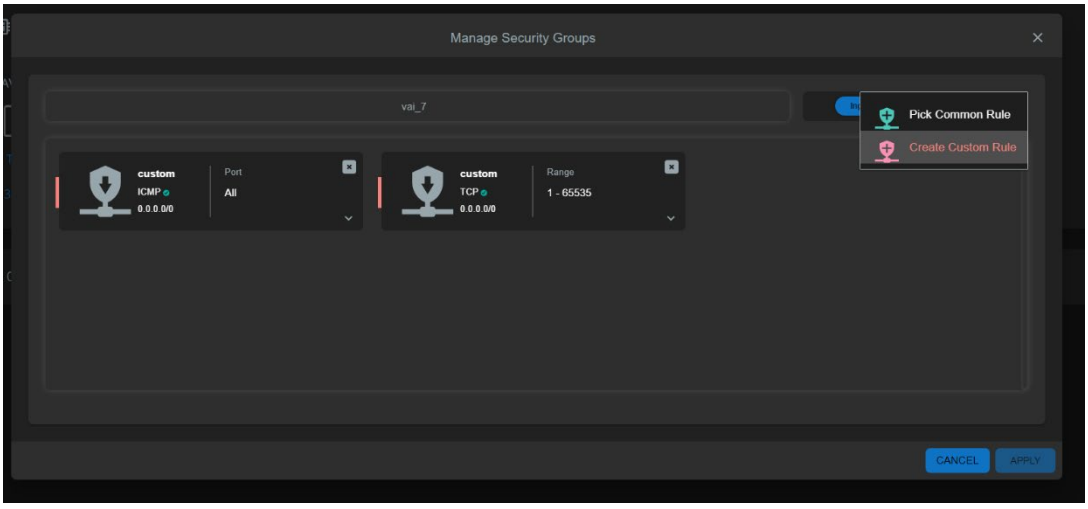


6. A new rule would display:
   - Name – Give any name to the rule.
   - Type – Type can be either **Range** or **Port**.
   - Select Protocol – Currently three protocols are supported **TCP**, **UDP**, and **ICMP**

- Description – Give some description to the Rule.

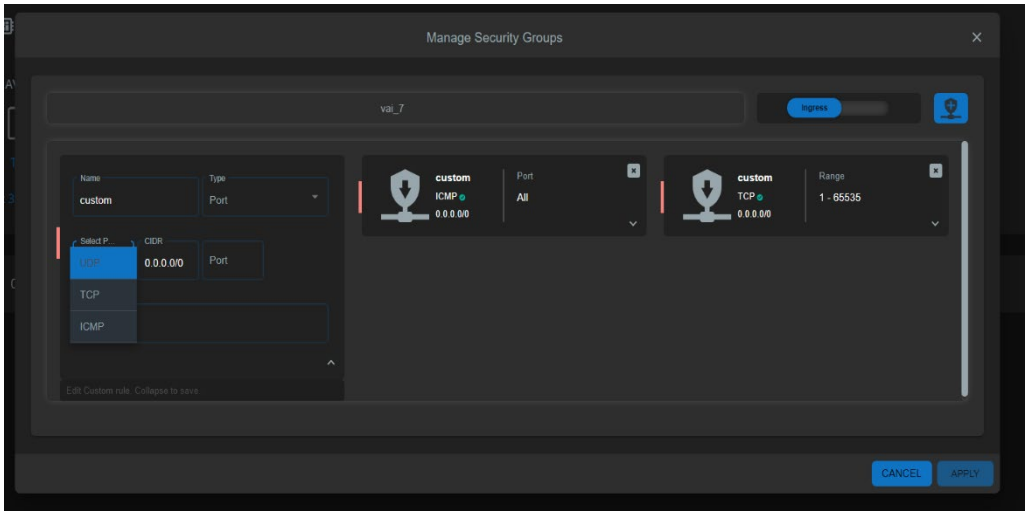


7. Once the rule details are complete, click on **APPLY**. A confirm dialog box would show to confirm the new rule, click on **CONFIRM** to apply.

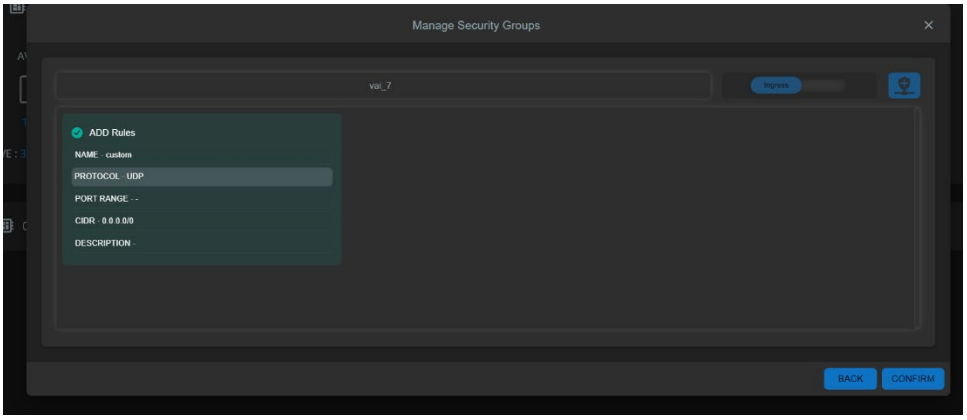

*Deleting rules*

1. To delete any rule, click on the cross icon and select **Apply.**

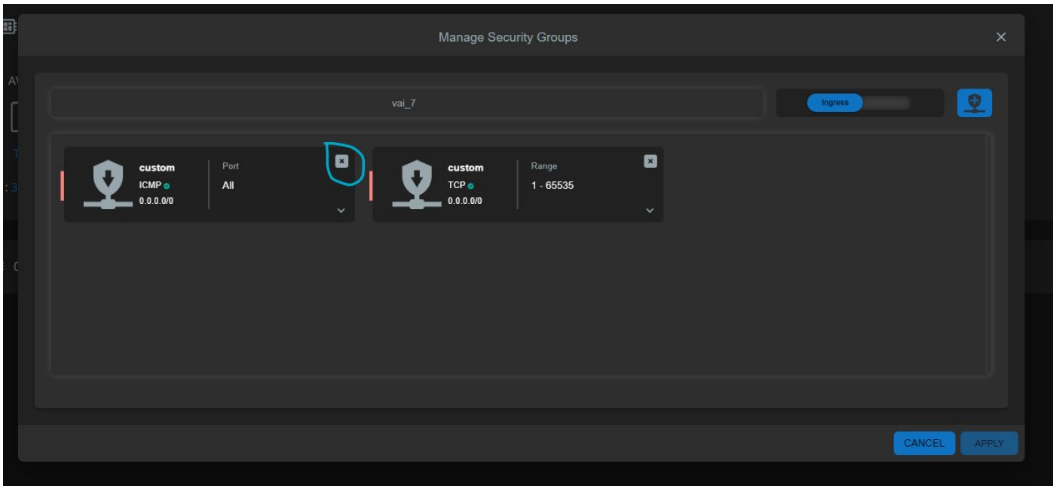


2. A confirm dialog box displays to confirm the deletion. Click **CONFIRM** to apply.

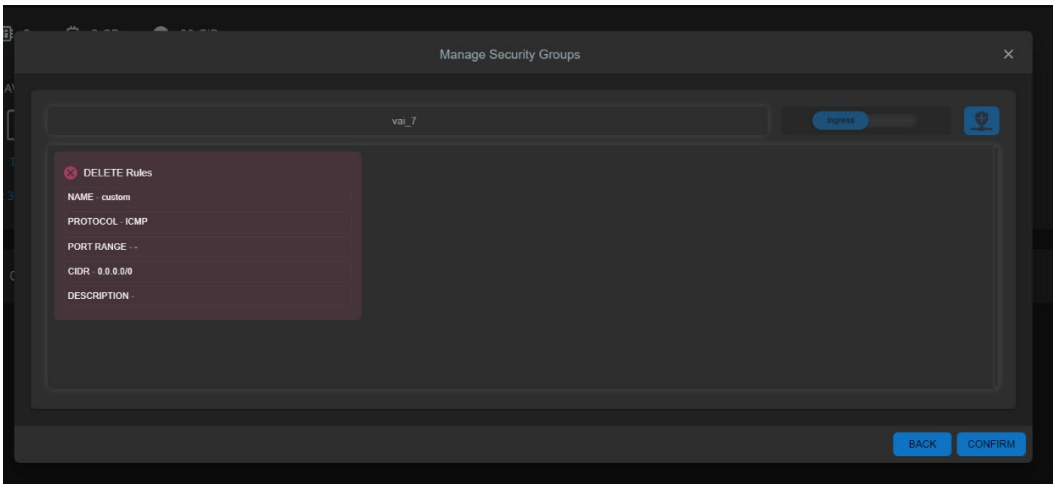


# Health and Monitoring

## Events & Logs

You can view the critical events and the logs by clicking on the **Bell** and **Log** icons, as shown in the top-right corner of the application.

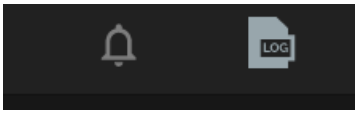

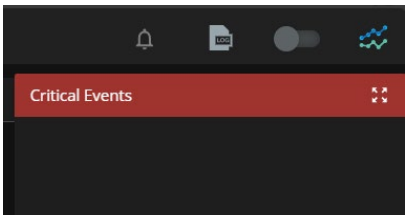


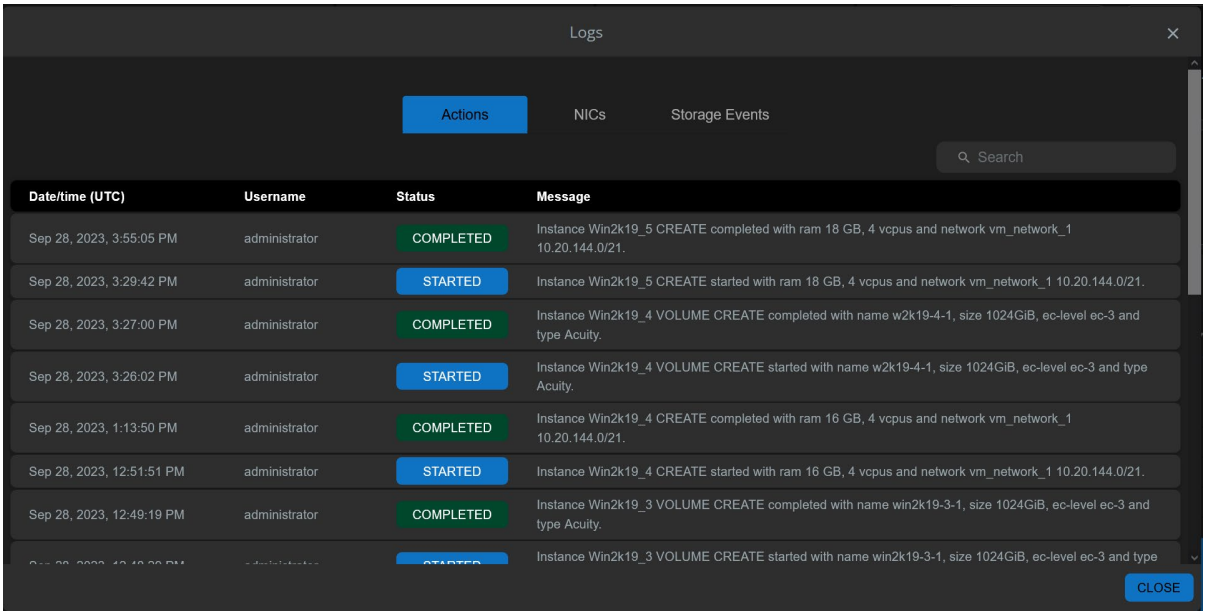


## Viewing Events & Tasks

To view the events and tasks, click the **Log** icon in the top right corner of the management application to launch the **Logs** dialog.

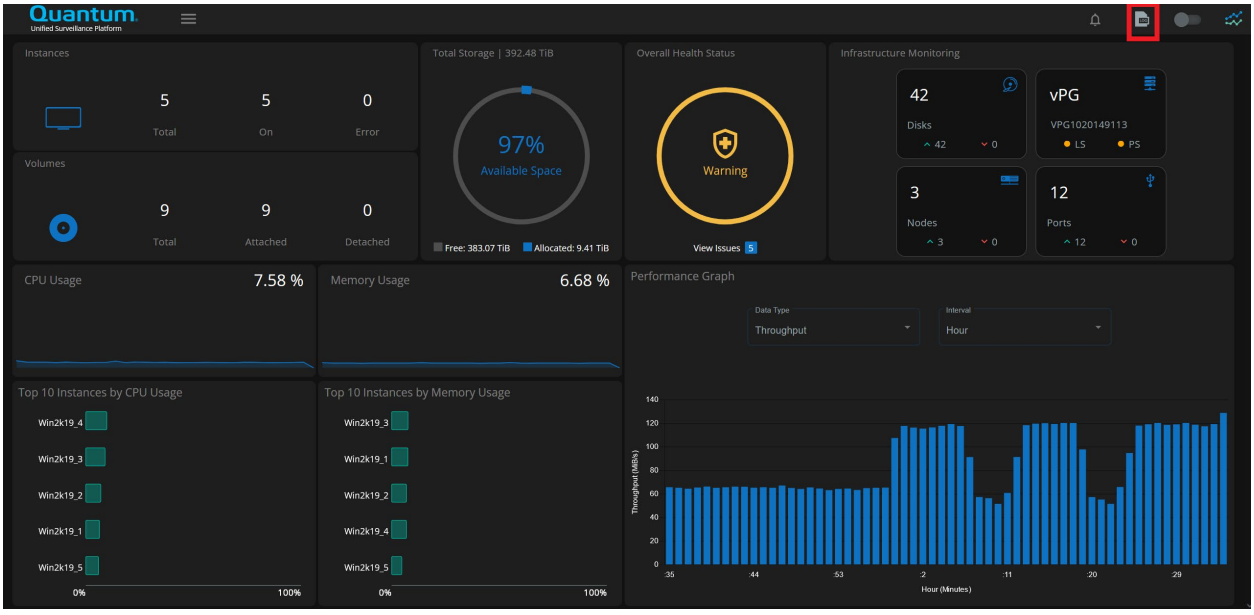

## *Actions*

Shows the user-initiated actions like creating instances.

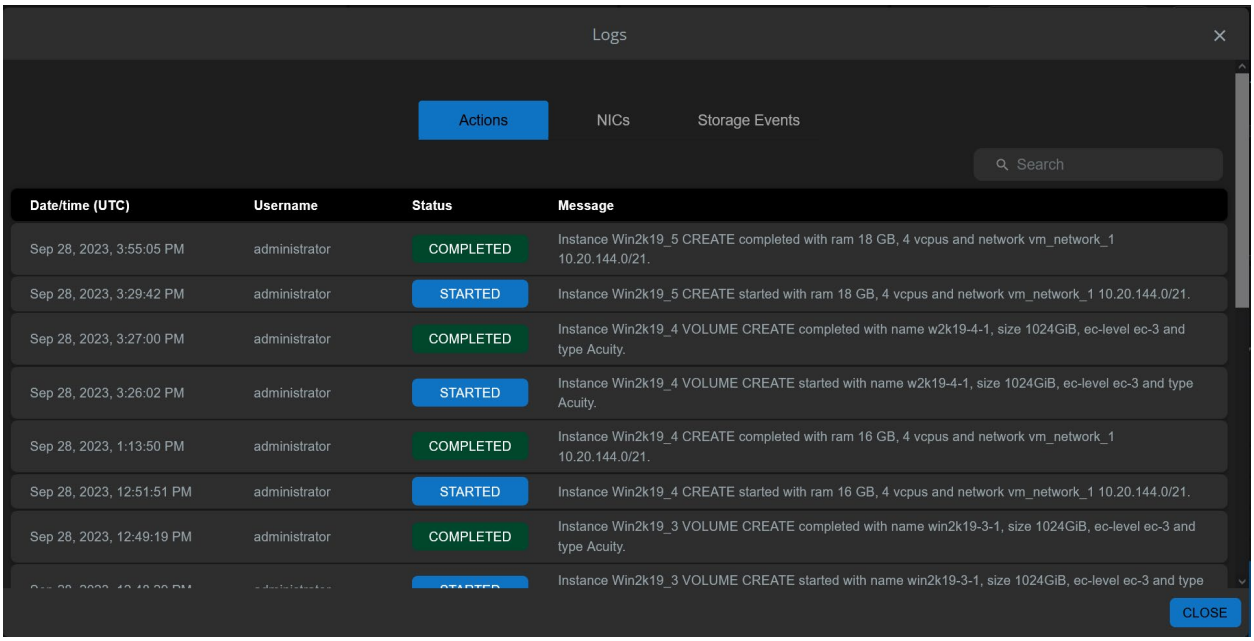


## *NICs*

Shows the physical NIC related events.

