

Dell EMC Metro node 7.0.1

Online Help

Version 7.0.1

Notes, cautions, and warnings

 **NOTE:** A NOTE indicates important information that helps you make better use of your product.

 **CAUTION:** A CAUTION indicates either potential damage to hardware or loss of data and tells you how to avoid the problem.

 **WARNING:** A WARNING indicates a potential for property damage, personal injury, or death.

Figures.....	8
Chapter 1: Welcome.....	9
Metro node overview.....	9
Welcome to metro node.....	10
What's New?.....	10
Chapter 2: Using the GUI.....	11
About the GUI.....	11
Sorting data.....	12
Selecting objects.....	12
Properties panel.....	12
Using the global search.....	12
Accessing online help.....	13
Chapter 3: Configuring GUI default settings.....	14
Configuring the Idle Timeout.....	14
Client Logging.....	14
Configure Settings.....	14
Chapter 4: Using storage hierarchy maps.....	16
About storage hierarchy maps.....	16
Single cluster maps.....	16
Distributed maps.....	17
Storage view maps.....	20
Viewing a map.....	21
Zooming in or out of a map and expand or collapse of all virtual volumes.....	22
Searching for a storage object in a map.....	22
Identifying unhealthy storage objects in a map.....	22
Exporting a map to an image file.....	23
Printing a map.....	23
Chapter 5: Viewing system status.....	25
The System Status view.....	25
Metro node cluster status.....	27
Metro node director status.....	28
Product version.....	29
Chapter 6: Monitoring the system.....	30
Performance.....	30
The Performance Monitoring dashboard.....	30
Viewing a chart.....	30
Modifying a dashboard layout.....	31
Creating a custom dashboard.....	31

Removing a chart.....	31
Moving a chart.....	31
Back-end Bandwidth Chart.....	32
Back-end Throughput chart.....	33
Back-end Errors chart.....	34
Back-end Latency chart.....	35
CPU utilization chart.....	36
Heap Usage chart.....	37
Front-end Queue Depth chart.....	38
Front-end Bandwidth chart.....	39
Front-end Latency chart.....	40
Front-end Throughput chart	41
Front-end Aborts chart	43
Write Latency Delta chart.....	44
WAN Port Performance chart	44
WAN Latency chart.....	46
Rebuild Status dashboard.....	47
Virtual Volumes dashboard.....	49
Front End Ports dashboard.....	54
System Health.....	55

Chapter 7: Provisioning storage..... 57

Guide.....	57
Provisioning from storage volumes.....	57
Provision Job properties.....	58
Distributed storage.....	58
Storage arrays.....	59
Supported storage arrays.....	59
Using the Storage Arrays view.....	59
Rediscovering a storage array.....	60
Viewing logical units for an array.....	60
Forgetting a storage array.....	61
Storage array connectivity status.....	61
Storage volumes.....	61
About storage volumes.....	61
The Storage Volumes view.....	62
Claiming storage volumes.....	63
The Claim Storage wizard.....	64
Setting a storage volume rebuild type	65
Unclaiming a storage volume.....	65
Viewing ITLs for a storage volume.....	66
Renaming a storage volume.....	66
Storage Volume status.....	66
Devices.....	67
About devices.....	67
Using the Devices view.....	68
The Create Devices wizard.....	69
The Add Local/Remote Mirror wizards.....	70
Viewing the status of IO to a device.....	71
Creating a device.....	71

Renaming a device.....	71
Deleting a device.....	72
Mirroring a device.....	72
Device status.....	72
Device component properties.....	73
Device properties.....	73
Distributed device properties.....	74
Add capacity to virtual volumes.....	75
Extent properties.....	75
Extents.....	76
About extents.....	76
Using the Extents view.....	76
The Create Extents wizard.....	77
Creating an extent.....	78
Renaming an extent.....	78
Deleting an extent.....	79
Extent status.....	79
Distributed devices.....	80
About distributed devices.....	80
The Distributed Devices view.....	80
The Create Distributed Device from Claimed Storage Volumes wizard.....	81
Distributed device rule sets.....	82
Changing the rule set for a distributed device.....	82
Creating a distributed device.....	83
Deleting a distributed device.....	83
Renaming a distributed device.....	84
Distributed Device status.....	84
Virtual volumes.....	85
About virtual volumes.....	85
The Virtual Volumes view.....	86
The Distributed Virtual Volumes view.....	88
Creating a virtual volume.....	89
About virtual volume expansion.....	90
Expanding a virtual volume using storage volumes.....	90
Enabling or disabling remote access for a volume.....	90
Manually assigning LUN numbers to volumes	91
Deleting a volume.....	91
Renaming a volume.....	92
Tearing down a volume.....	92
Virtual Volume status.....	92
Pool properties.....	93
Virtual volume properties.....	94
Show ITLs dialog box.....	95
Logical unit properties.....	95
ALUA Support field values.....	96
Visibility field values.....	96
Extent or Device mobility job properties.....	96
Metro node port properties.....	96
Storage array properties.....	97
Storage view properties.....	97

Storage volume properties.....	97
Create Virtual Volumes dialog box.....	98
Consistency group.....	98
About consistency groups.....	98
Using the Consistency Groups view.....	99
Distributed Consistency Groups view.....	100
Create Consistency Group wizard.....	101
Types of consistency groups.....	101
Creating a consistency group.....	102
Adding a volume to a consistency group.....	102
Removing a volume from a consistency group.....	103
Deleting a consistency group.....	103
Consistency Group status.....	104
Consistency group properties.....	105
Step 1: Select or create a consistency group for the virtual volume.....	105
Step 1: Create a consistency group.....	105
Step 2: Select volume options.....	106
Step 3: Select a storage pool.....	106
Step 3: Select a pool for each mirror on the second cluster.....	107
Step 3: Select a pool for each mirror in the cluster.....	108
Step3: Create thin virtual volumes.....	109
Select a storage view for the virtual volume(s) (optional).....	109
Step 5: Review your selections.....	109
Step 6: View results.....	109
Step 2: Select volume options.....	110
Step 2: Select volume options.....	110
Step 3: Select a storage volume to create the virtual volume.....	110
Step 3: Select a source and target storage volume.....	110
Step 3: Create thin volumes.....	111
Step 3: Select a target storage volume on the remote cluster.....	111
Step 3: Select target storage on the remote cluster.....	111
Step 6: View results.....	111
Show Logical Units.....	111
Chapter 8: Exporting storage.....	112
Initiators and metro node ports.....	112
About Initiators and metro node ports.....	112
Using the Ports view.....	112
Using the Initiators view.....	113
Enabling or disabling a port.....	115
Registering an initiator.....	115
Rediscovering an initiator.....	116
Manually registering an initiator.....	116
Unregistering an initiator.....	117
Metro node port status.....	117
Initiator properties.....	117
Add capacity to virtual volumes.....	117
Storage views.....	118
About storage views	118
Using the Storage Views screen.....	119

The Create Storage View wizard.....	120
Creating a storage view.....	120
Deleting a storage view.....	121
Renaming a storage view.....	121
Adding or removing initiators from a storage view.....	122
Adding virtual volumes to a storage view.....	122
Removing virtual volumes from a storage view.....	123
Adding or removing metro node ports from a storage view.....	123
Storage view status.....	124
Storage group properties.....	124
Director properties.....	124
Cluster properties.....	125
Chapter 9: Moving data.....	126
Mobility.....	126
Move Data Within Cluster.....	127
Move Data Across Clusters.....	128
Create Mobility Job wizards.....	129
Mobility job transfer size.....	130
Creating a mobility job.....	130
Viewing job details.....	131
Committing a job.....	131
Canceling a job.....	131
Pausing a job.....	132
Resuming a job.....	132
Removing the record of a job.....	132
Changing a job transfer size.....	132
Searching for a job.....	133
Mobility job status.....	133
Chapter 10: Notifications.....	134
Platform alerts view.....	134
Hardware alerts view.....	136
Configure alerts view.....	137
Configure notifications.....	138
Configure email notifications.....	142
Test Alerts.....	143
Test Email.....	143
Platform view properties.....	144
Hardware view properties.....	145

1	Simple device configuration.....	68
2	Complex device configuration.....	68
3	Extent configurations.....	76
4	Distributed device configurations.....	80

Welcome

This section gives overview on metro node, new features, and known issues.

Topics:

- [Metro node overview](#)
- [Welcome to metro node](#)
- [What's New?](#)

Metro node overview

The metro node family is a solution for federating Dell EMC and non Dell EMC storage arrays. Metro node resides between the servers and heterogeneous storage assets supporting a variety of arrays from various vendors. Metro node simplifies storage management by allowing LUNs provisioned from various arrays to be managed through a centralized management interface. Metro node removes physical barriers within, across and between data centers.

- Metro node Local provides simplified management and non-disruptive data mobility across heterogeneous arrays.
- Metro node Metro provides data access and mobility between two metro node clusters within synchronous distances.