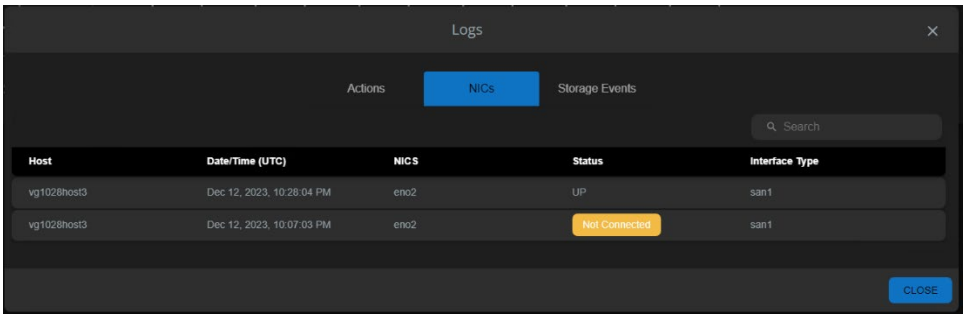

## *Acuity Storage Events*

Shows the Acuity backend-storage events.

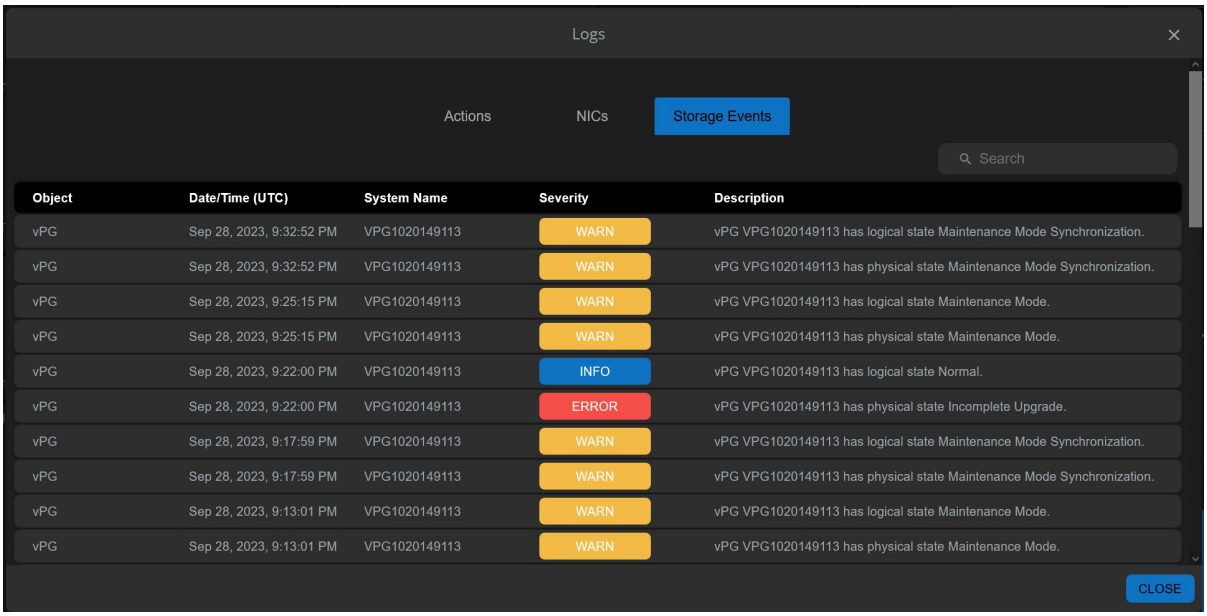


| Object | Date/Time (UTC) | System Name | Severity | Description |
|---|---|---|---|---|
| vPG | Sep 28, 2023, 9:32:52 PM | VPG1020149113 | WARN | vPG VPG1020149113 has logical state Maintenance Mode Synchronization. |
| vPG | Sep 28, 2023, 9:32:52 PM | VPG1020149113 | WARN | vPG VPG1020149113 has physical state Maintenance Mode Synchronization. |
| vPG | Sep 28, 2023, 9:25:15 PM | VPG1020149113 | WARN | vPG VPG1020149113 has logical state Maintenance Mode. |
| vPG | Sep 28, 2023, 9:25:15 PM | VPG1020149113 | WARN | vPG VPG1020149113 has physical state Maintenance Mode. |
| vPG | Sep 28, 2023, 9:22:00 PM | VPG1020149113 | INFO | vPG VPG1020149113 has logical state Normal. |
| vPG | Sep 28, 2023, 9:22:00 PM | VPG1020149113 | ERROR | vPG VPG1020149113 has physical state Incomplete Upgrade. |
| vPG | Sep 28, 2023, 9:17:59 PM | VPG1020149113 | WARN | vPG VPG1020149113 has logical state Maintenance Mode Synchronization. |
| vPG | Sep 28, 2023, 9:17:59 PM | VPG1020149113 | WARN | vPG VPG1020149113 has physical state Maintenance Mode Synchronization. |
| vPG | Sep 28, 2023, 9:13:01 PM | VPG1020149113 | WARN | vPG VPG1020149113 has logical state Maintenance Mode. |
| vPG | Sep 28, 2023, 9:13:01 PM | VPG1020149113 | WARN | vPG VPG1020149113 has physical state Maintenance Mode. |

## Viewing Current Issues

If there is a problem with the cluster, you can launch the **Issues** dialog to see all issues in one location. Press **View Issues** on the dashboard to launch the Issues dialog.

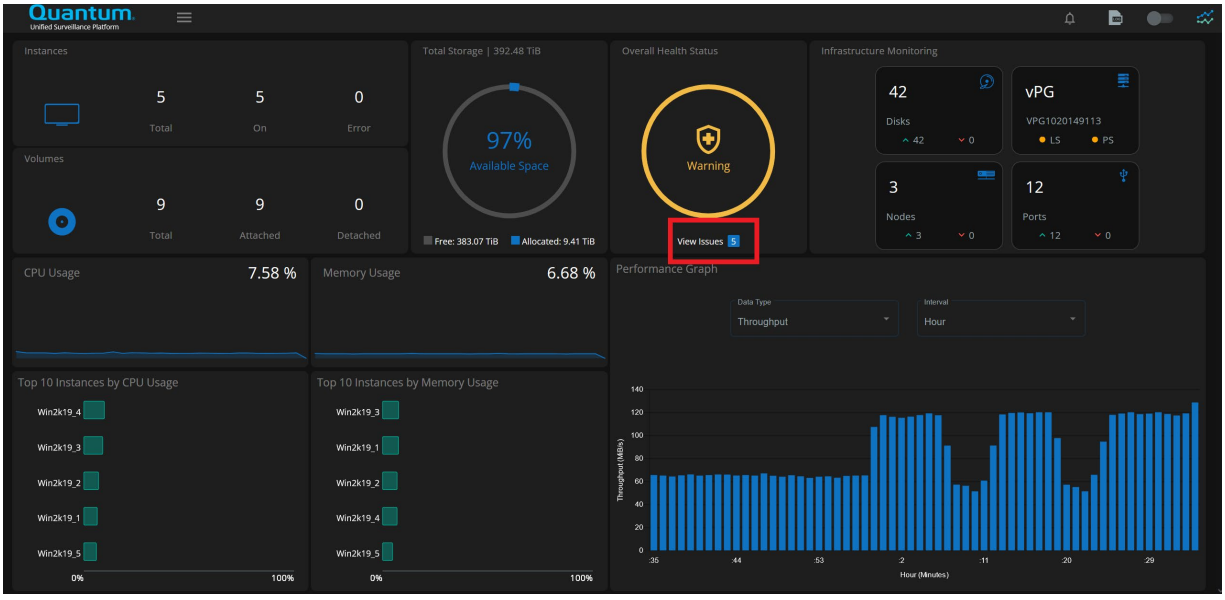

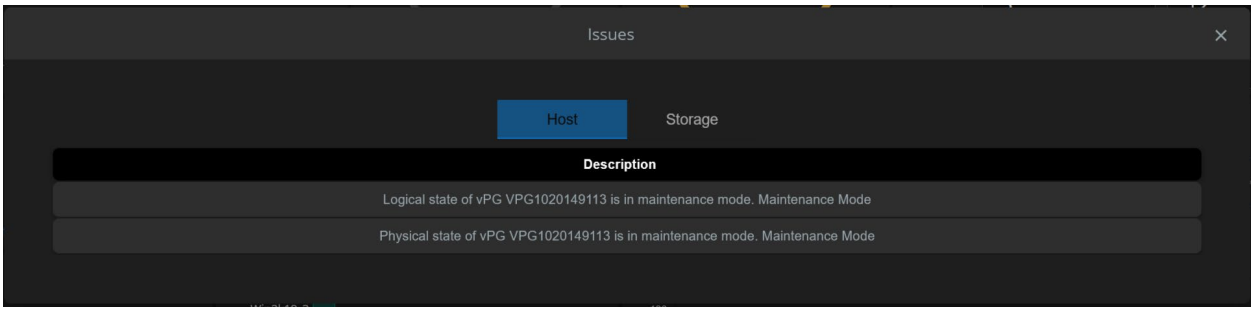


## Upgrade USP

Upgrading to USP 5.1.0 or later is a two-step process.

1. Upgrade using the interim package named upgrade_package.tgz.
2. Upgrade using the USP package named USP_Upgrade_<version>-<build>.zip

Follow these instructions:

1. Extract the USP upgrade package zip file on your local system. It will contain two packages:
   - USP_Upgrade_<version>-<build>.tgz
   - upgrade_package.tgz
2. Launch the upgrade dialog using the navigation bar.

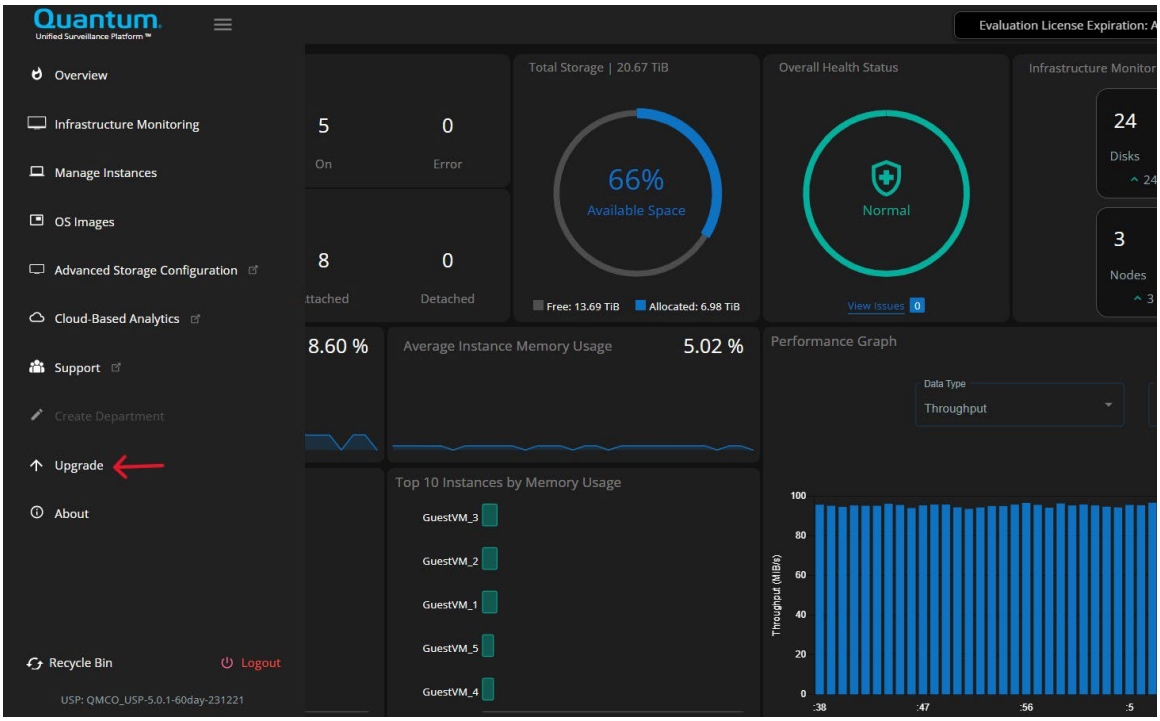


3. Click on **Upgrade** on the top right and upload the upgrade package named upgrade_package.tgz.

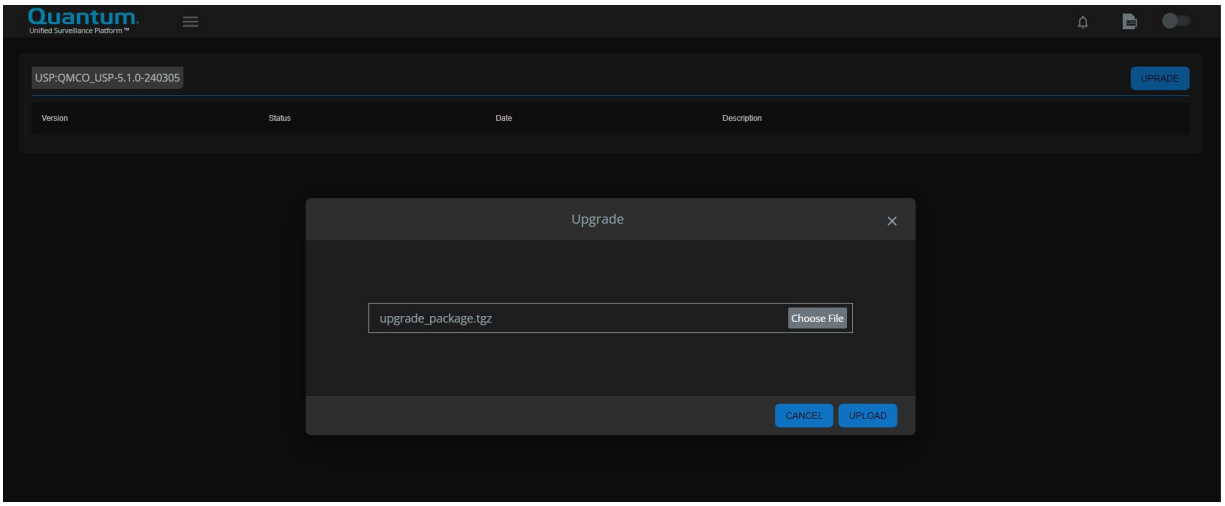


4. Enter the server password on next screen and click on **Confirm.**

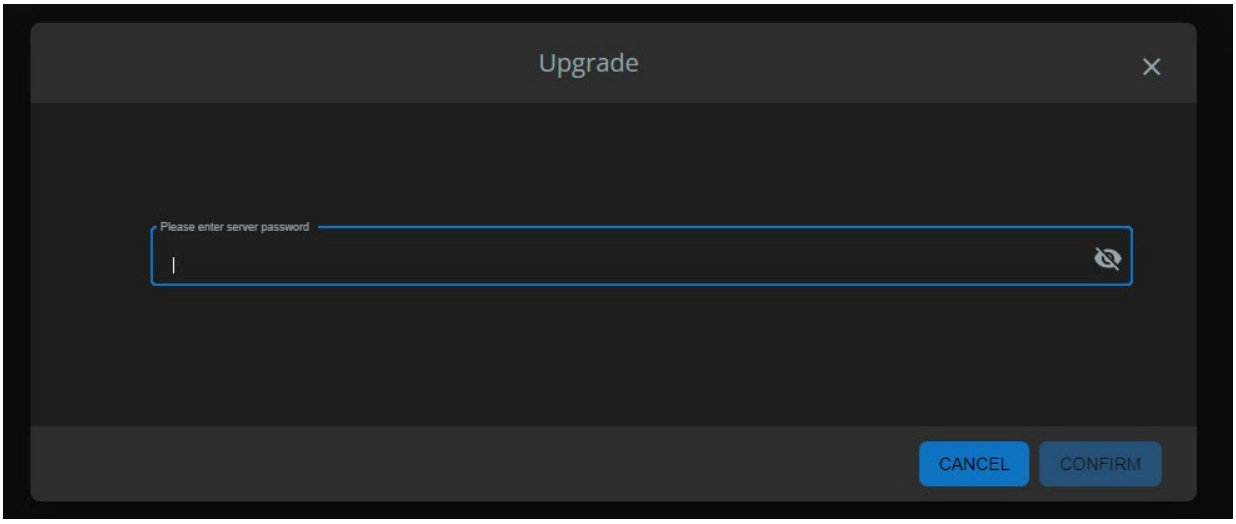


5. The progress dialog indicates that the upgrade is in progress. The dialog will disappear when the upgrade is complete.

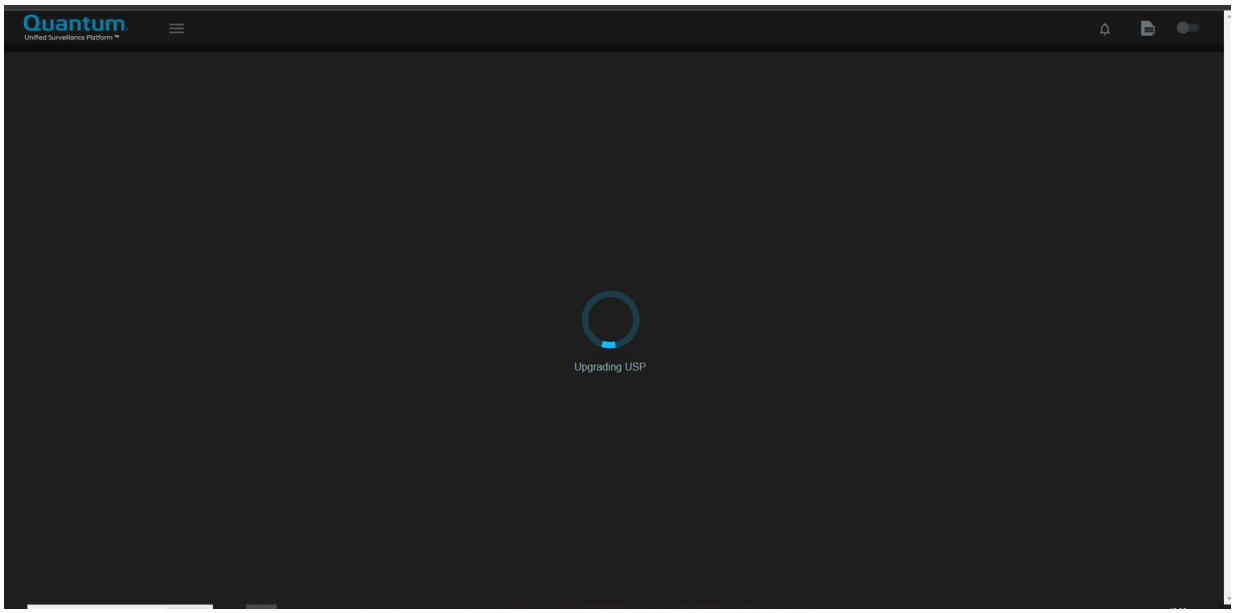


6. Once the upgrade completes. you will need to hard refresh the browser (Ctrl+F5).
7. Repeat steps 1-6 using the USP_Upgrade_<version>-<build>.tgz package.
8. Once the upgrade completes, you will need to hard refresh the browser (Ctrl+F5).
9. Navigate back to the upgrade dialog to view the results of the upgrade.

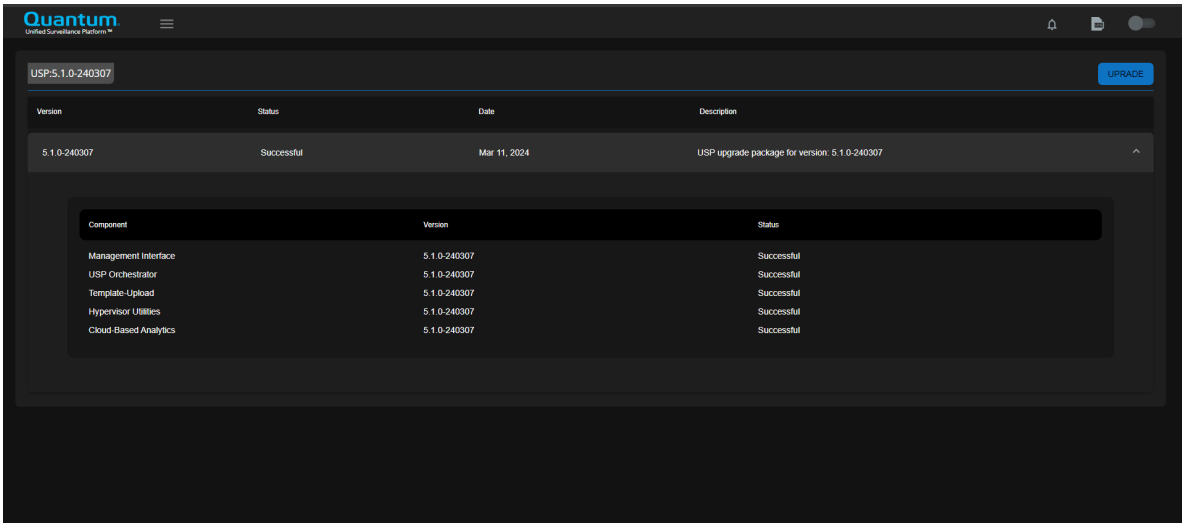

## Launching the Acuity Advanced Storage Configuration Utility

To launch the Acuity Advanced Storage Configuration Utility, navigate to the USP Menu and select Advances Storage Configuration.

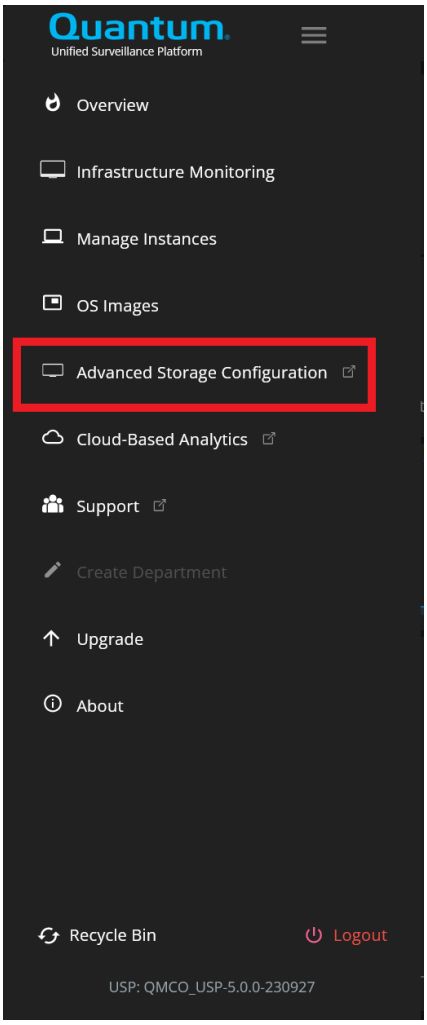


Log on to the Advanced Storage Configuration Utility using the same credentials you used to access the Quantum USP Management Application.

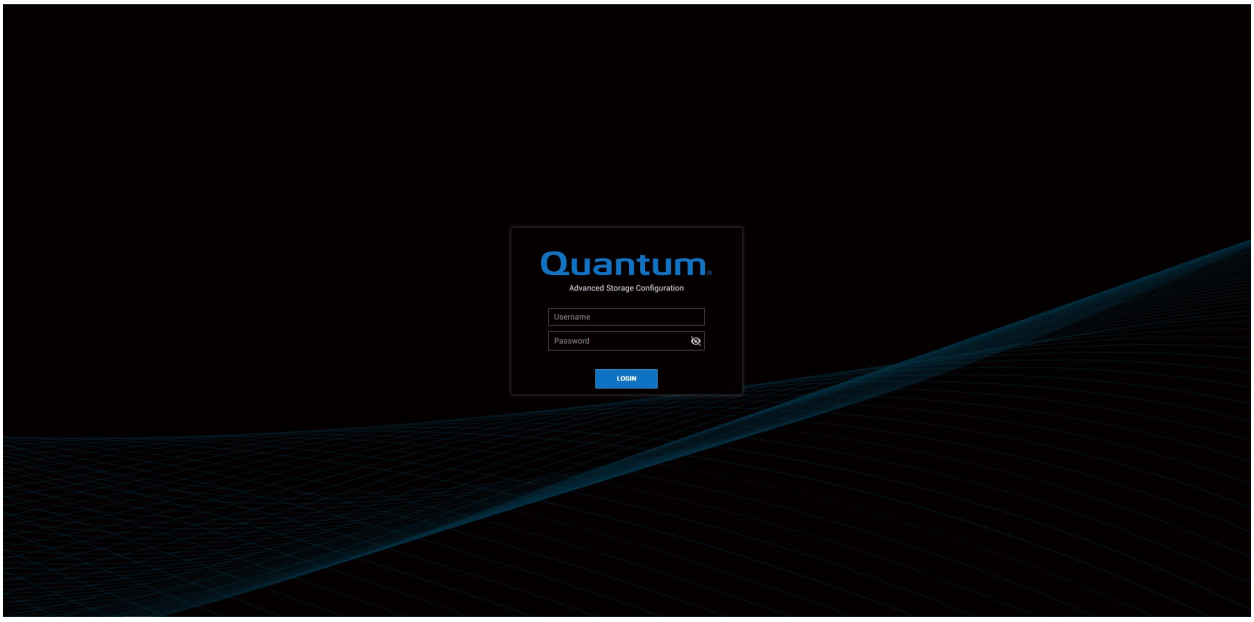
Quantum
Advanced Storage Configuration
Username
Password
LOGIN

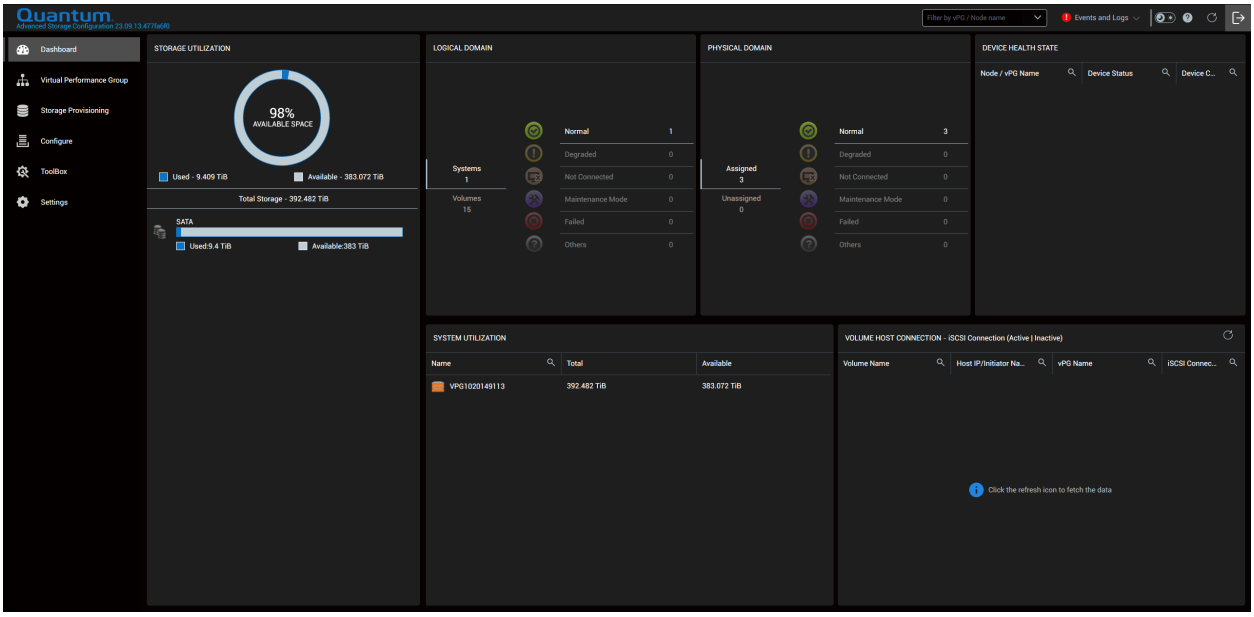
Quantum
Advanced Storage Configuration 23.09.13.477fa6f0
Filter by vPG / Node name
Events and Logs
Dashboard
Virtual Performance Group
Storage Provisioning
Configure
ToolBox
Settings
STORAGE UTILIZATION
98%
AVAILABLE SPACE
Used - 9.409 TiB
Available - 383.072 TiB
Total Storage - 392.482 TiB
SATA
Used:9.4 TiB
Available:383 TiB
LOGICAL DOMAIN
Systems
1
Volumes
15
Normal 1
Degraded 0
Not Connected 0
Maintenance Mode 0
Failed 0
Others 0
PHYSICAL DOMAIN
Assigned
3
Unassigned
0
Normal 3
Degraded 0
Not Connected 0
Maintenance Mode 0
Failed 0
Others 0
DEVICE HEALTH STATE
Node / vPG Name
Device Status
Device C...
SYSTEM UTILIZATION
Name
Total
Available
VPG1020149113
392.482 TiB
383.072 TiB
VOLUME HOST CONNECTION - iSCSI Connection (Active | Inactive)
Volume Name
Host IP/Initiator Na...
vPG Name
iSCSI Connec...
Click the refresh icon to fetch the data

## Viewing the Acuity VM Console

1. Log on to the USP Host's CentOS console.

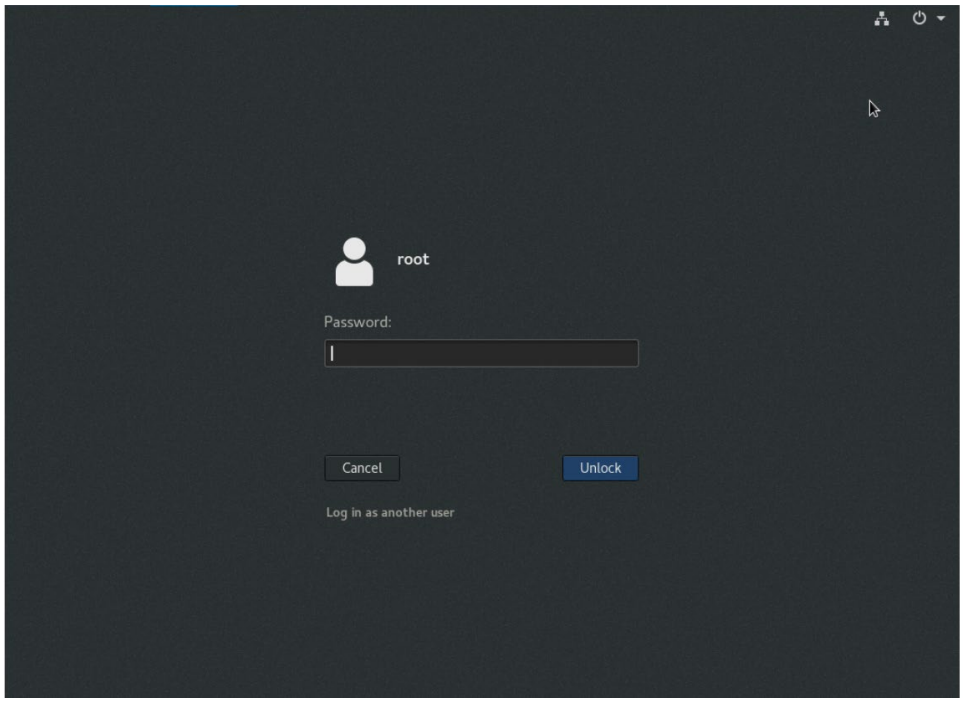


2. Click the **Activities** button, located in the upper left-hand corner.
3. Search for **Virtual Machine Manager** and click on the icon to launch.