With a unique scale-up and scale-out architecture, metro node's advanced data caching and distributed cache coherency provides workload resiliency, automatic sharing, balancing and failover of storage domains, and enables both local and remote data access with predictable service levels.

System management

In addition to the web-based Graphical User Interface (UI), metro node also supports a command-line interface (CLI) for managing your system. The *CLI Guide for the metro node appliance* provides more information.

Basic workflow

To begin using metro node, provision and export storage so that hosts and applications can use the storage. For an overview of the provisioning process and step-by-step instructions for provisioning storage, see **Guide**. To view the Guide, click **Provision Storage**, and then expand the **GUIDE** at the upper-right corner of the screen.

The screenshot displays the Dell EMC VPLEX-Metro Provision Storage interface. The top navigation bar includes 'Dell EMC | VPLEX | VPLEX-Metro' and status indicators for 'Cluster-1' and 'Cluster-2'. The main menu shows 'Dashboard', 'Performance', 'Provision Storage' (selected), 'Mobility', and 'Jobs'. The 'Provision Storage' section shows 'Provision on: Cluster-1' and 'View By: Storage Volumes'. A list of 53 storage volumes is shown, with one volume selected: VPD83T3:6000144000000100184d8bca06ce1d4, with a capacity of 80.00 GB and a health status of green. The interface also includes a sidebar with navigation options like 'Dashboard', 'Performance', 'Provision Storage', 'Mobility', and 'Jobs'. A 'Guide' button is highlighted in the top right corner.

Welcome to metro node

Metro node is a Distributed Storage Federation solution that provides Local or Distributed Federation within a single data center, or between two clusters within synchronous distances.

- Mobility—Move and relocate Virtual Machines, applications, and data over distance.
- Availability—Create and maintain high-availability storage infrastructure by mirroring across locations.
- Collaboration—Enable concurrent read/write access to data across locations.

For more information:

- [Dell EMC Online Support](#)
- [Release Notes for metro node feature](#)
- [CLI Guide for metro node feature](#)
- [Administration Guide for metro node feature](#)
- [Security Configuration Guide for metro node feature](#)
- [Configuration Guide for metro node feature](#)
- [SolVe Online](#)

Your suggestions will help us continue to improve the accuracy, organization, and overall quality of the user publications. Send your comments to: vplex.doc.feedback@dell.com

What's New?

Notifications

The Notification page is the consolidated area of all the appliance alerts and events for the cluster. Events provide information about changes to the system, which indicate there is a problem with the system. Alerts are categorized by severity which indicates the urgency of the alert.


On the notifications tab, there are two alert tabs :

Platform Alerts: You can monitor the status of metro node platform which includes alerts at director level and cluster level. In the Live alerts drop-down, you can view alerts that were created during the last 48 hours.

Hardware Alerts: You can monitor the status of metro node hardware which includes alerts that are generated at hardware level. In the Live alerts drop-down, you can view alerts that were created during the last 48 hours.

In the details of each alert, additional information can be found including the Severity, Message, Corrective Action, and other properties. This information is useful in troubleshooting scenarios and allows users to remediate issues seen on the system.

When alerts are no longer relevant or are resolved, the system automatically clears the alerts with no user intervention. This action ensures that cleared alerts are hidden from the default view so that only relevant issues are displayed. Acknowledging an alert does not indicate that the issue is resolved. Acknowledged alerts can be viewed as well through table filtering options.

 **NOTE:** Historical alerts are not supported for internal partner drop.

Configure Alerts

Under settings, go to **Notifications**. The user can also configure alerts at various levels.

- **Enable or disable alert at alert condition ID level**
 - For example: 0x30004 is the condition ID for **Storage Volume Unreachable** alert and if the customer wants to disable this alert, select that particular alert and disable 0x30004. With this, whenever metro node hits 0x30004, it will not be listed in GUI, and also a dial home event will not be generated.
- **Configure alerts at component level**
 - When the user click **Configure notifications** button, then side panel shows the options to disable the alerts at components level. For example: If **System Volume** is disabled, then whenever metro node presses any system volume-related alerts, it will not be listed in GUI, and also a dial home event will not be generated.

Using the GUI

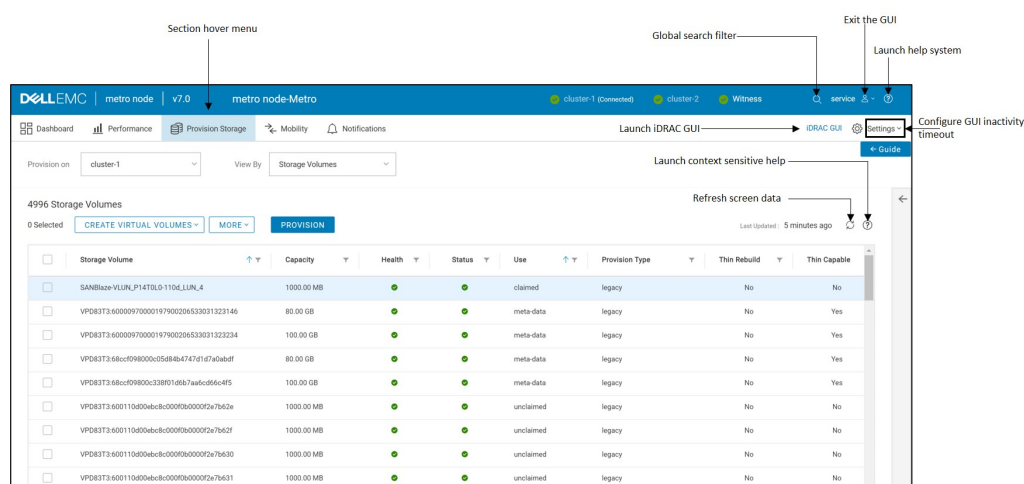
This section describes how to use the GUI.

Topics:

- About the GUI
- Sorting data
- Selecting objects
- Properties panel
- Using the global search
- Accessing online help

About the GUI

Management Console screen





Section hover menus

The main menu on the Management Console screen provides the following menu sections:

- **Dashboard** — Provides a quick access to System Status (default home page) and System Health.
- **Performance** — Provides a visual summary of performance statistics information, the health status of major metro node components, and the status of provisioning jobs.
- **Provision Storage** — Provides an overview of provisioning and launches the wizards that allow you to provision from storage pools or storage volumes. From this menu, you can also view storage information for each cluster, view distributed storage devices, and learn how to provision storage on metro node.
- **Mobility** — Provides a wizard for moving data and applications. From this menu, you can move data within a cluster or across clusters.

Sorting data

When columns of data are listed in a view, an arrow displays in the column that is sorted by default. The arrow indicates the direction of the sort; ascending () or descending (). You can also sort in the following ways:

- Create a sort order — Click in a column header to create a sort order.
- Change a sort order — Click the column header to sort the data in the column, and then click a second time to change the direction of the sort.
- Clear a sort order — Click a sorted column header to remove the number and clear the sort order.
- Rearrange columns — Click the column header, and then drag-and-drop the column in the desired location.

Selecting objects

In any storage object view or dialog box that shows a list of objects, you can select or unselect in the following ways:

- Select a single object — Click anywhere in the object row.
- Select adjacent objects — Click the first row, press and hold **Shift**, and then click the last row of the group to be selected.
- Select non-adjacent objects — Press and hold **Ctrl**, and then click in each row you want to select.
- Add or remove objects from the current selection — Press and hold **Ctrl**, and then click in the row you want to add or remove.


Properties panel

In the GUI screen, you can view the properties of storage objects in the **Properties** panel at the right side of the screen.

Using the global search

The search filter on the right in the GUI title bar allow you to enter the full or partial name of a storage object and search across all metro node-managed objects.

You can also use the search filter to search across a specific object type. For example, to search for a specific virtual volume, select virtual volumes from the drop-down list, and then enter the volume's full or partial name in the text box. Only virtual volumes matching the search criteria will display in the search results. If **All** is displayed in the drop-down, the search will be performed across all storage objects, and will include other types of storage objects that match the search criteria.