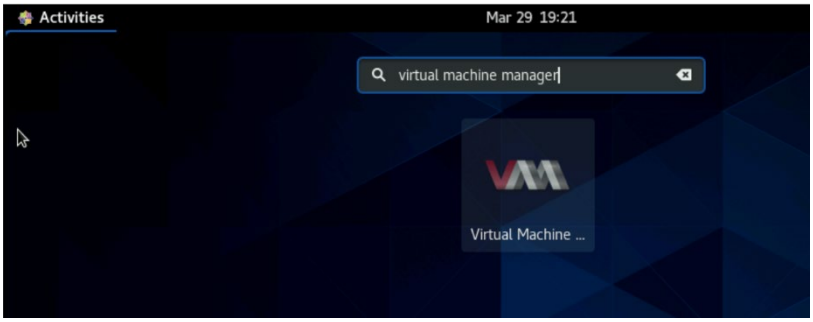

4. Select the **Acuity-VM** instance and click on **Open** to launch the console.

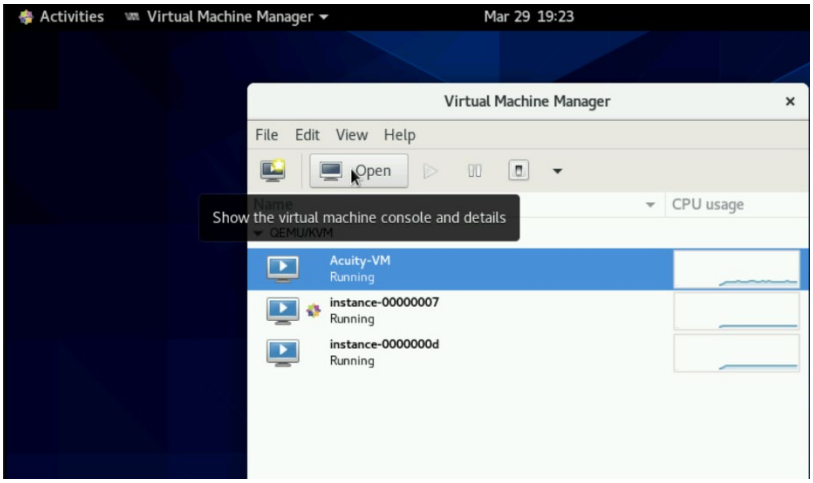


## Cloud-Based Analytics (CBA) Portal

The USP product will automatically report to Quantum's CBA portal. Log in to the Quantum CBA Portal UI from any browser and type in your user details and credentials. If you do not have access to the portal, request access through the Quantum CBA Portal UI by clicking the ***Request Site Access*** link.

Once you log in, the CBA portal UI displays the solutions and nodes that are mapped to your account. The portal can also display your account information and association token.

The Quantum CBA portal displays the options to view **Solutions** and **Nodes**. To monitor the VS-NVR servers, click **Nodes**.


The list of nodes that are associated with your account along with some basic information is displayed.

| Health | Customer | Node | Family | Model | Version | Serial | Solution | Actions |
|---|---|---|---|---|---|---|---|---|
| | Quantum Customer | dev-snc-daiquiri-n2 | Quantum StorNext | StorNext Software | 7.0.1 | SV1535CKH00011 | Denver - StorNext | |
| | Quantum Customer | kvm2-esosvm1-b | Quantum F-Series | F2000 | 1.0.0 | USWSJ00918OA... | Denver F-Series | |
| | Quantum Customer | kvm3-vm3-f1k-pr-a | Quantum F-Series | F1000 | 1.0.0 | 1.0.0-Build7... | Denver F-Series | |
| | Quantum Customer | kvm5-esosvm2-a | Quantum F-Series | F2000 | 1.0.0 | USWSJ00918OA... | Denver F-Series | |
| | Quantum Customer | kvm5-esosvm2-b | Quantum F-Series | F2000 | 1.0.0 | USWSJ00918OA... | Denver F-Series | |
| | Quantum Customer | kvm5-esosvm3-a | Quantum F-Series | F2000 | 1.0.0 | USWSJ00918OA... | Denver F-Series | |
| | Quantum Customer | kvm5-esosvm3-b | Quantum F-Series | F2000 | 1.0.0 | USWSJ00918OA... | Denver F-Series | |

You can filter the nodes by name, model, version, and so on. Select the node that you want to monitor.

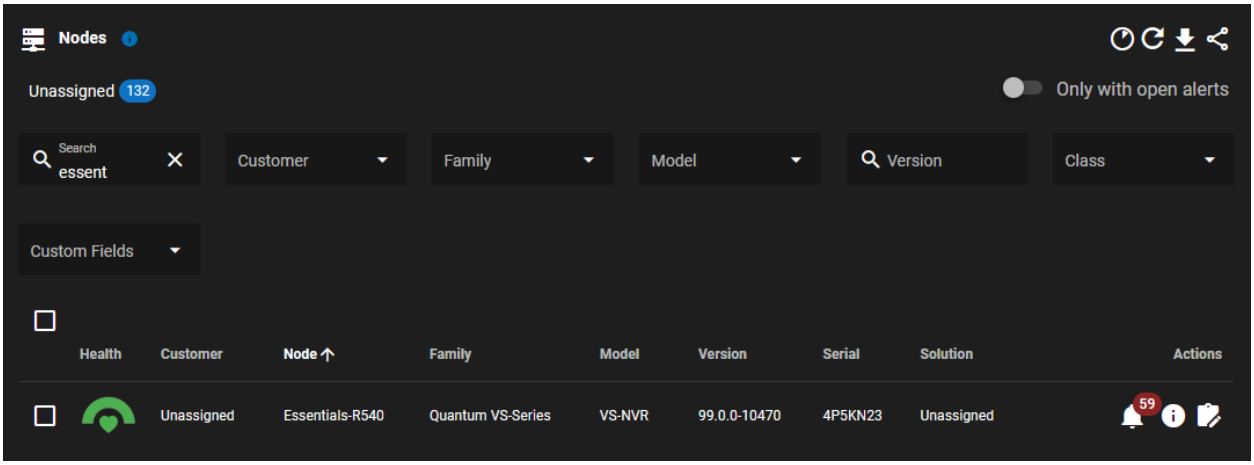


Once you clicked the required node, you can see the following information about the node.

- Overview
- Alerts
- Performance:
    - System Resources
    - System Interfaces
    - Storage Disk
    - VMS
- History: Reports

## *Overview*

This section describes the overall status and basic information about the node.

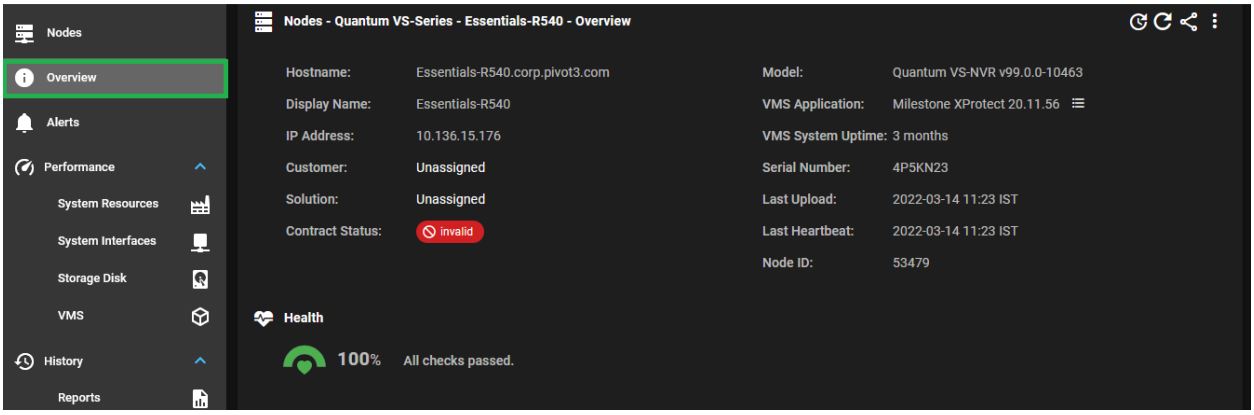

The **Overview** section also lists the VMS applications that are installed in the VS-NVR system.

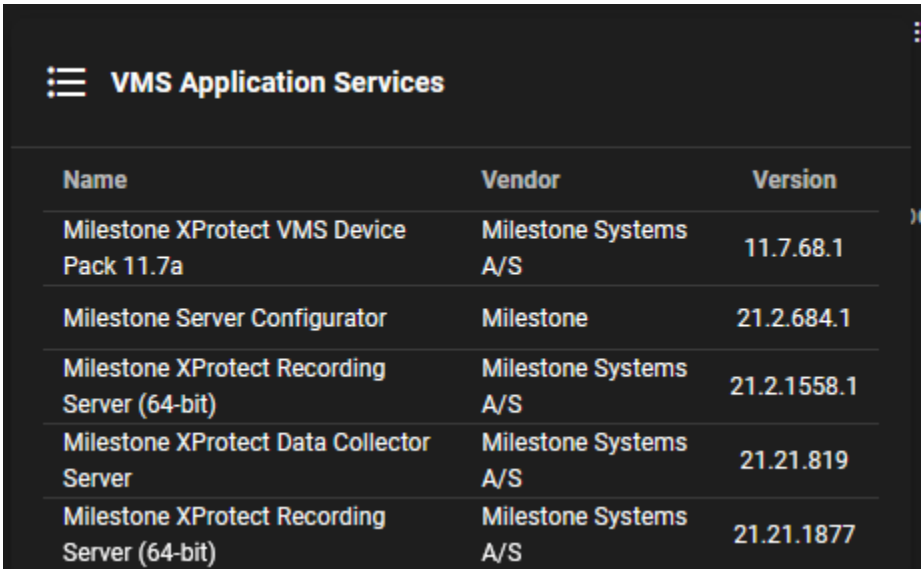

VMS Application Services

| Name | Vendor | Version |
|---|---|---|
| Milestone XProtect VMS Device Pack 11.7a | Milestone Systems A/S | 11.7.68.1 |
| Milestone Server Configurator | Milestone | 21.2.684.1 |
| Milestone XProtect Recording Server (64-bit) | Milestone Systems A/S | 21.2.1558.1 |
| Milestone XProtect Data Collector Server | Milestone Systems A/S | 21.21.819 |
| Milestone XProtect Recording Server (64-bit) | Milestone Systems A/S | 21.21.1877 |

## *Alerts*

The **Alerts** section displays the events that are specific to the selected node. The Alert table lists each event with the time it occurred, message, severity of the event, and current status of the event. The status message displays whether the alert has been acknowledged.

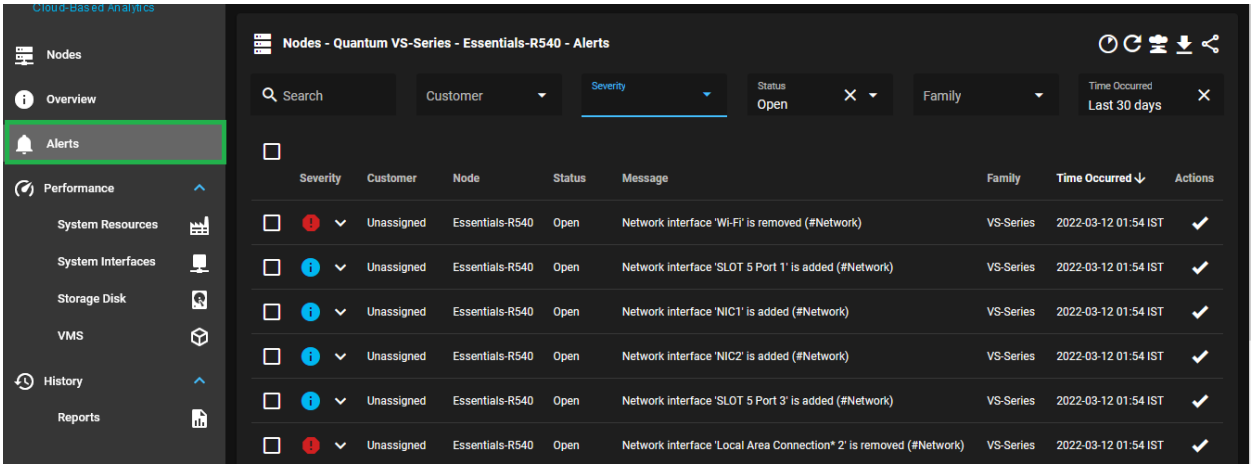


## *Performance*

The **Performance** metrics display historical data about various system resources, such as CPU, Memory, Network, and Storage Disk.

**System Resources**

The **System Resources** option displays the CPU (Max and Average) and Memory (Bytes and utilization) performance charts.

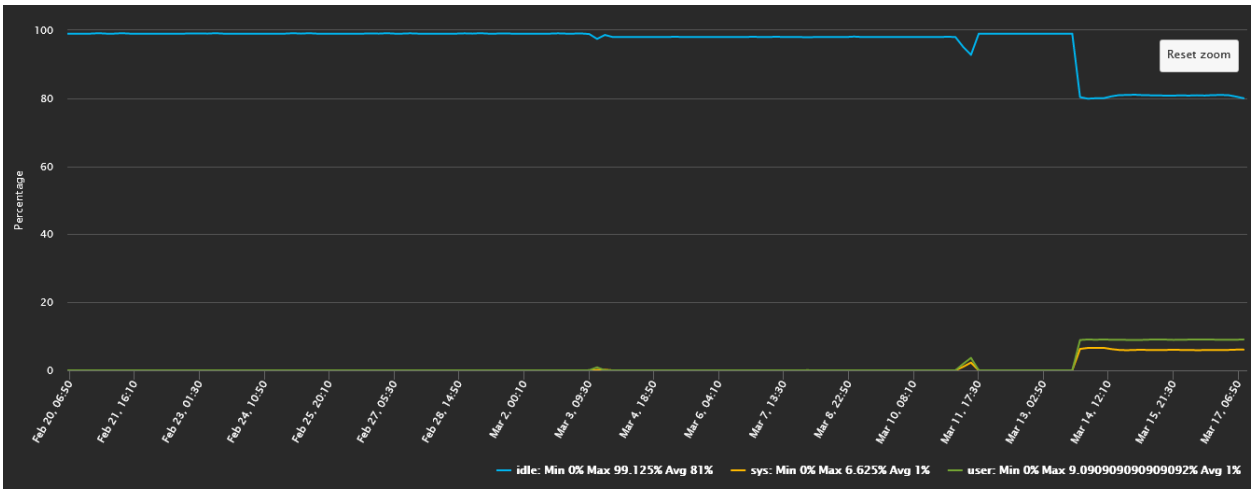


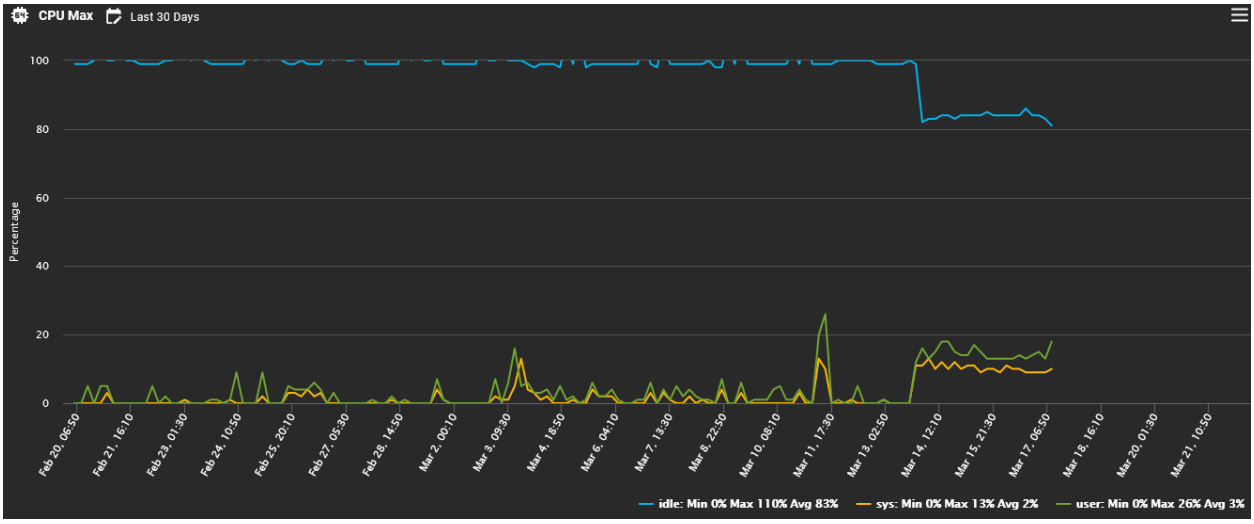

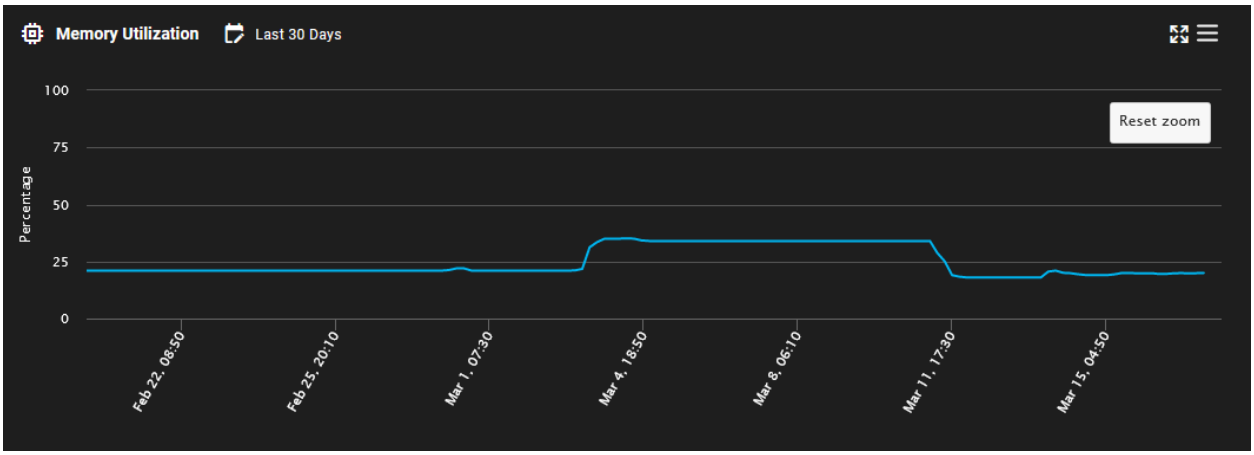
Memory Utilization
Last 30 Days
Reset zoom
Percentage
100
75
50
25
0
Feb 22, 08:50
Feb 25, 20:10
Mar 1, 07:30
Mar 4, 18:50
Mar 8, 06:10
Mar 11, 17:30
Mar 15, 04:50

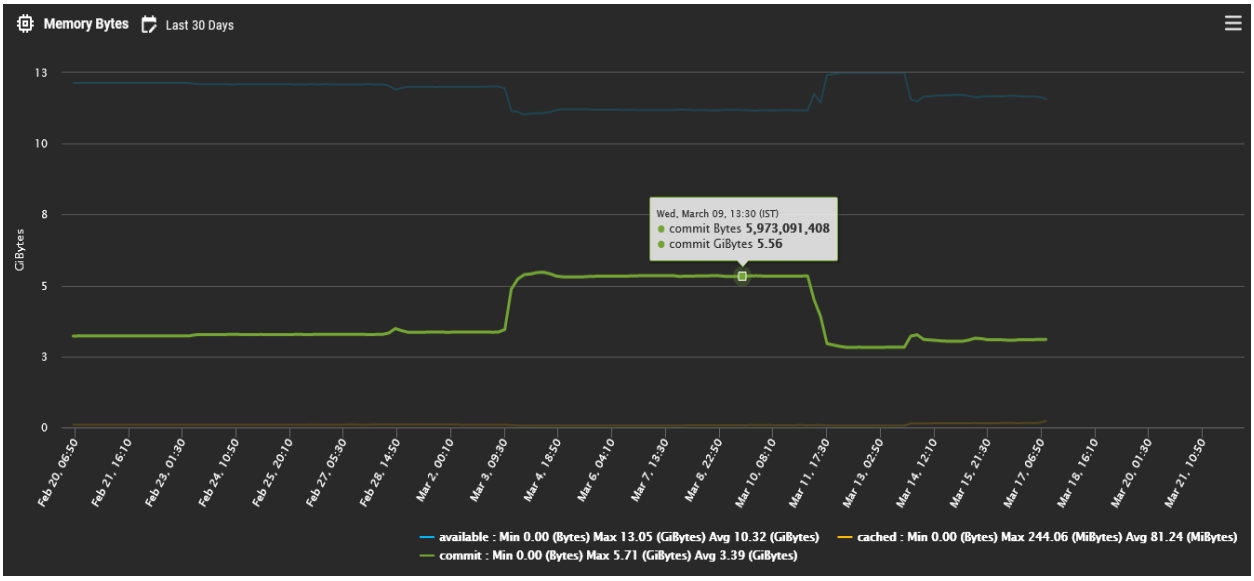
Memory Bytes
Last 30 Days
GiBytes
13
10
8
5
3
0
Wed, March 09, 13:30 (IST)
commit Bytes 5,973,091,408
commit GiBytes 5.56
Feb 20, 06:50
Feb 21, 16:10
Feb 23, 01:30
Feb 24, 10:50
Feb 25, 20:10
Feb 27, 05:30
Feb 28, 14:50
Mar 2, 00:10
Mar 3, 09:30
Mar 4, 18:50
Mar 6, 04:10
Mar 7, 13:30
Mar 8, 22:50
Mar 10, 08:10
Mar 11, 17:30
Mar 13, 02:50
Mar 14, 12:10
Mar 15, 21:30
Mar 17, 06:50
Mar 18, 16:10
Mar 20, 01:30
Mar 21, 10:50
available : Min 0.00 (Bytes) Max 13.05 (GiBytes) Avg 10.32 (GiBytes)
cached : Min 0.00 (Bytes) Max 244.06 (MiBytes) Avg 81.24 (MiBytes)
commit : Min 0.00 (Bytes) Max 5.71 (GiBytes) Avg 3.39 (GiBytes)

**System Interfaces**

The **System Interface Throughput** option displays charts related to network performance, such as network throughput and packets of each network interface.

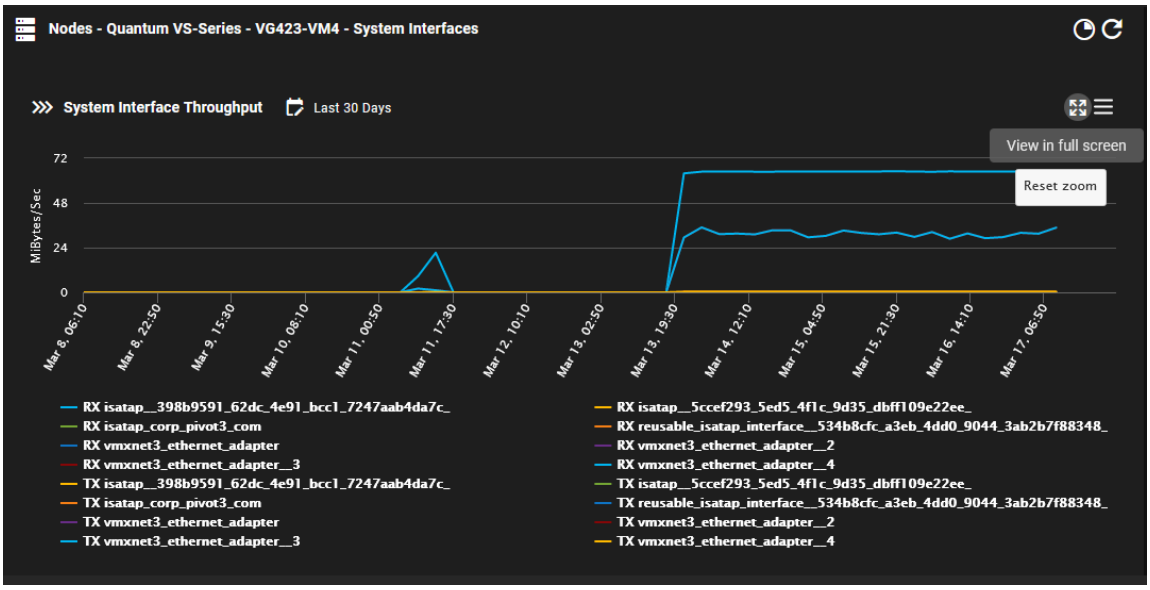


**Storage Disk**

This **Storage Disk** option provides you with the historical disk performance (disk throughput and utilization) in chart view.

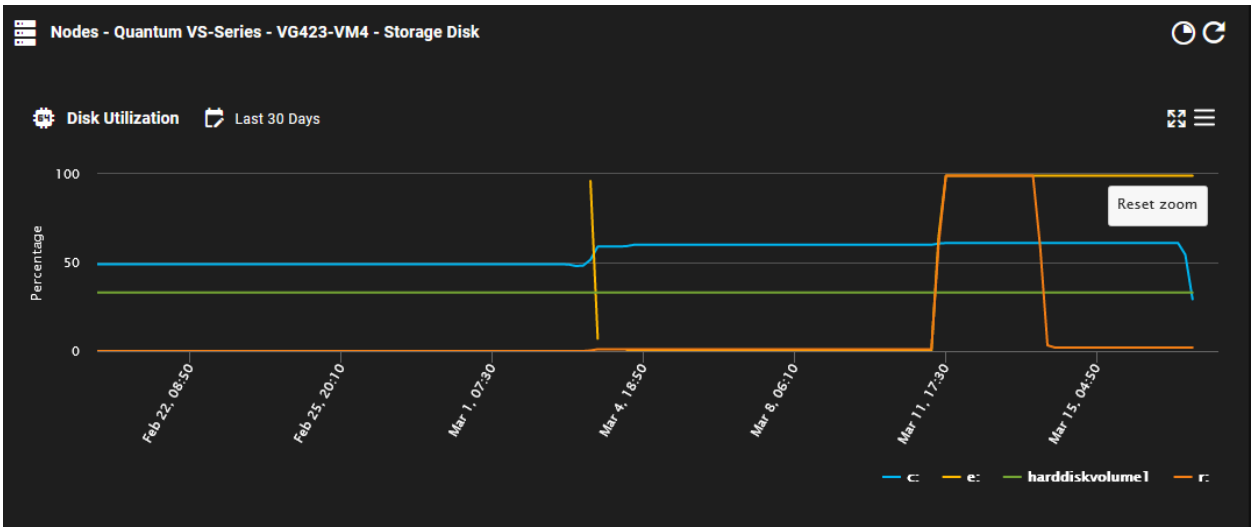


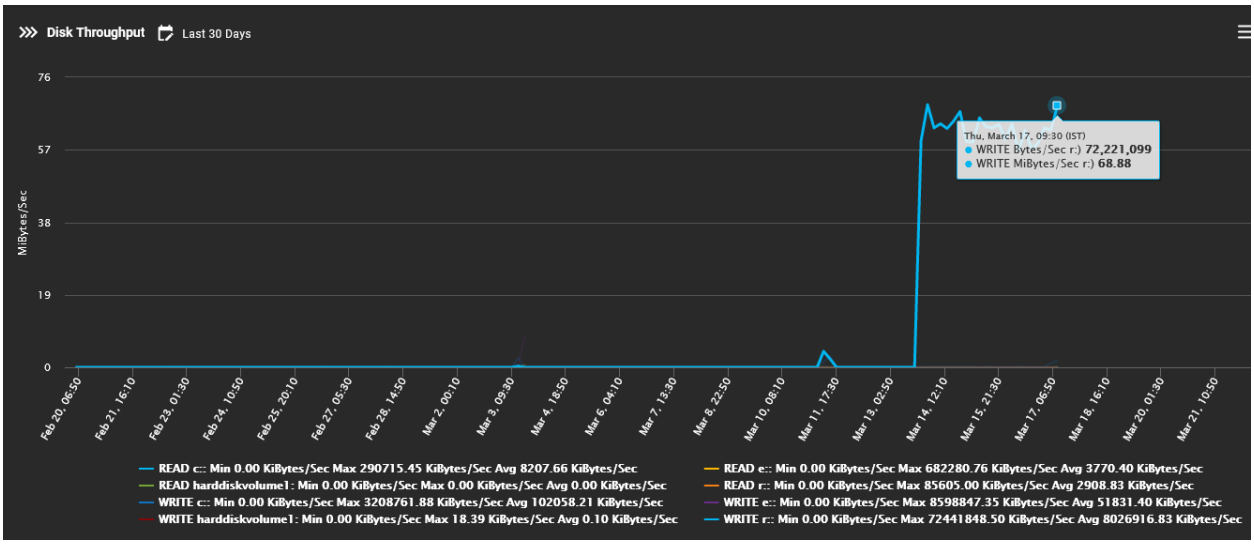


## User Association Token

You use the Association token to associate servers/agents to your account. To associate a server/agent to your account, provide your account Association token during USP installation. To receive an Association token, log in to the CBA portal and click the ***user*** icon, as highlighted in the following image.

When you click the ***user*** icon, the following menu is displayed with the appropriate options. Click **Association Token** for the Association token for your account.

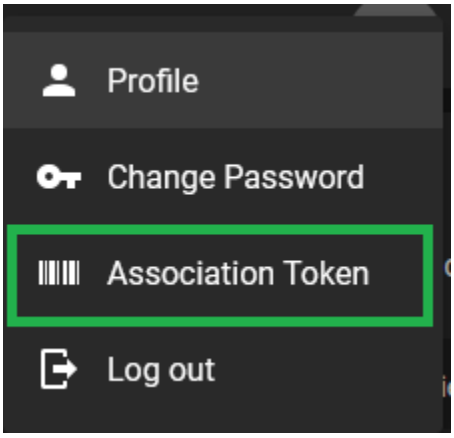


Your unique Association token is displayed in a pop-up dialog box. Click **Copy to Clipboard** to copy the Association token to your clipboard.

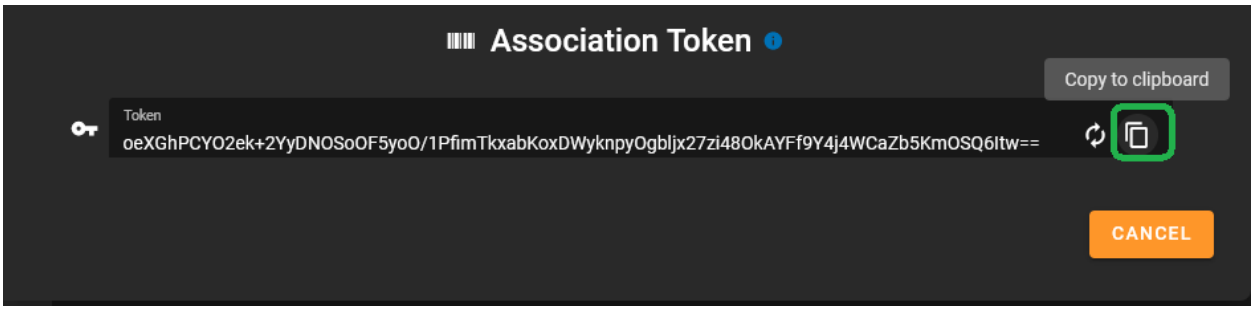


You can use the Association token during the USP installation.