 **NOTE:** Global search accepts the wildcard (*) character, and storage object names are case-sensitive.


You can search for the following metro node-managed objects:

- Virtual volumes
- Devices (including distributed devices and underlying devices)
- Extents
- Storage volumes
- Initiators
- Ports
- Storage arrays
- Consistency groups
- Storage views

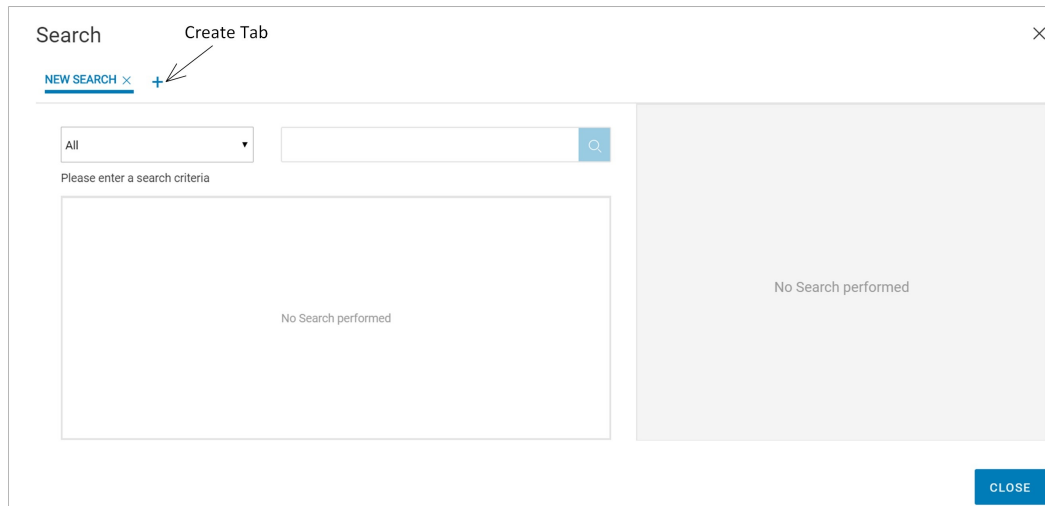
Using the Search Results dialog box

The **Search Results** dialog box shows the object name, the cluster(s) to which it belongs, and the type of object. You can perform new searches, save or delete the current search results, view the properties of an object, or go directly to the object's screen to perform any of these actions:

- To view the properties of an object listed in the search results, click the object name link.

- To go directly to the object screen to manage that object, hover over the object name, and then click **Manage** in the tool tip.
- To delete the current search results and perform additional searches, type or select the new search criteria in the appropriate box and then click . The current results will update accordingly.
- To save the current search results and perform new searches, hover over the area to the right of the current tab, and then click **+** to create a new tab. In the new tab, select or type the new search criteria. You can move back and forth between tabs.
- To delete a tab, click **x** in the tab.

Creating a new tab for another search



Accessing online help

Click the **Help** icon from the following locations to access online help:

- Main screen — Launches the online help system.
- Any storage object view — Displays help for the current view.
- Dialog box title bar — Displays help for the dialog box.

Configuring GUI default settings

This section describes how to change the GUI default settings.

Topics:

- [Configuring the Idle Timeout](#)
- [Configure Settings](#)

Configuring the Idle Timeout

About this task

A GUI session locks after 10 minutes of inactivity, and you must reenter your password to resume the session. When you re-enter your password, any work in progress resumes at the point of the timeout. You can set the idle timer to up to 12 hours or 720 minutes. The idle timeout is based on the URL used to access the GUI. If you access the GUI on cluster 1 and set the idle timeout, it will apply only to cluster 1. If you connect to cluster 2, the idle timeout dialog will not have the same value as that of cluster 1, because the URL is different.

 **CAUTION:** Setting idle time to greater than 30 minutes may expose a potential security risk.

Steps

1. Select **Settings > Configure Settings** from the upper-right corner of the GUI title bar. The **Configure Settings** dialog box is displayed.
2. In the **Idle Timeout** field, click the drop-down , and select either minutes or hours.
3. Click **SAVE** to save the information.

Client Logging

About this task

Enabled Client Logging catches the JavaScript error on a browser which is not supported, and then client-side exceptions are sent to the server side which help in debugging. By default, Client Logging is enabled.


Steps

1. Select **Settings > Configure Settings** from the upper-right corner of the GUI title bar.
2. To disable the client logging, clear the **Enable Client-Side Logging** check box, and then click **Save**.

Configure Settings


Use this dialog box to configure the inactivity timeout for your connection to the GUI, or to define the data reload policy when navigating through provisioning storage object views.

GUI idle timeout

 **WARNING:** A potential security risk can be exposed by setting idle time to greater than 30 minutes. Therefore, if you set the idle timeout to greater than 30 minutes, a warning will display on this dialog box.

The idle timeout is based on the URL used to access the GUI. Therefore when you access the GUI on cluster 1 and set the idle timeout, it will apply only to cluster 1. If you log into cluster 2, the idle timeout dialog will not have the same value as that of cluster 1, because the URL is different.

The value of the timeout field can be an integer value between 1 and 12 if the field unit value is set to hours. The value of the Timeout field can be an integer value between 1 and 720 if the field unit value is set to minutes.


 **NOTE:** The value for Timeout is stored in a file named using the URL with which you accessed the GUI. If you access the GUI for a cluster using different URLs (For example, using the DNS name of the cluster or using the IP address), the timeout value for each URL used will be different.

Provision storage data reload policy

Select one of the following options to define how often data is reloaded when navigating through provisioning views.

- **Always** — Always reload data upon navigation to a storage object view. This is the default setting.
- **Never** — Never reload data upon navigation to a storage object view.
- **After** — The number of hours or minutes after which to reload data when you navigate to a storage object view. Enter a number in the field and then select the time interval. Once you navigate to an object view, the data will automatically reload at the specified time interval.

The data reload policy setting applies only to the URL and client used to access the GUI when you set the policy. Therefore, if you log in to the GUI again using a different client (laptop versus desktop) or a different URL (cluster 1 versus cluster 2), the reload policy setting will not be applied to that GUI session.

To manually refresh the data, click the Refresh icon  in the upper-right corner of the screen title bar or the **Last Refreshed** link at the bottom of the screen.

Using storage hierarchy maps

This section describes how to use storage hierarchy maps.

Topics:

- [About storage hierarchy maps](#)
- [Single cluster maps](#)
- [Distributed maps](#)
- [Storage view maps](#)
- [Viewing a map](#)
- [Zooming in or out of a map and expand or collapse of all virtual volumes](#)
- [Searching for a storage object in a map](#)
- [Identifying unhealthy storage objects in a map](#)
- [Exporting a map to an image file](#)
- [Printing a map](#)

About storage hierarchy maps

Storage hierarchy maps allow you to select a storage object and view a graphical display of all other objects that are connected to the selected object. It includes underlying objects, and top-level objects that use the selected object. You can then analyze further to see the health and properties of each connected object. When a map is open, you can [print](#), [export](#), [search for an object](#), and [zoom in and out of a map](#).

There are two types of maps:

- **Storage view map**—Shows the objects in the storage view (initiators, ports, and virtual volumes), and allows you to expand or collapse a virtual volume to see underlying devices, extents, and so forth. You can also expand or collapse the list of initiators, ports, and virtual volumes in the storage view. [Storage view maps](#) provides detailed information.
- **Storage hierarchy map**—Shows the complete storage hierarchy from the cluster down to the back-end storage array. For distributed, remote, and globally visible storage objects, maps show the relationships between these objects across clusters. [Distributed maps](#) provides detailed information.

Single cluster maps

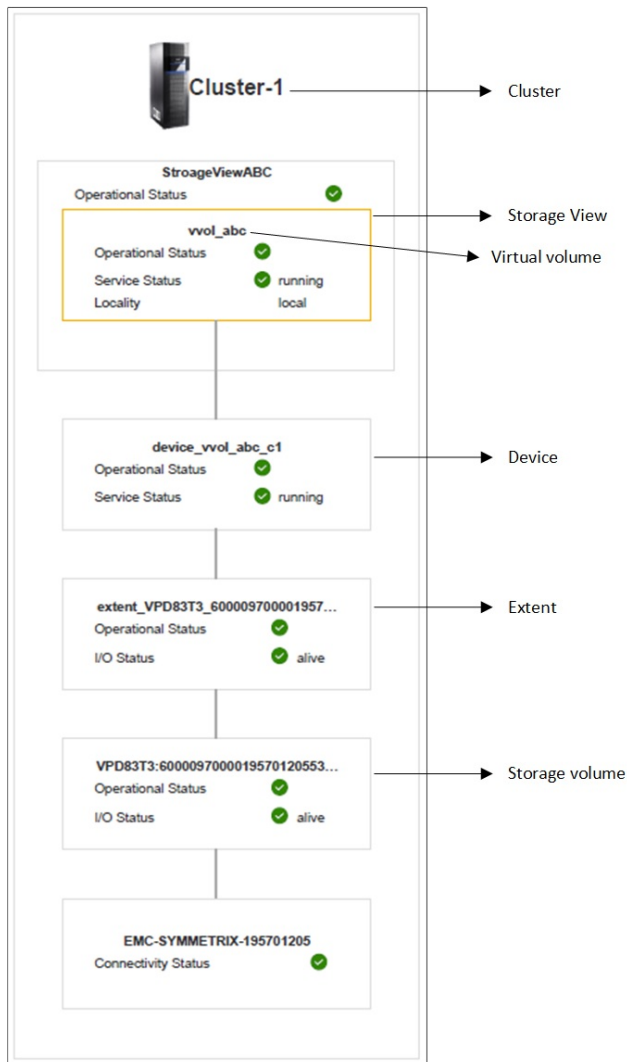
A storage hierarchy map in a single cluster shows the complete storage hierarchy from the cluster down to the back-end storage array. Select a storage volume name, and then click **View Map** to view its underlying object (array), and the objects created on top of the storage volume (extent, device, virtual volume). An orange border around a storage object identifies the subject of the map.

Maps also use color to quickly identify unhealthy storage objects. [Identifying unhealthy storage objects in a map](#) provides more information.

Storage hierarchy maps are available for the following storage objects within a single cluster:

- Virtual volumes
- Devices
- Extents
- Storage volumes

Single cluster map



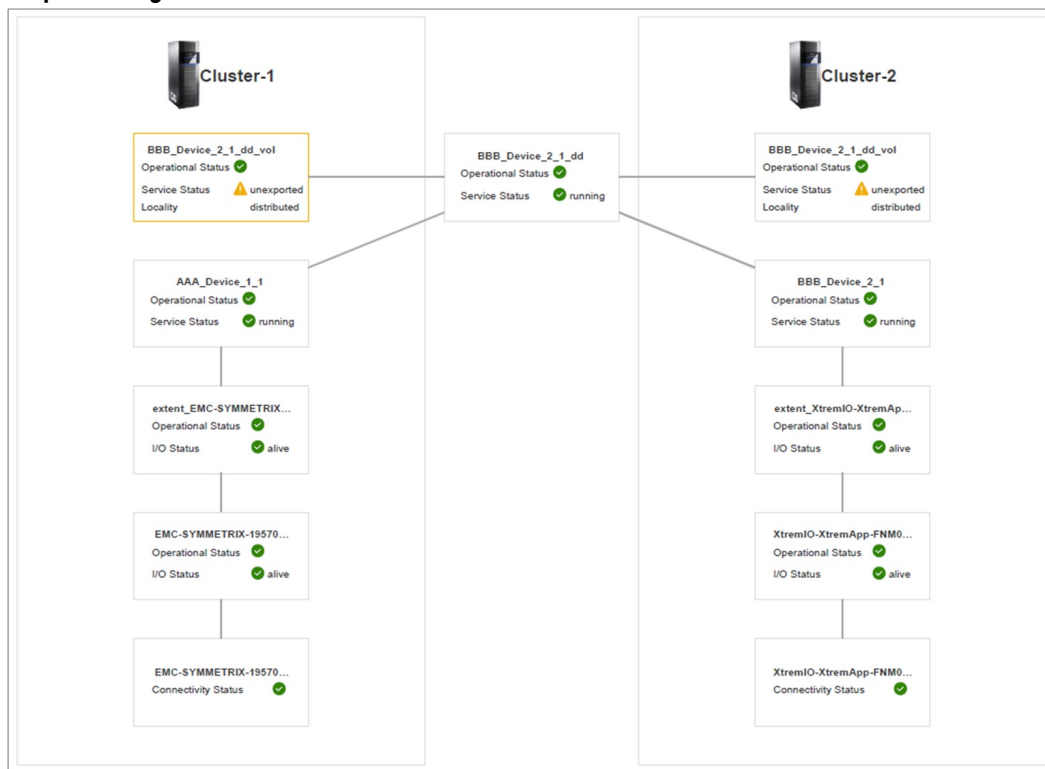
Distributed maps

Storage hierarchy maps are available for storage objects that span clusters. For these storage objects, maps show the relationships between these objects across clusters.

Distributed devices and virtual volumes

Distributed devices appear outside of both clusters because they are not associated with a specific cluster, but instead spans both clusters. A virtual volume on top of a distributed device appears in both clusters and may appear in different storage views depending on the view membership on either cluster.

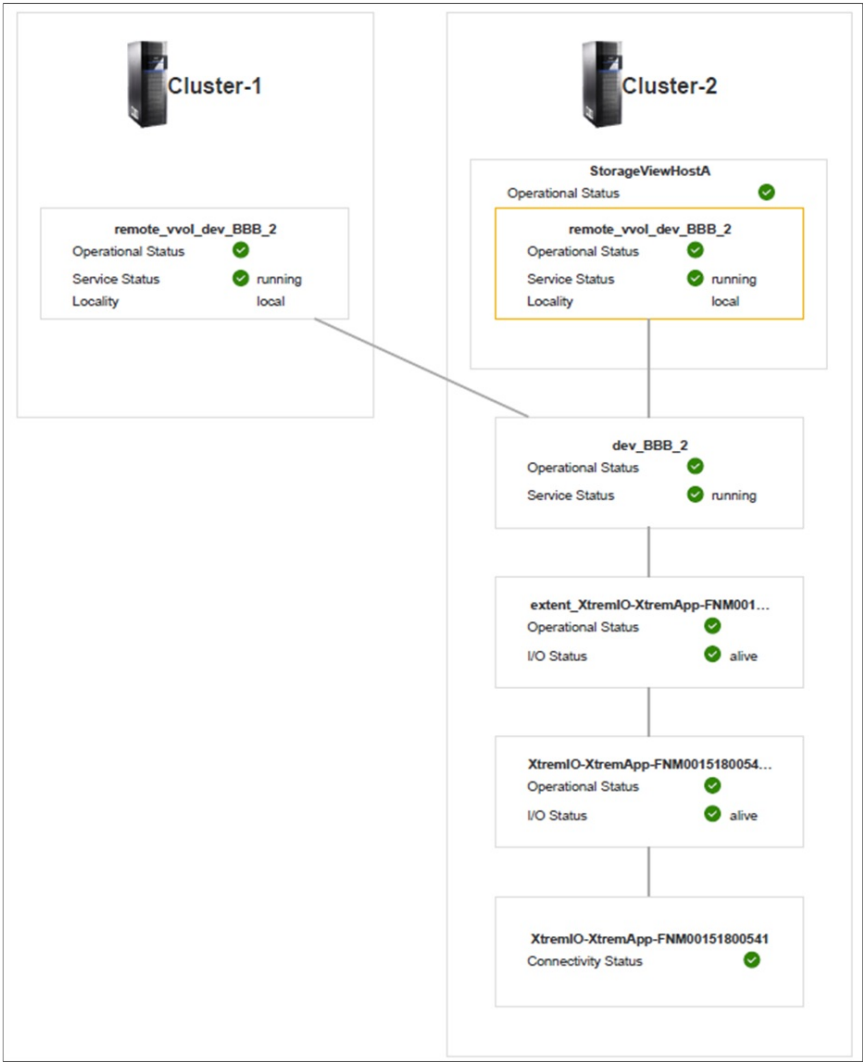
Map showing a distributed device and associated virtual volume



Remote virtual volumes

A virtual volume that has remote locality means that its supporting device is on a remote cluster. The associated map shows the supporting device on the other cluster along with the underlying storage objects.

Map showing a remote virtual volume



Globally visible devices

A virtual volume that is local to a cluster but is on top of a global supporting device is shown as a remote virtual volume on a different cluster.

Map showing a globally visible device.



Storage view maps

A storage view map shows an overview of the components in a storage view, and allows you to drill down further to see the health and properties of all underlying storage objects, down to the array level. You can [print](#), [export](#), [search for an object](#) and [zoom in an out of a map](#). This topic describes features specific to storage view maps.


NOTE: Storage view maps do not support multi-cluster storage objects. The map will show the objects only on one cluster.

The initial map view

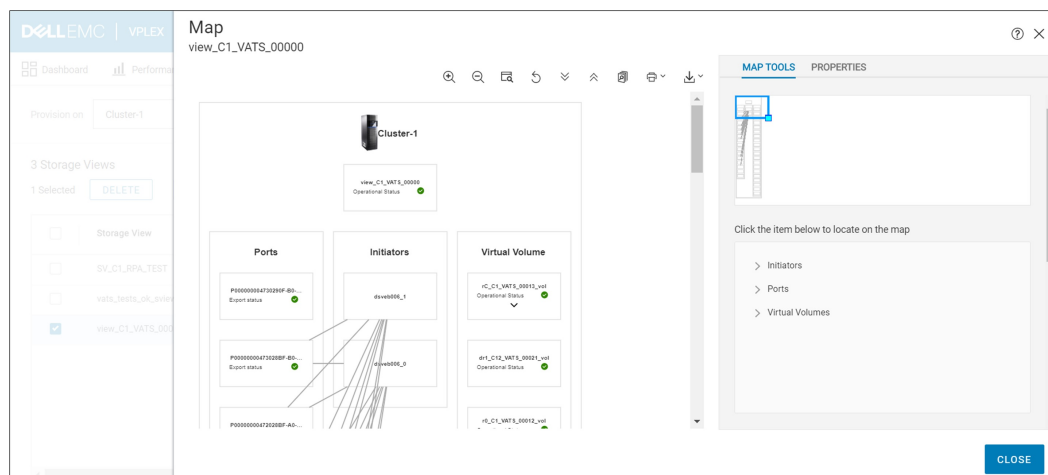
When you launch a storage view map, the initial view shows the entire map. All healthy virtual volumes are shown collapsed (no connecting line to supporting devices). Virtual volumes with an Operational status other than OK are shown expanded (showing connecting line to supporting devices) and are colored red or yellow, depending on the severity of the status. [Identifying unhealthy storage objects in a map](#) provides more information on the use of color to indicate unhealthy storage objects.