## Maintenance Mode

You can put a node of the cluster in maintenance mode. It advisable to put a node in maintenance mode before doing any hardware operations, or any other maintenance activity.

Use the following commands to put the node in maintenance mode:

- cd /root/enclouden/dev/stack_orchestrator/
- Enter maintenance mode.
    - python3 maintenance_mode -a enter
- Exit maintenance mode.
    - python3 maintenance_mode -a exit

# Cluster Management

## Adding a Node

User can add new nodes to an existing cluster. Follow the below steps:

### *Pre-Requisites*

1. Install the same version of the USP5.0 iso that is running on the existing cluster.
2. Complete the ISO installation and assign the management IP to the server on the preferred NIC.

### *Node Addition Template*

1. You can edit the same template file that you used for the Deployment of this cluster.
2. Remove the node details from the existing template file and add the details of only the new nodes you are planning to add to the cluster.

### *Adding the New Node(s) to the cluster*

1. Open the Node Addition UI. You can open the Node Addition UI from a browser using the IP of the first node of the cluster. The URL is: http://<IP of 1st node>/#/node-add

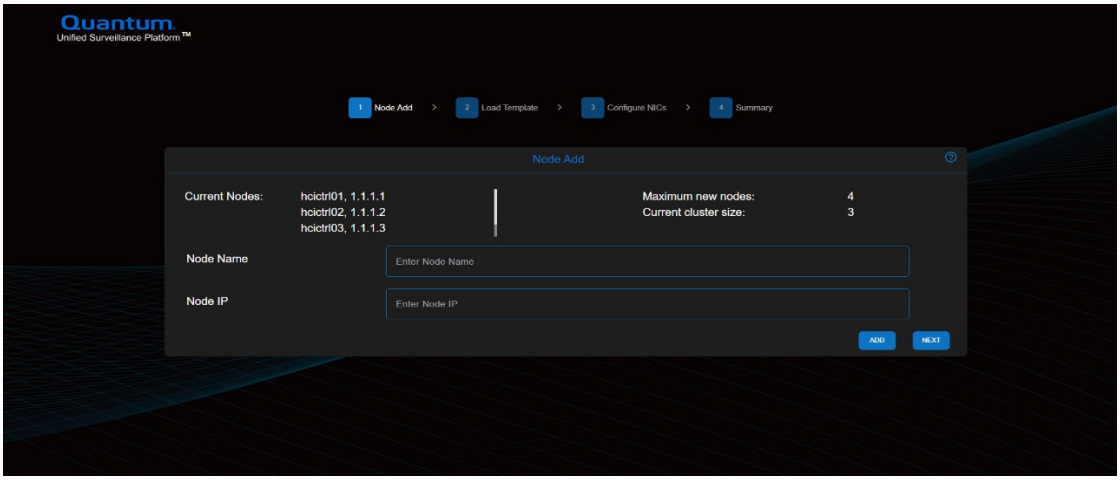


2. From the “Node Add” page.
   a. This page shows the existing nodes in the cluster and their IP addresses. **“Maximum new nodes”** is the number of nodes that can be added to the cluster. “Current cluster size” is the number of nodes present in the cluster.
   b. Enter the Hostname and IP of the new node. Press **Add** to add the node to the cluster.
   c. Repeat the steps to add more nodes. Once details of all new nodes are added, Click on **NEXT**.
3. From the “Load Template” page.
   a. Now you can load the template. The “Load Template” page allows you to upload the configuration template containing information about the new nodes. Click **Next** to proceed to the next step.

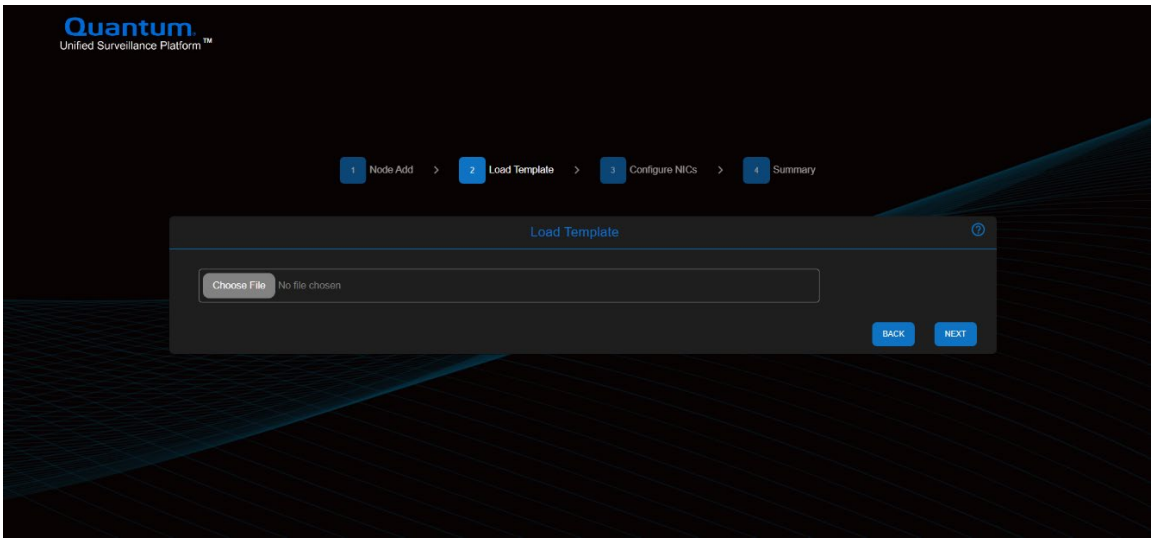


4. From the “Configure NICs” page.
    a. Configure the logical to physical mapping for each network of the new nodes. Press **Identify Nic** to blink the LED of the selected NIC.

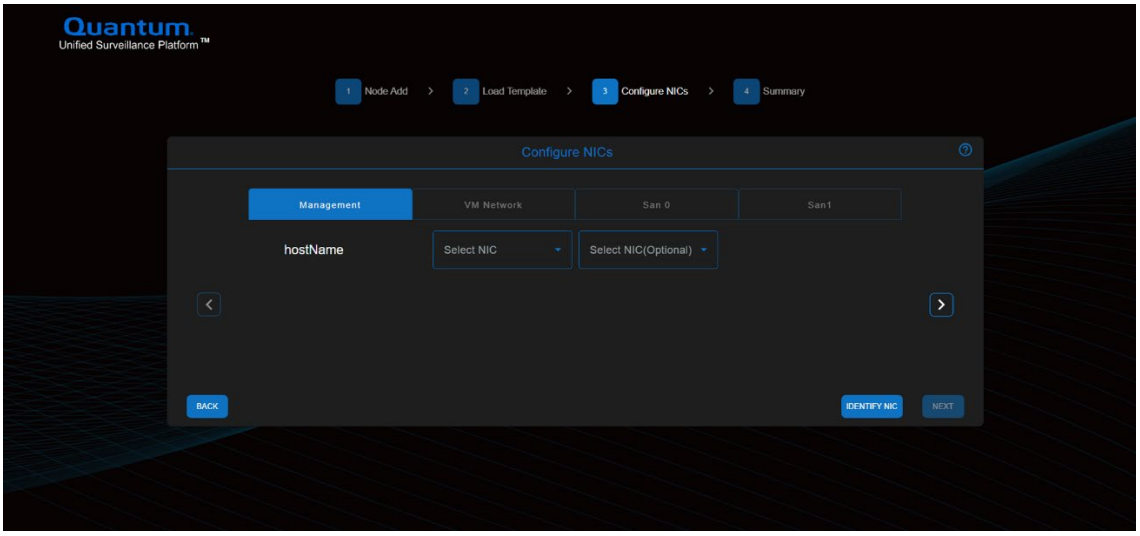


    b. Click on **NEXT** once all details are given.
5. From the “Summary” page.
    a. Confirm that all of the details are correct. Press **Back** to make changes. Press **Download** to save a copy of the summary for future reference.

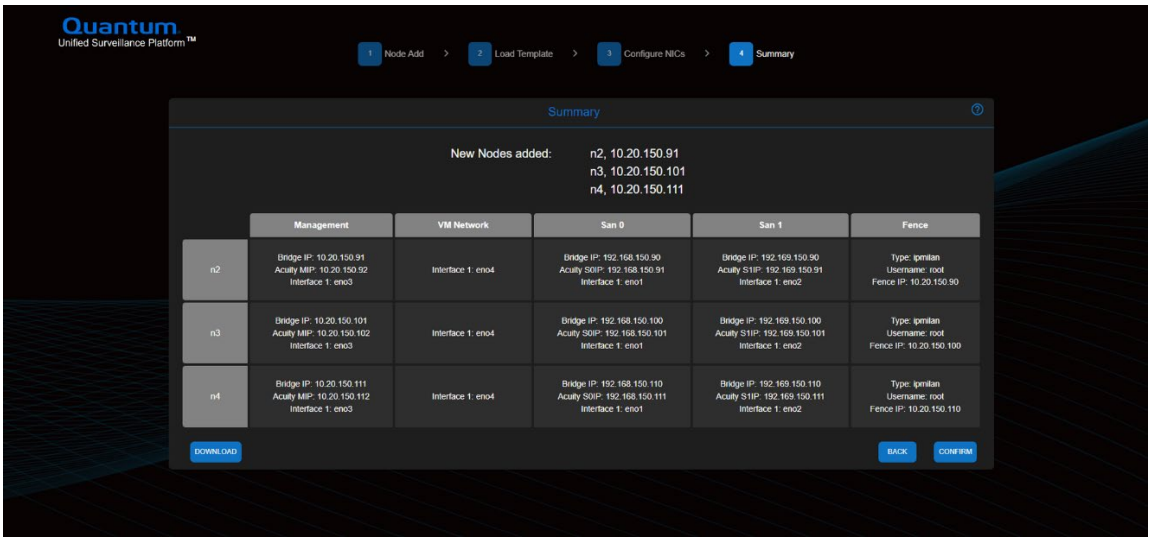


b. Once you verify the details, press **Confirm**.

c. Click on **Confirm** on the popup screen to start the deploy process.

6. The log/progress of the steps during the deployment displays.

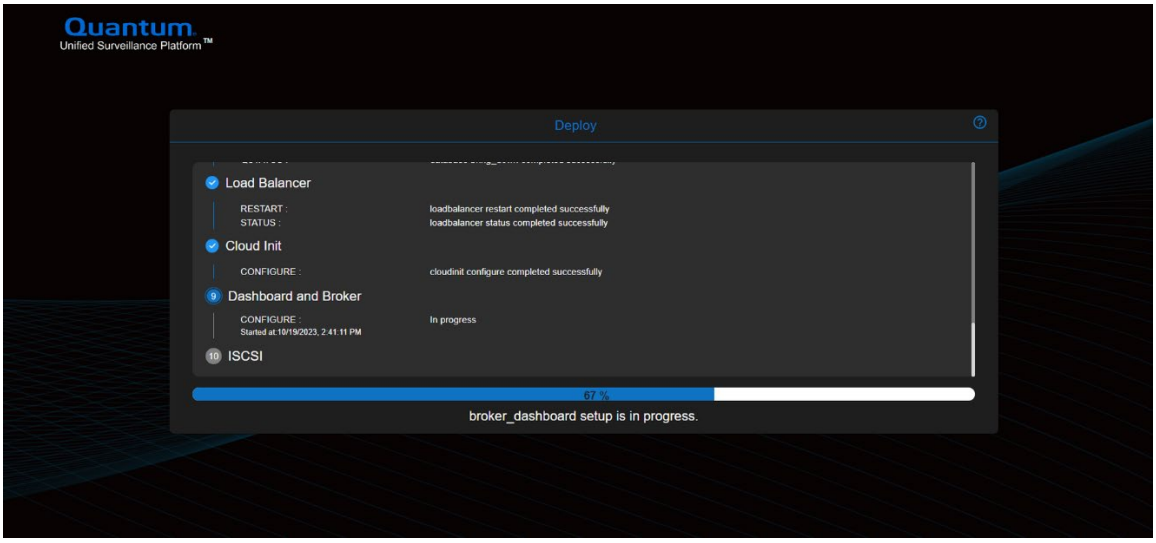


7. The following page displays to indicate that the Node Addition has completed. You can click on the provided link to go to USP 5.0 dashboard.

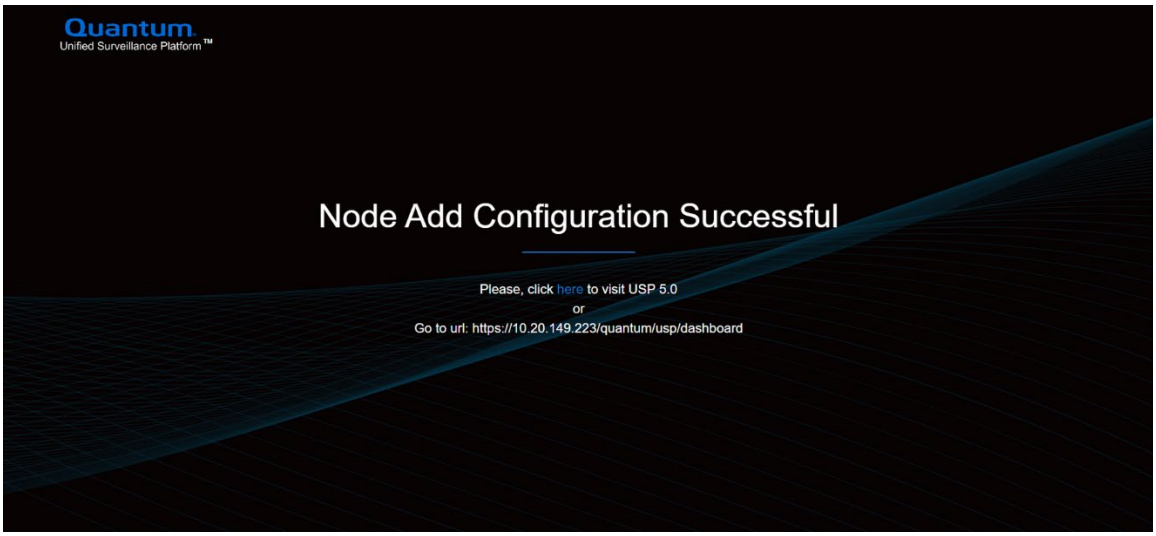


## Replacing a Node

You can replace a failed node in an existing cluster.

### *Pre-Requisites*

1. Install the same version of the USP5.0 iso that is running on the existing cluster.
2. Complete the ISO installation and assign the management IP to the new node on the preferred NIC. (Assuming the node to be replaced has failed and is Powered OFF.)

### *Node Addition Template*

1. User can use the Node Replacement Template file for this process.
2. Enter the details of the new node to the template file.

### *Replacing with New Node*

1. Open Node Replacement UI. You can open the Node Replacement UI page from a browser using the IP of the first node of the cluster. The URL is: http://<IP of anode>/#/node-replace

CRITICAL: *It is recommended to open UI on the 1st node. But if you are replacing the 1st node, you can open the UI on 2nd or 3rd node.* ***The UI should not be opened on the node getting replaced or the new node.***

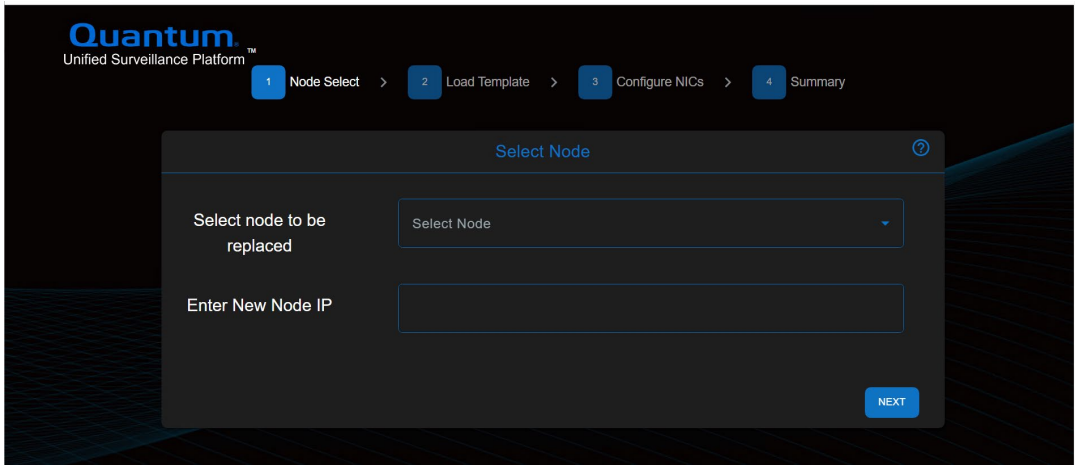


2. From the “Node Select” page.
    a. From the drop-down list in first field, select the failed node which needs to be replaced.
    b. In the second field, enter the management IP assigned to the new node.
    c. Click **NEXT**.