Storage view maps can be very large depending on the number of objects in the view. To help you navigate the map, there is a map overview area in the upper portion of the panel that allows you to pan and zoom using a birds-eye view of the map. Under

the Map Overview is a section that lists the storage objects in the storage view. Expanding a list and selecting an object from this list will cause the map to zoom and pan to the selected object.

After zooming in or out of specific areas of the map, you can click  to refresh the map and reset it to the initial view.

Initial map view



Storage view components

The components of the storage view (initiators, ports and virtual volumes) are shown at the bottom of the left panel. Click a tab (**Initiators**, **Ports**, or **Virtual Volumes**) to see a list of those objects in the view, and the operational status of each object. You can expand volumes to show their underlying objects such as supporting device, extent, storage volume, and storage array.



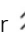
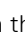
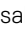
You can also click an object name in the panel on the left to automatically locate and highlight the object in the map. To resize the panel, click and drag the right edge of the panel.

Expanding or collapsing a virtual volume

Each virtual volume box in a map contains icons that allow you to instantly expand or collapse the virtual volume. When you expand a virtual volume, the map shows the underlying supporting devices, down to the array level. When you collapse a virtual volume, no supporting devices display for the virtual volume.

NOTE: The GUI does not support expanding or collapsing virtual volumes with remote or distributed locality, or global visibility. As a result, these virtual volume boxes do not show the expand/collapse icon. To view the locality or visibility of these virtual volumes in a storage view map, place the mouse pointer anywhere in the virtual volume box to display a tool tip that includes this information.

Expand or collapse virtual volumes in the following ways:

- A single virtual volume — Click  or  in a virtual volume box to expand or collapse the virtual volume. If you reset the map with a virtual volume expanded or collapsed, the initial map view zooms in on the expanded or collapsed virtual volume. Expanded virtual volumes are repositioned at the bottom of the virtual volumes area of the map.
- All virtual volumes — Click  or  in the map title bar to expand or collapse all virtual volumes in the map. This operation may take several minutes, depending on the number of virtual volumes in the storage view. If a storage view contains more than 50 virtual volumes,  is disabled.

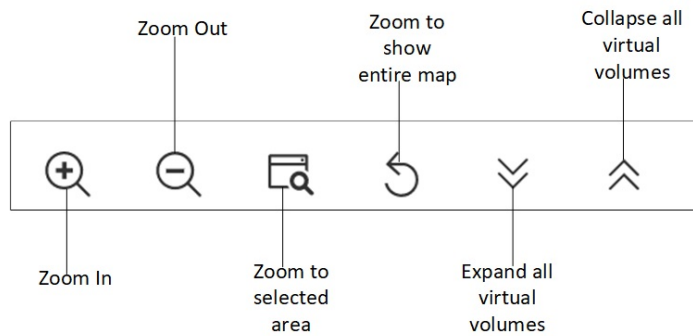
Viewing a map

In the object view, click the object's name, and then click the **View Map** in the object's properties panel.

Zooming in or out of a map and expand or collapse of all virtual volumes

Zoom in or out of a map and expand or collapse of all virtual volume in the following ways:


- Incrementally zoom in or out of a map
- Zoom into a selected area of a map
- Zoom the map to show the entire screen
- Expand or collapse of all virtual volumes



- Move the entire map — In the Map view, go to **MAP TOOLS**, click and hold anywhere in the map (the pointer changes to a blue color box), and then drag in the direction you want to move the map. The area of the map you are viewing is shown in the map overview in the left panel.

Searching for a storage object in a map

Steps

1. Click  in the map's navigation bar to toggle between showing or hiding the **Find** text box.
2. In the **Find** text box, type the full or partial object name, and then click **Find**.
You can use the **Find** filters to narrow the search results.
3. If you are viewing a storage view map, click the object name in the panel on the left to automatically find and highlight the object in the map. The object displays with a blue border. [Storage view maps](#) provide more information.

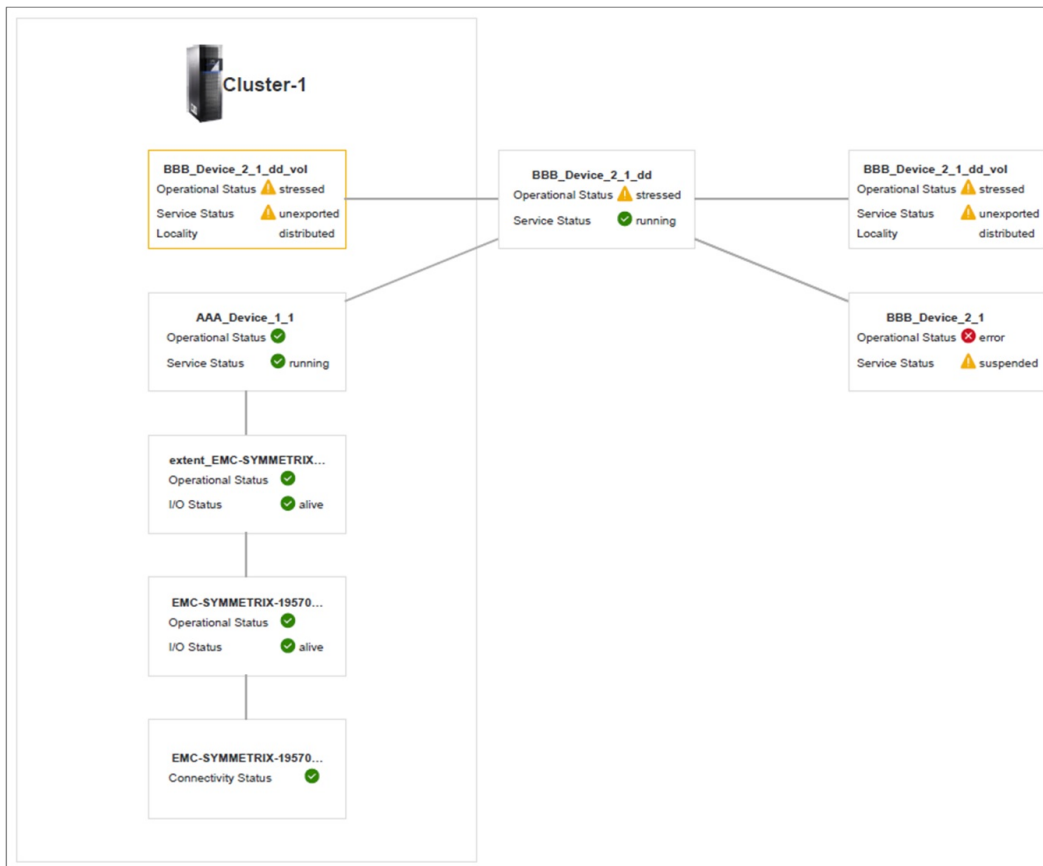
Identifying unhealthy storage objects in a map

The objects in a map are color-coded to allow you to quickly locate a faulty object. The following color scheme indicates unhealthy objects and the level of severity:

- Green—The component is functioning normally.
- Yellow—Warning or minor failure.
- Red—Critical or major failure.

An orange border around an object indicates the subject of the map.

Unhealthy storage objects




Exporting a map to an image file

About this task

You can export an entire map, or the currently visible portion of a map to a .PNG image file, which you can print.

Steps


1. Click  in the map's navigation bar.
2. In the drop-down list, select one of the following:
 - **Export Entire Map** — Exports the entire map to an image file.
 - **Export In View** — Exports the currently visible portion of the map to an image file.
3. In the **Select location for download** dialog box, type a filename and select a location for the file.
4. Click **Save**.
5. Navigate to the file location and doubleclick to open or print the file.

Printing a map

About this task

You can print an entire map, the visible portion of a map, or print a map to span multiple pages.

Steps

1. Click  in the map's navigation bar.

2. In the drop down list, select one of the following:
 - **Print Entire Map** — Prints the entire map.
 - **Print Visible Map** — Prints the currently visible portion of the map. Select a print orientation that most closely matches the dimensions of the map area you want to print. In most cases, landscape orientation works best.
3. In the **Print** dialog box, select the printer and print settings, and then click **Print**.

Viewing system status

This section describes how to view the status of the system and other components.