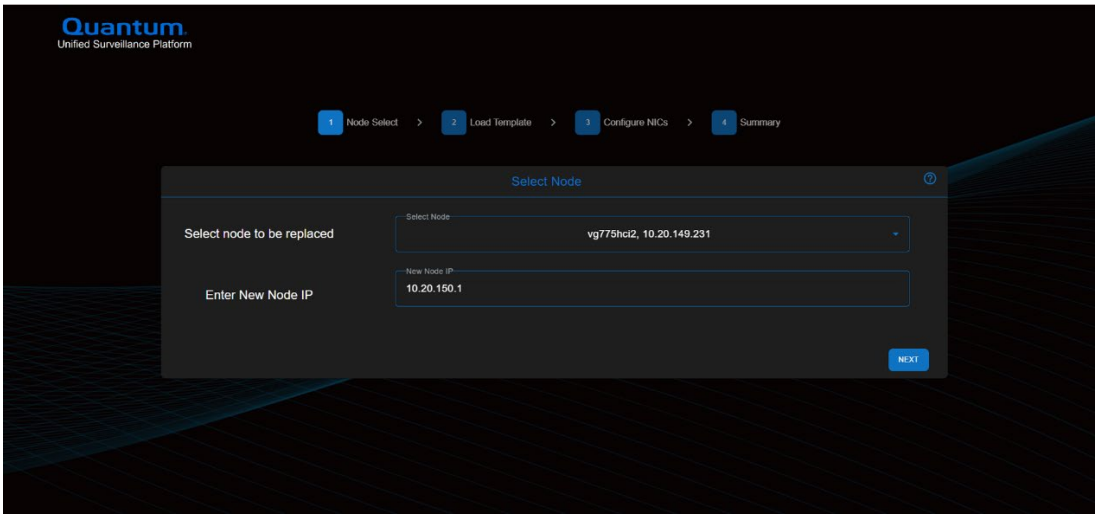


3. From the Upload Template page: You can upload the Excel file, which contains information about the new node, in this page. Then click **NEXT**.

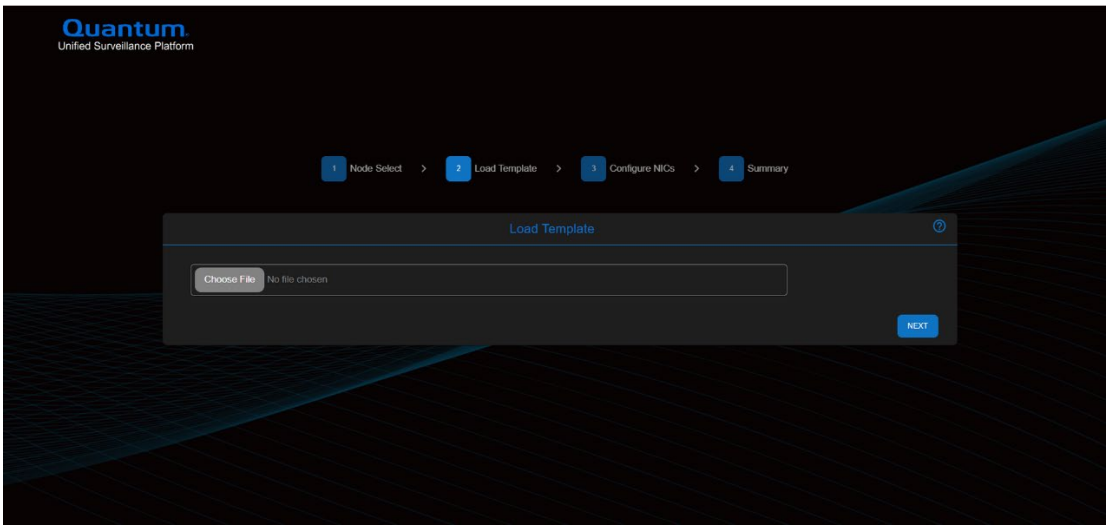


4. From the “Configure NICs” page: You can configure the NICs of the new node for all networks. You can also Identify the NIC by clicking on **IDENTIFY NIC** button, which will blink the LED for that NIC selected. Once you are done, Click **NEXT**.

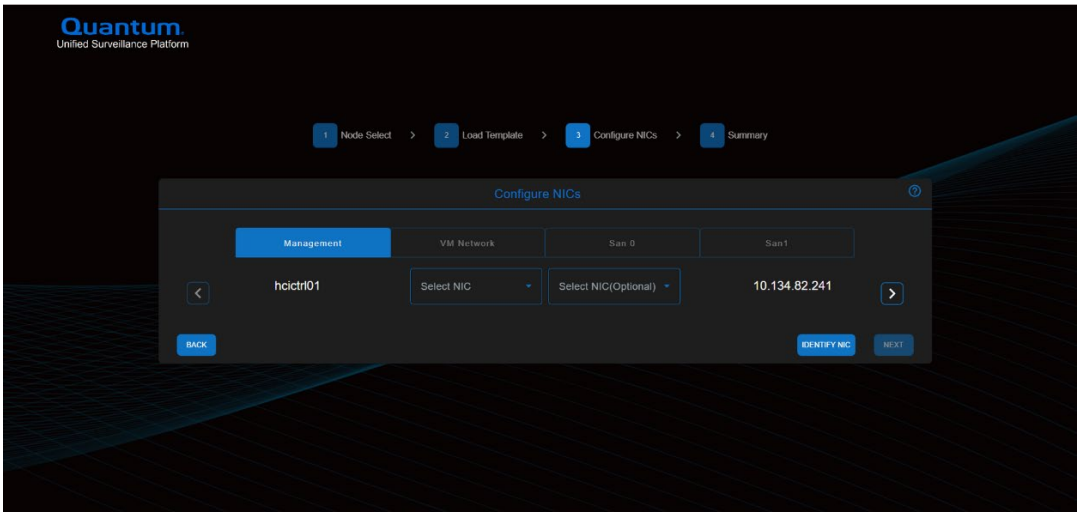


5. From the “Summary” page: You can see the summary of all the information uploaded until now. You can also download this summary in YAML format by clicking the Download button.

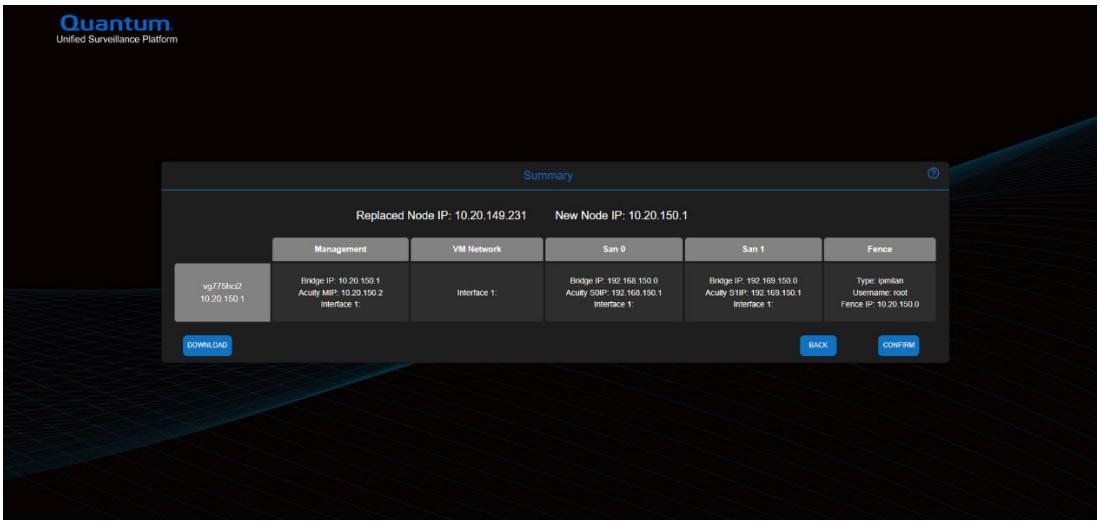


a) Once You verify the details, click the **CONFIRM** button.
b) Click on **Confirm** on the popup screen to start the deploy process.

6. The log/progress of the steps during the deployment displays.

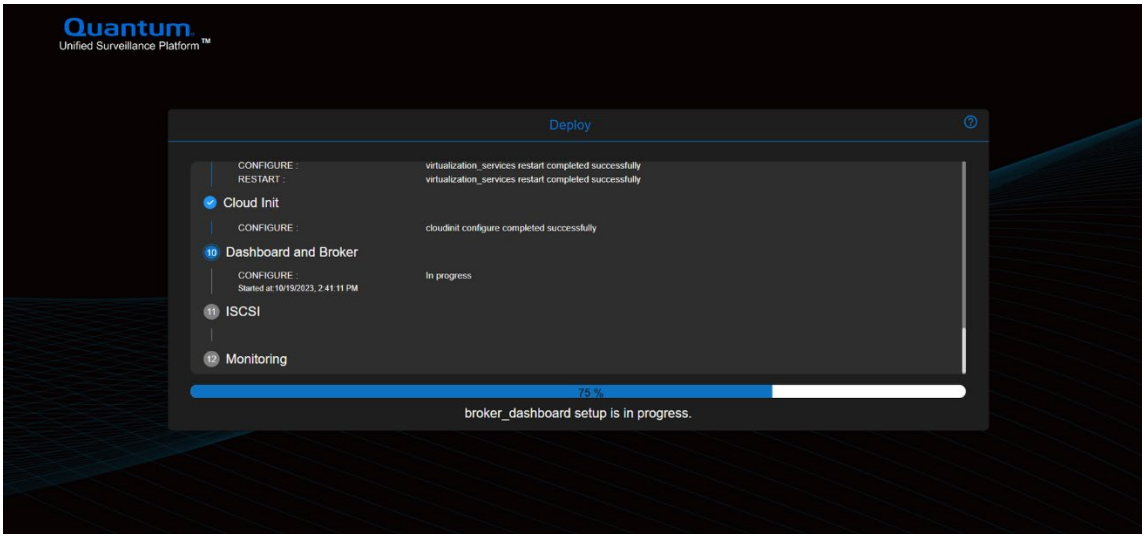


7. The following page displays to indicate the Node Addition has completed. You can click on the provided link to go to USP 5.0 dashboard.

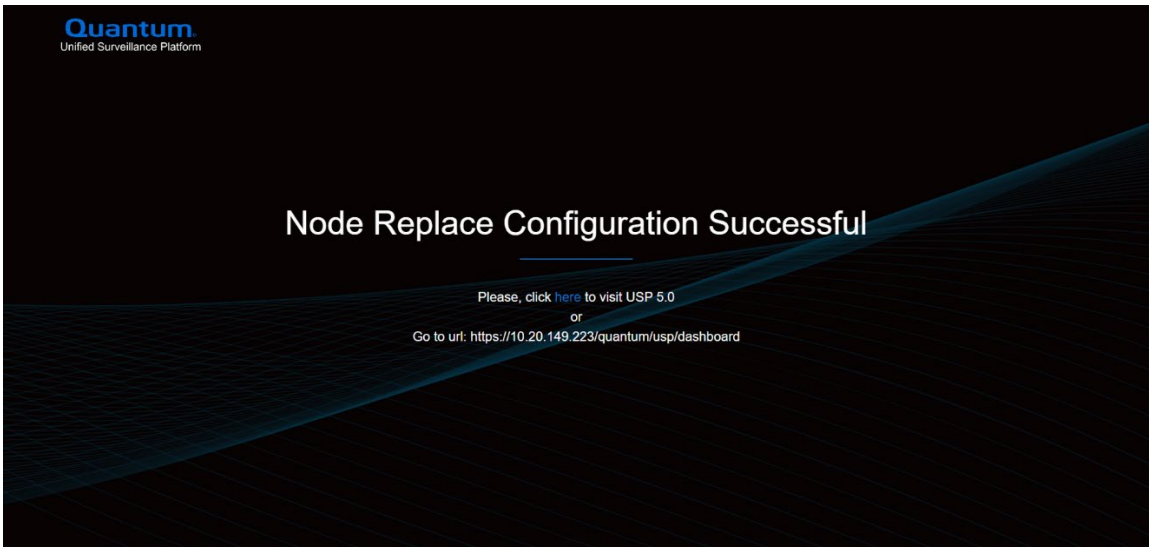


*Troubleshooting*

If you see the following error during node replacement, restart the node replacement process and it should complete.

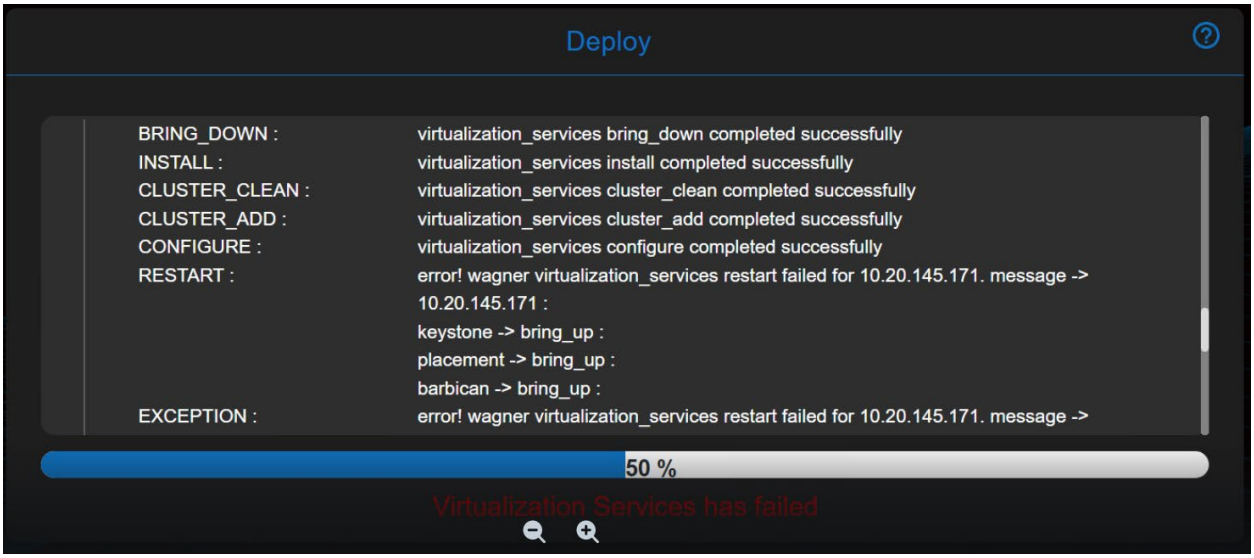


If you see the following error, wait some time and restart the node replacement process.