Topics:

- [The System Status view](#)
- [Metro node cluster status](#)
- [Metro node director status](#)
- [Product version](#)

The System Status view

The **System Status** view is the default view that displays when you connect to the GUI. It shows a graphical representation of your metro node configuration. Use this view to quickly see the status of your system and some of its major components such as directors, storage arrays, and storage views. This view also shows the status of the licenses on the clusters. Colored status indicators next to a cluster indicate the number of components in that cluster that are in a specific state.

If you are managing a metro node Local configuration, you see a figure of a single cluster with status indicators. If you are managing a metro node Metro configuration, you see two clusters, each with their own set of status indicators. The time of the last data refresh displays at the upper-right of the screen, and automatically refreshes every 30 s. You can click the Refresh

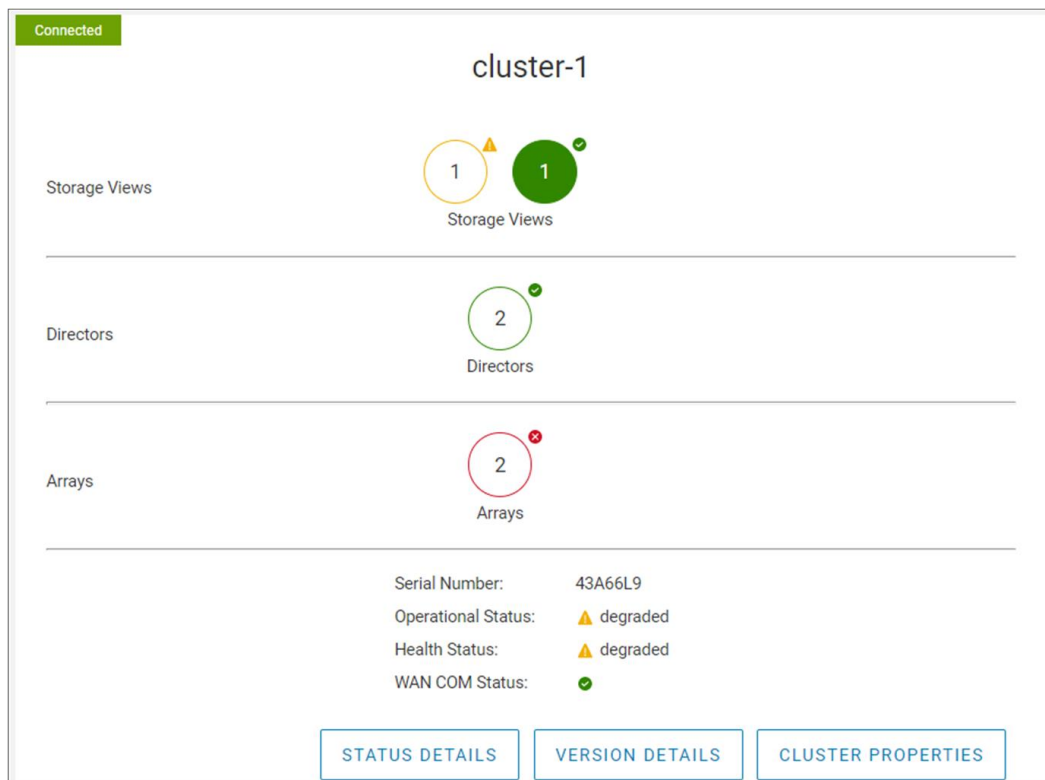
icon  in the upper-right corner of the screen to manually refresh the data.

Colored status indicators

The numbers and colors in each status indicator indicate how many components are in a particular state. The following colors indicate the state of one or more components:

- Green—OK, the storage object is functioning.
- Yellow—Warning, the storage object is changing state or needs user attention.
- Red—Error, an error has occurred that requires user action.
- Grey—Unknown status, metro node cannot determine the status of the component.

In the following example, the Cluster-1 contains two functioning (green) storage view:



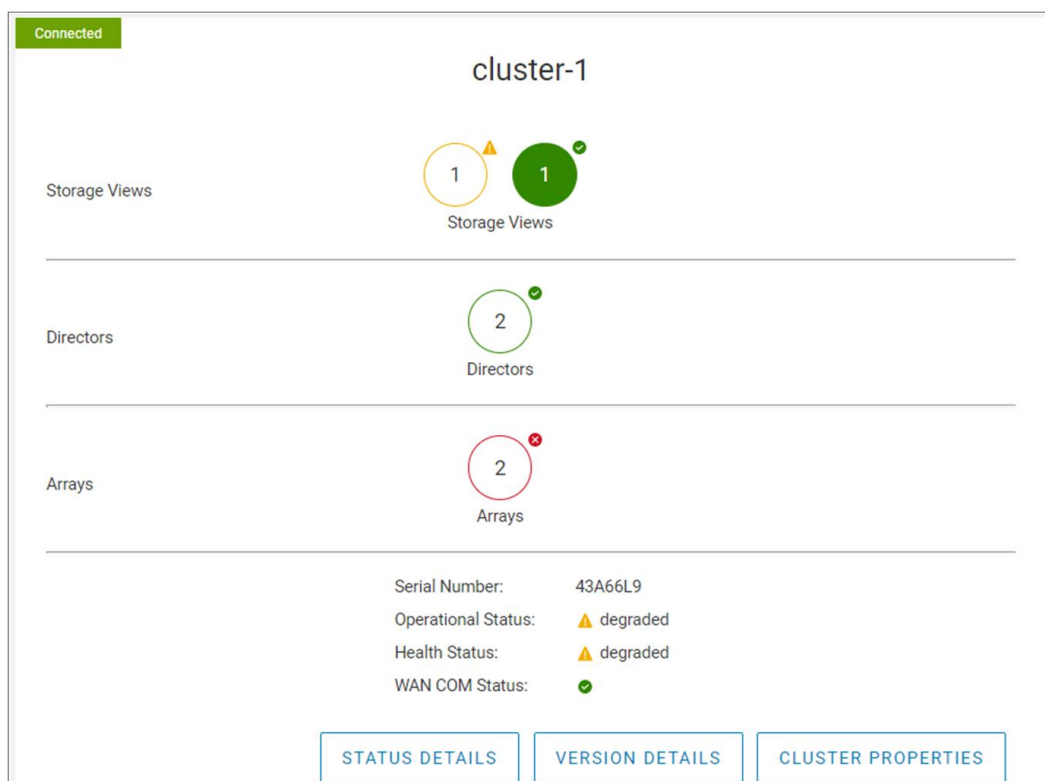
Metro node clusters

The Connect icon **Connected** next to a cluster name indicates the cluster that you are connected to.

To view information about the cluster, click the **Cluster Properties** at the bottom display. For help on any of the cluster properties, click **Help**.

To view version details, click the **Version Details** at the bottom of the status display.

The cluster status displays at the bottom of the cluster image. To view status details, click **Status Details** at the bottom of the status display.



Metro node cluster status

The following table defines the values for the metro node cluster states.

Status Type	Values	Definition
Operational Status	OK	The cluster is functioning normally.
	Degraded	The cluster is not operating as configured. This may indicate one or more of the following problems: degraded redundancy level (for example, a director is dead), suspended virtual volumes or exported volumes, storage volumes not visible from all directors, or the metadata volume is not yet processed.
	Unknown	Metro node encountered a problem determining the operational status of the cluster. This may indicate a degraded state, since it usually means that at least one of the directors is not responding or is communicating abnormally.
	Transitioning	The cluster is reacting to external events and may not be operating as configured. I/O may be suspended during the transition period.
	Isolated	The cluster is not communicating with any other clusters.
	Not-running	All directors in any one of the clusters are in Stopped state and when the VPN is stopped between clusters.
Transitioning Indication	Suspended exports	Cluster lost local quorum during the unvaulting state. This is possibly due to a disk not visible from all directors, no active meta-volume, delayed cluster transition, metadata, not all directors are present, or no quorum.
	Suspended volumes	The distributed device is not accepting new I/O; pending I/O requests are frozen.
Health state	OK	The cluster is functioning normally.
	Degraded	The cluster is not functioning at an optimal level. This may indicate non-functioning remote virtual volumes, unhealthy devices or storage

Status Type	Values	Definition
		volumes, suspended devices, conflicting director count configuration values, out-of-date devices, and so forth.
	Unknown	Metro node cannot determine the cluster's health state, or the state is invalid.
	Critical failure	The cluster is not functioning and may have failed completely. This may indicate a complete loss of back-end connectivity through all directors in the cluster.
	Major failure	<p>The cluster is failing and some functionality may be degraded or unavailable. This may indicate a loss of back-end connectivity through one or more (but not all) directors. When any director has issues with field replacement or if the required number of uninterrupted power supplies is not detected. If any one of the directors is waiting to gain quorum or unvaulting. Example:</p> <pre>Major-failure: engine-2-1 : director-2-1-A : Director has not gained quorum Major-failure: director is unvaulting</pre>
	Minor failure	The cluster is functioning, but some functionality may be degraded. This may indicate one or more unreachable storage volumes.
LOCAL_COM	OK	All LOCAL-COM links have the expected connectivity: this port-group is operating correctly.
	warning	Some links have unexpected connectivity. This port-group is operational but not properly configured. Performance may not be optimal.
	error	Some connectivity is missing from this port-group. It is not operating correctly.
	fail	All connectivity is missing from this port-group. LOCAL-COM is not operational.
WAN_COM	OK	The WAN COM switch is operating normally
	full	All port-groups have a status of either ok or warning. wan-com connectivity is complete through minor configuration errors may still exist. See individual port-group statuses.
	partial	Some port-groups have a status of error or fail, but at least one port-group has a status of OK or warning. wan-com is operating (possibly minimally) through at least one channel. Performance is degraded.
	none	All port-groups have a status of either error or fail. wan-com is not operational.
	not-applicable	The system is a single-cluster (i.e. metro node Local) system. Validating wan-com connectivity is not applicable.

Metro node director status

The following tables define the metro node director Operational and Communication states.

Operational state	Definition
OK	The director is functioning normally.
Stopping	The director is shutting down cleanly.
Starting	The director is starting up.
Stopped	The director has shut down cleanly and is now inactive.

Operational state	Definition
Unknown	Metro node cannot detect the operational state of the director, or the state is invalid.

Communication state	Definition
OK	The director is functioning normally.
Busy	The director's firmware is currently too busy to respond.
No contact	The management server is unable to connect to the director firmware.
Lost communication	The management server has lost the connection to the director firmware.
Unknown	Metro node cannot determine the director's communication state, or the state is invalid.

Product version

The **Product Version** dialog box shows the following product version details for the cluster:

Product Version	The product version of GeoSynchrony running on the cluster.
SMSv2	The version of SMSv2 running on the management server.
Mgmt server base	The version of base software running on the management server.
Mgmt server software	The version of management server software running.
Cluster Witness Server Software	The version of Cluster Witness server software running (if Cluster Witness is installed).
Directors	Lists the name and version number of each director in the cluster.

Monitoring the system

This section describes how to monitor system performance.


Topics:

- [Performance](#)
- [System Health](#)

Performance

The Performance Monitoring dashboard

The performance monitoring dashboard provides a customized view into the performance of your metro node system and its storage objects. You decide which aspects of the system's performance to view and compare. Virtual volumes are the only storage objects for which you can view performance statistics.


 **NOTE:** The charts display data only for the cluster to which you are connected. To simultaneously view charts for another cluster, open a second browser session and connect to the second cluster.

The performance monitoring dashboard provides a tabbed view of the performance of your system, enabling you to view default views of system performance, and customized views. The interface provides the following default dashboards:

- System Resources
- End to End
- Front-end
- Back-end
- Rebuild
- WAN
- Virtual Volumes

You can display up to 12 charts on each dashboard tab, and create up to five dashboard tabs.

Using dashboard tabs

- To close (delete) any tab within the dashboard, click the **X** for that tab. In the **Delete** dialog box, click **Delete**.
- To change the name of an existing tab, select the tab name, and then click . In the Rename Dashboard dialog box, enter the new name for the tab. Click **OK**.
- To move a tab, drag it to the correct location.

Viewing a chart



About this task

You can view a chart in normal or full-screen mode.

Steps

1. From the GUI main menu, select **Performance**.
The Performance dashboard is displayed.
2. Click the tab in which you want to display the chart (or [create a custom tab](#)).
3. Click **+ADD CONTENT**.

The **+Add Content** section is displayed and shows the available charts to choose from.

4. Select the chart that you want to add.
The chart is added to the dashboard.
5. When you are finished adding charts, click **+ADD CONTENT** to close the **+Add Content** section.
The selected charts are added to the dashboard.
6. To view a chart in full-screen mode, click the Maximize icon  in top-right corner of the chart.
The chart expands to fit the entire screen. To return to normal mode, click the Restore icon  in the top-right corner of the expanded chart.

Modifying a dashboard layout

About this task


Any changes that you make to a dashboard tab are automatically saved when you do the change.

Steps

1. To add performance charts to a performance monitor tab, click **+ADD CONTENT** and select the performance chart that you want to add.
By default, performance monitor tabs show performance charts in two columns.
2. To view more columns of performance charts on each row, select **3** or **4** from the **Columns** drop-down.
You can also select **1** to display only one column of charts. But, some performance monitors require two columns for display.

Creating a custom dashboard

Steps

1. Click **+** and select **Add Custom Dashboard**.
The **New Dashboard** tab is added.
2. For the performance monitor you would like to display on the dashboard, click **+** and select from the drop-down.
3. To provide a unique name for the new dashboard, click .
The **Rename Dashboard** dialog box is displayed.
4. Enter a new name for the New Dashboard tab.
5. Click **OK**.
6. To change the number of monitors displayed in each row of the dashboard, select the appropriate number from the **Columns** menu on the top-right corner of the dashboard.

Removing a chart

About this task

On the performance monitor tab containing the chart, in the upper-right corner of the chart, click .

Moving a chart


Steps

1. Hover over the title bar.
2. Hold the left mouse button down while dragging the chart to the new location on the dashboard tab.
3. Release the mouse button.
The chart is moved.

Back-end Bandwidth Chart

The Back-End Bandwidth chart on the Performance Dashboard shows the quantity of back-end reads and writes per second over time for directors. Generally, bandwidth (measured in KB/s or MB/s) is associated with large block I/O (64KB or greater I/O requests).

Each array type, model, and underlying hard disk drives is different. Ensure you know your array's capabilities, for example, response time, IOPS, and bandwidth.

 **NOTE:** The chart displays data only for the cluster to which you are currently connected. To simultaneously view back-end bandwidth charts for another cluster, open a second browser session and connect to the second cluster.

Guidelines

- Having a baseline of array performance is important so you know what your specific array and what its setup is capable of. When adding metro node to an existing environment, know what your host to storage array (native) performance should be. Although metro node tends to boost read-intensive workloads because of additional caching, metro node can only perform as well as this native performance.
- The underlying issue could be poor storage array performance.
- Unexpectedly low bandwidth could indicate array saturation or back-end fabric issues.
- Keep in mind that many things affect storage array bandwidth performance such as:
 - IO request size
 - Read vs. write request
 - Underlying disk type (SSD, FC, SATA)
 - RAID type — 0, 1, 0+1, 5, 6
 - FAST VP — Fully Automated Storage Tiering across 3 distinct tiers: Flash, enterprise hard disk drives (10K and 15K rpm), and high-capacity SATA HDDs
 - Thin/thick pools
 - Storage array cache settings and size
 - Running snapshots/clones
- For arrays with write-back caching, writes will generally be faster than reads. If the read data is cached by the array, then the latency will be comparable to that of the writes.
- Back-end bandwidth can be negatively affected while running snapshots/clones
- For Symmetrix arrays, performance can be impacted by Write Pending (WP) limits when there is a lack of available free cache slots to accept incoming writes forcing the array to proactively flush pages to disk. Running SRDF sessions might also affect performance.
- For CLARiiON arrays, performance can be impacted by forced flushes due to lack of available cache to buffer writes.

Corrective actions

- Check back-end errors: These indicate that metro node had to abort and re-try operations. Could be a back-end fabric and/or storage array health issue.
- Examine the back-end fabric for its overall health state, recent changes, reported errors, and properly negotiated speeds.
- Examine the back-end storage array for general health state, and that performance best practices are followed for disk/RAID layout when needed.
- Use available performance monitoring tools from the storage array vendor to confirm the array's performance. Additional metrics available to the array but not visible to metro node can confirm the problem.
- Be sure to run the recommended storage array firmware version. Check for newer software releases and known bug fixes.
- Consult your storage array vendor's performance specialists if the problem persists.

Changing the view

Use the following appropriate selection criteria to filter the data:

- **Director** — Allows you to select all directors or a specific director in the cluster.
- **Read** and **Write** check boxes — Allows you to select one or both check boxes to filter bandwidth for Reads and Writes.


Viewing the Back-end Bandwidth chart

1. From the GUI main menu, click **Performance**.
2. In the Performance Dashboard, select the tab in which you want to display the Back-end Bandwidth chart (or create a custom tab).
3. Click **+Add Content**.
4. Click the Back-end Bandwidth icon.

Back-end Throughput chart

The Back-end Throughput chart shows the back-end I/Os per second over time for directors. Generally, throughput (or more commonly referred to as IOPS) is associated with small block I/O (4KB or 16KB I/O requests.)

Each array type, model, and underlying hard disk drive is different. Ensure to know your array's capabilities, for example, response time, IOPS, and bandwidth. Having a baseline of array performance is important so you know what your specific array and its setup is capable of.

 **NOTE:** The chart displays data only for the ports in the cluster to which you are currently connected. To simultaneously view back-end throughput for another cluster, open a second browser session and connect to the second cluster.