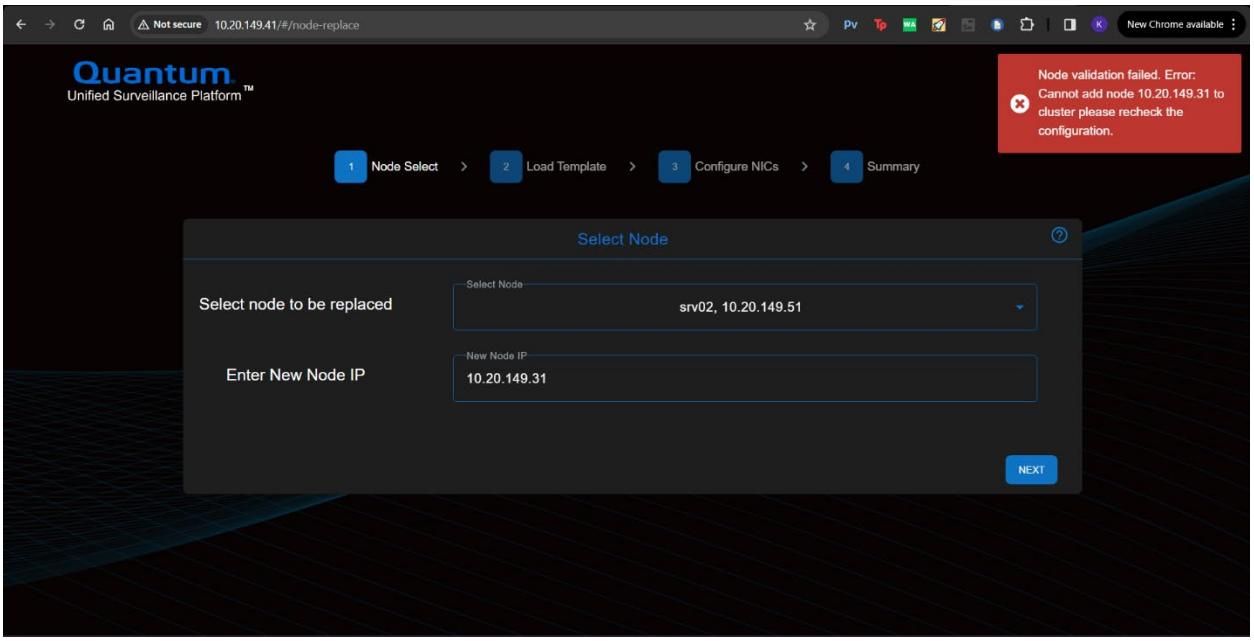


## Bringing Down the Services Gracefully

Before shutting down the entire cluster to do any cluster wide hardware activity or any maintenance, you should bring down the software services. Once the bring down is run, you can shut off all the nodes of the cluster.

Log in to the first node of the cluster and execute the following commands from the command line:

- cd /root/enclouden/dev/stack_orchestrator/
- python3 usp_down.py

Once the cluster is down, a cluster wide shutdown should be run. To bring the cluster back online, power on each server using the out of band management interface.

## Changing Passwords

### *Dashboard*

1. SSH to the first Node in the Cluster to change the administrative password.
2. cd /var/lib/enclouden-sysadmin-dashboard
3. Run: python manage.py change_admin_password --password <new passsword>
4. Login to the Dashboard with the new password.

### *KVM Host*

1. SSH to each host.
2. issue command 'passwd root'
3. Enter the new password.
4. Retype the password to complete the process.

```
Using username "root".
root@10.20.149.111's password:
Last login: Tue Dec 12 18:56:50 2023 from 10.20.149.111
[root@vg1028host1 ~]# passwd root
Changing password for user root.
New password:
Retype new password:
passwd: all authentication tokens updated successfully.
[root@vg1028host1 ~]#
```

## *Out-of-Band Management Interface*

The following steps will need to be followed to change the out-of-band management credentials after the USP cluster has been deployed.

CRITICAL: *Only change the credentials on a single host at a time. If you need to change the credentials on multiple hosts, follow the entire procedure for a single host before moving on to the next host.*

1. Change the out-of-band management credentials using the vendors instructions
2. SSH in to the first host in the USP cluster
3. Issue the command below to list the hosts in the cluster:
   - **pcs status**

```
[root@b1 ~]# pcs resource cleanup
Cleaned up all resources on all nodes
[root@b1 ~]# pcs status
Cluster name: hciclstr01
Cluster Summary:
  * Stack: corosync
  * Current DC: b3.quantumusp.com (version 2.1.4-5.el8-dc6eb4362e) - partition with quorum
  * Last updated: Mon Oct 16 13:40:40 2023
  * Last change:  Mon Oct 16 13:40:17 2023 by root via cibadmin on b1.quantumusp.com
  * 3 nodes configured
  * 30 resource instances configured

Node List:
  * Online: [ b1.quantumusp.com b2.quantumusp.com b3.quantumusp.com ]

Full List of Resources:
  * ipmilan_b1—(stonith:fence_ipmilan):         Started b2.quantumusp.com
  * ipmilan_b2—(stonith:fence_ipmilan):         Started b1.quantumusp.com
  * ipmilan_b3—(stonith:fence_ipmilan):         Started b1.quantumusp.com
  * Clone Set: locking-clone [locking]:
    * Started: [ b1.quantumusp.com b2.quantumusp.com b3.quantumusp.com ]
  * Clone Set: vg_glance_shared-clone [vg_glance_shared]:
    * Started: [ b1.quantumusp.com b2.quantumusp.com b3.quantumusp.com ]
  * Clone Set: vg_nova_shared-clone [vg_nova_shared]:
    * Started: [ b1.quantumusp.com b2.quantumusp.com b3.quantumusp.com ]
  * Clone Set: vg_cinder_shared-clone [vg_cinder_shared]:
    * Started: [ b1.quantumusp.com b2.quantumusp.com b3.quantumusp.com ]
  * openstack-cinder-volume     (systemd:openstack-cinder-volume):      Started b3.quantumusp.com
  * fence-nova-hciclstr01       (stonith:fence_compute):        Started b2.quantumusp.com
  * nova-evacuate       (ocf::openstack:NovaEvacuate):   Started b3.quantumusp.com

Daemon Status:
  corosync: active/enabled
  pacemaker: active/enabled
```

4. The names in the right column under “Full List of Resources” correspond to the host name. Locate the host that you changed the password on and gather the name from the left column. In the example below we will use ipmilan_b2.
5. Issue the command below to update the cluster with the new password. Note that “ipmilan_b2” will match your hostname.

- **pcs stonith update ipmilan_b2 username='username' password='newpassword'**

6. Now issue the command below. If there are no errors listed, then everything worked as expected.
    - wagner cluster status

## Configure Cluster for GPU Passthrough

To configure passthrough GPU devices on your instances, you must first register the GPU devices with the cluster. This is done by executing a script on the first node in the cluster. Once the GPUs are registered on each host you can attach or detach them from the cluster.

1. SSH in to the first node of the cluster.

NOTE: *This step only needs to be executed on a single host, and it does not matter if the host executing the script has a GPU device.*

2. Change directories by issuing the following command:
    - cd /root/enclouden/dev/stack_orchestrator/
3. Execute the setup script by running the command below. This will configure the GPUs on each host in the cluster.
    - python3 gpu_setup.py -a setup
4. There will be several prompts asking which GPUs to configure for passthrough.
5. Once the process is complete, you will see a “success” message.

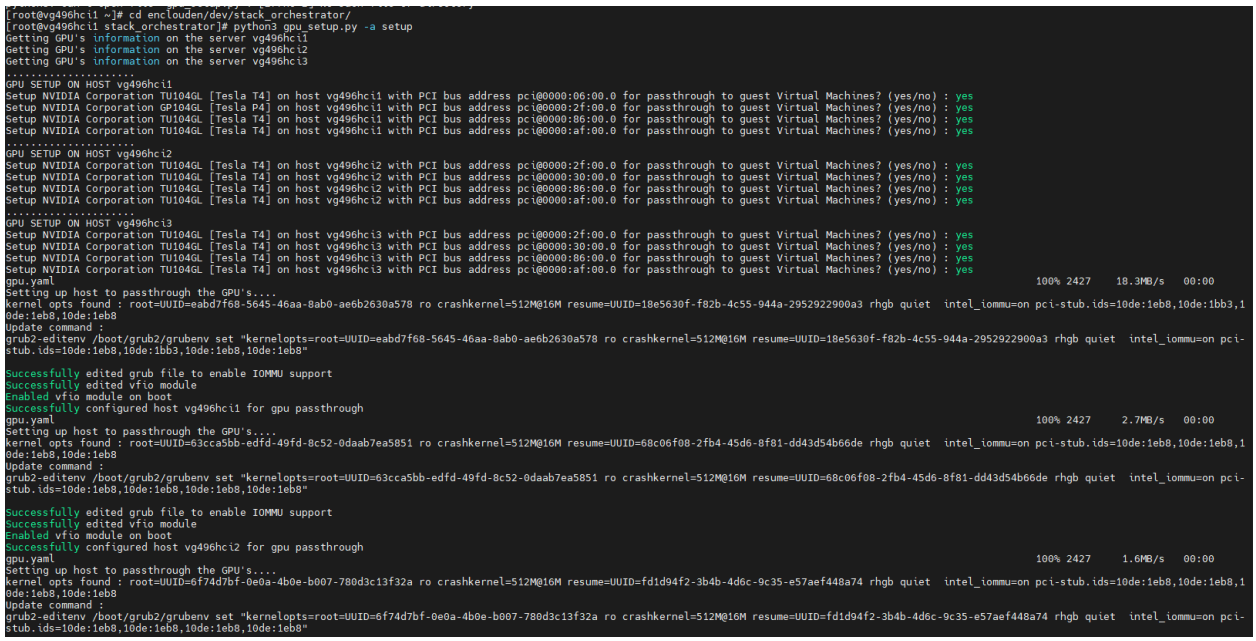

```
[root@vg496hci1 ~]# cd enclouden/dev/stack_orchestrator/
[root@vg496hci1 stack_orchestrator]# python3 gpu_setup.py -a setup
Getting GPU's information on the server vg496hci1
Getting GPU's information on the server vg496hci2
Getting GPU's information on the server vg496hci3
.....................
GPU SETUP ON HOST vg496hci1
Setup NVIDIA Corporation TU104GL [Tesla T4] on host vg496hci1 with PCI bus address pci@0000:06:00.0 for passthrough to guest Virtual Machines? (yes/no) : yes
Setup NVIDIA Corporation GP104GL [Tesla P4] on host vg496hci1 with PCI bus address pci@0000:2f:00.0 for passthrough to guest Virtual Machines? (yes/no) : yes
Setup NVIDIA Corporation TU104GL [Tesla T4] on host vg496hci1 with PCI bus address pci@0000:86:00.0 for passthrough to guest Virtual Machines? (yes/no) : yes
Setup NVIDIA Corporation TU104GL [Tesla T4] on host vg496hci1 with PCI bus address pci@0000:af:00.0 for passthrough to guest Virtual Machines? (yes/no) : yes
.....................
GPU SETUP ON HOST vg496hci2
Setup NVIDIA Corporation TU104GL [Tesla T4] on host vg496hci2 with PCI bus address pci@0000:2f:00.0 for passthrough to guest Virtual Machines? (yes/no) : yes
Setup NVIDIA Corporation TU104GL [Tesla T4] on host vg496hci2 with PCI bus address pci@0000:30:00.0 for passthrough to guest Virtual Machines? (yes/no) : yes
Setup NVIDIA Corporation TU104GL [Tesla T4] on host vg496hci2 with PCI bus address pci@0000:86:00.0 for passthrough to guest Virtual Machines? (yes/no) : yes
Setup NVIDIA Corporation TU104GL [Tesla T4] on host vg496hci2 with PCI bus address pci@0000:af:00.0 for passthrough to guest Virtual Machines? (yes/no) : yes
.....................
GPU SETUP ON HOST vg496hci3
Setup NVIDIA Corporation TU104GL [Tesla T4] on host vg496hci3 with PCI bus address pci@0000:2f:00.0 for passthrough to guest Virtual Machines? (yes/no) : yes
Setup NVIDIA Corporation TU104GL [Tesla T4] on host vg496hci3 with PCI bus address pci@0000:30:00.0 for passthrough to guest Virtual Machines? (yes/no) : yes
Setup NVIDIA Corporation TU104GL [Tesla T4] on host vg496hci3 with PCI bus address pci@0000:86:00.0 for passthrough to guest Virtual Machines? (yes/no) : yes
Setup NVIDIA Corporation TU104GL [Tesla T4] on host vg496hci3 with PCI bus address pci@0000:af:00.0 for passthrough to guest Virtual Machines? (yes/no) : yes
gpu.yaml                                                                                                                    100% 2427    18.3MB/s   00:00
Setting up host to passthrough the GPU's....
kernel opts found : root=UUID=eabd7f68-5645-46aa-8ab0-ae6b2630a578 ro crashkernel=512M@16M resume=UUID=18e5630f-f82b-4c55-944a-2952922900a3 rhgb quiet  intel_iommu=on pci-stub.ids=10de:1eb8,10de:1bb3,1
0de:1eb8,10de:1eb8
Update command :
grub2-editenv /boot/grub2/grubenv set "kernelopts=root=UUID=eabd7f68-5645-46aa-8ab0-ae6b2630a578 ro crashkernel=512M@16M resume=UUID=18e5630f-f82b-4c55-944a-2952922900a3 rhgb quiet  intel_iommu=on pci-
stub.ids=10de:1eb8,10de:1bb3,10de:1eb8,10de:1eb8"

Successfully edited grub file to enable IOMMU support
Successfully edited vfio module
Enabled vfio module on boot
Successfully configured host vg496hci1 for gpu passthrough
gpu.yaml                                                                                                                    100% 2427     2.7MB/s   00:00
Setting up host to passthrough the GPU's....
kernel opts found : root=UUID=63cca5bb-edfd-49fd-8c52-0daab7ea5851 ro crashkernel=512M@16M resume=UUID=68c06f08-2fb4-45d6-8f81-dd43d54b66de rhgb quiet  intel_iommu=on pci-stub.ids=10de:1eb8,10de:1eb8,1
0de:1eb8,10de:1eb8
Update command :
grub2-editenv /boot/grub2/grubenv set "kernelopts=root=UUID=63cca5bb-edfd-49fd-8c52-0daab7ea5851 ro crashkernel=512M@16M resume=UUID=68c06f08-2fb4-45d6-8f81-dd43d54b66de rhgb quiet  intel_iommu=on pci-
stub.ids=10de:1eb8,10de:1eb8,10de:1eb8,10de:1eb8"

Successfully edited grub file to enable IOMMU support
Successfully edited vfio module
Enabled vfio module on boot
Successfully configured host vg496hci2 for gpu passthrough
gpu.yaml                                                                                                                    100% 2427     1.6MB/s   00:00
Setting up host to passthrough the GPU's....
kernel opts found : root=UUID=6f74d7bf-0e0a-4b0e-b007-780d3c13f32a ro crashkernel=512M@16M resume=UUID=fd1d94f2-3b4b-4d6c-9c35-e57aef448a74 rhgb quiet  intel_iommu=on pci-stub.ids=10de:1eb8,10de:1eb8,1
0de:1eb8,10de:1eb8
Update command :
grub2-editenv /boot/grub2/grubenv set "kernelopts=root=UUID=6f74d7bf-0e0a-4b0e-b007-780d3c13f32a ro crashkernel=512M@16M resume=UUID=fd1d94f2-3b4b-4d6c-9c35-e57aef448a74 rhgb quiet  intel_iommu=on pci-
stub.ids=10de:1eb8,10de:1eb8,10de:1eb8,10de:1eb8"
```

# Known Issues and Limitations

## Instance Operations During Image Upload

Instance Operations are not supported while an image is uploading. Performing virtual machine operations while uploading an image can lock the department and requires assistance from Quantum Support.

If a guest VM is left in an error state, but is still functional, stop the guest VM to clear the error state and then power it back on for use.

Quantum technology, software, and services provide the solutions that today's organizations need to make video and other unstructured data smarter – so their data works for them and not the other way around. With over 40 years of innovation, Quantum's end-to-end platform is uniquely equipped to orchestrate, protect, and enrich data across its lifecycle, providing enhanced intelligence and actionable insights. Leading organizations in cloud services, entertainment, government, research, education, transportation, and enterprise IT trust Quantum to bring their data to life, because data makes life better, safer, and smarter. Quantum is listed on Nasdaq (QMCO) and the Russell 2000® Index. For more information visit www.quantum.com.

www.quantum.com | 800-677-6268