Guidelines

- When adding metro node to an existing environment, know what your host to storage array (native) performance since although metro node tends to boost read intensive workloads because of additional caching, often times metro node can only perform just as well as this native performance.
- Metro node could simply be the victim of poor storage array performance.
- Unexpectedly low IOPS could indicate array saturation or back-end fabric issues. Recommended to check back-end aborts/retries.
- Keep in mind many things affect storage-array IOPS performance such as:
 - IO request size
 - Read vs. write request
 - Underlying disk type (SSD, FC, SATA)
 - RAID type — 0, 1, 0+1, 5, 6
 - FAST VP — Fully Automated Storage Tiering across 3 distinct tiers: Flash, enterprise hard disk drives (10K and 15K rpm), and high-capacity SATA HDDs
 - Thin/thick pools
 - Storage array cache settings and size
 - Running snapshots/clones
- For arrays with write-back caching, writes will generally be faster than reads. If the read data is cached by the array, then the latency will be comparable to that of the writes.
- Performance can be negatively affected while running snapshots or clones.
- For Symmetrix arrays, performance can be impacted by Write Pending (WP) limits when there is a lack of available free cache slots to accept incoming writes, forcing the array to proactively flush pages to disk. Running SRDF sessions might also affect performance.
- For CLARiiON arrays, performance can be impacted by forced flushes due to the lack of available cache to buffer writes.

Corrective actions

- Check back-end errors: These indicate that metro node had to abort and re-try operations. Could be a back-end fabric or storage-array health issue.
- Examine the back-end fabric for its overall health state, recent changes, reported errors, and properly negotiated speeds.
- Examine the back-end storage array for general health state, and that performance best practices are followed for disk/RAID layout when needed.
- Utilize available performance monitoring tools from the storage array vendor to confirm the array's performance. Additional metrics available to the array but not visible to metro node might shed light on the problem.
- Ensure to run the recommended storage array firmware version. Check for newer software releases and known bug fixes.
- Consult your storage array vendor's performance specialists if the problem persists.

Changing the view

Use the following appropriate selection criteria to filter the data:

- **Director** — Allows you to select all directors or a specific director in the cluster.
- **Read** and **Write** check boxes — Allows you to select one or both check boxes to filter throughput for Reads and Writes.

Viewing the Back-end Throughput chart

1. From the GUI main menu, click **Performance**.
2. In the Performance Dashboard, select the tab in which you want to display the Back-end Throughput chart (or create a custom tab).
3. Click **+Add Content**.
4. Click the Back-end Throughput icon.

Back-end Errors chart

The **Back-end Errors** chart displays the back-end I/O errors to and from the storage array. There are three categories of back-end errors:

- **Aborts** — Indicate the metro node back-end gave up and aborted the I/O operation to the storage array, the array itself decided to abort the I/O operation, or another SCSI initiator (metro node director, or host) connected to the array caused the I/O to abort.
- **Timeouts** — Indicate the metro node back-end saw an I/O operation to a storage volume that has not completed within 10 seconds.
- **Resets** — Logical Unit resets issued by the metro node back-end to a storage volume as corrective action when after 20 seconds of no response for any I/O by the storage volume (Logical Unit on the storage array). The metro node back-end re-tries all outstanding I/O to the storage volume.

NOTE:

The chart displays data only for the cluster to which you are currently connected. To simultaneously view back-end errors for another cluster, open a second browser session and connect to the second cluster.

Guidelines

- Back-end errors typically indicate back-end fabric and/or storage array issues.
- For a normal healthy system, there should be no aborts, timeouts, or resets.
 - Timeouts might happen during bursts of I/O to a storage array. Seeing a few of these is generally not bad for performance, however, frequent or periodic timeouts are not normal.
 - Aborts and Resets likely indicate major performance issues on the storage fabric or storage array.
 - Investigate the cause for back-end errors immediately.

Corrective actions

- Look closely at the latency related categories (front-end read/write latency, and back-end read/write latency) for any high averages or large spikes. Try to correlate the spikes to the errors.
- Examine the back-end fabric for changes, reported errors, proper negotiated speeds, and health state.
- Examine the back-end storage array for general health state, and best practices disk/volume layout.
- Check the metro node firmware log for events indicating command timeouts, retries, or other general back-end health issues.

Changing the view

Use the following appropriate selection criteria to filter the data:

- **Director** — Allows you to select all directors or a specific director in the cluster.

- **Aborts** — Select only this box to view only the aborts. Select this box with other boxes to view the aborts and other statistics. Clear this box to stop viewing aborts.
- **Resets** — Select only this box to view only the resets. Select this box with other boxes to view the resets and other statistics. Clear this box to stop viewing resets.
- **Timeouts** — Select only this box to view only the timeouts. Select this box with other boxes to view the timeouts and other statistics. Clear this box to stop viewing timeouts.

Viewing the Back-end Errors chart

1. From the GUI main menu, click **Performance**.
2. In the Performance Dashboard, select the tab in which you want to display the Back-end Errors chart (or create a custom tab).
3. Click **+Add Content**.
4. Click the Back-end Errors chart icon.

Back-end Latency chart

The Back-end Latency chart on the performance dashboard provides a time-based view of the back-end I/O latency. The chart allows you to:

- View back-end latency for a selected director or all directors
- Change the view to data received only, data sent only, or both (default)
- Place the mouse over any point on the Read or Write data latency indicators to view details for a specific director

Each array type, model, and underlying hard disk drive is different. Make sure you know your array's capabilities: response time, IOPS, and bandwidth. Having a baseline of array performance is important so you know what your specific array and its setup is capable of.

Keep in mind that many things affect storage array latency, such as:

- IO request size
- Read vs. write request

NOTE: The chart displays data only for the ports in the cluster to which you are currently connected. To simultaneously view back-end latency for another cluster, open a second browser session and connect to the second cluster.

Guidelines

- When adding metro node to an existing environment, know what your host to storage array (native) performance. Although metro node tends to boost read intensive workloads because of additional caching, metro node can perform only as well as this native performance.
- High latency could indicate poor storage array performance.
- High latency due to low bandwidth could indicate array saturation or back-end fabric issues.
- In a metro node Metro system, this metric represents the local storage volume write latency, and does not reflect the time spent sending data to or receiving data from a remote cluster.
- For arrays with write-back caching, writes will generally be faster than reads. If the read data is cached by the array, then the latency will be comparable to that of the writes.
- High latency may be present while running snapshots or clones
- Keep in mind that many things affect storage array IOPS performance, such as:
 - Underlying disk type (SSD, FC, SATA)
 - RAID type — 0, 1, 0+1, 5, 6
 - FAST VP — Fully Automated Storage Tiering across 3 distinct tiers: Flash, enterprise hard disk drives (10K and 15K rpm), and high-capacity SATA HDDs
 - Thin/thick pools
 - Storage array cache settings and size
 - Running snapshots/clones

VNX guidelines

- <10msec — Great
- 10-20 msec — Decent
- 20-100 msec — Not so great
- >100 msec — Poor performance
- High latency could cause forced flushes due to lack of available cache to buffer writes.

Symmetrix guidelines

- <1 msec — Great
- 1-10 msec — Decent
- 10-50 msec — Not so great
- >50 msec — Poor performance
- For Symmetrix arrays, performance can be impacted by Write Pending (WP) limits when there is a lack of available free cache slots to accept incoming writes forcing the array to proactively flush pages to disk. Running SRDF sessions might also affect performance.

Corrective actions

- Check back-end errors: These indicate that metro node had to abort and re-try operations. Could be a back-end fabric and/or storage array health issue.
- Examine the back-end fabric for its overall health state, recent changes, reported errors, and properly negotiated speeds.
- Examine the back-end storage array for general health state, and that performance best practices are followed for disk/RAID layout when needed.
- Utilize available performance monitoring tools from the storage-array vendor to confirm the array's performance: Additional metrics available to the array not visible to metro node may shed light on the problem.
- Be sure to run the recommended storage array firmware version. Check for newer software releases and known bug fixes.
- Engage your storage array vendor's performance specialists if the problem persists.

Changing the view

Use the following appropriate selection criteria to filter the data:

- **Read** — Displays latency statistics pertaining to reads only.
- **Write** — Displays latency statistics pertaining to writes only.
- **Director** — Displays data for all directors or a specific director in the cluster.

Viewing the Backend Latency chart


1. From the GUI main menu, click **Performance**.
2. In the Performance Dashboard, select the tab in which you want to display the Back-end Latency chart (or create a custom tab).
3. Click **+Add Content**.
4. Click the Back-end Latency chart icon.

CPU utilization chart

The CPU Utilization chart provides a time-based view of the utilization load on the director CPUs of your metro node system. By default, the chart shows an averaged view of the utilization loads of all of the CPUs on all the directors in your metro node system. This category is typically the first one to check for problems. When the CPU reaches 100% busy, its I/O processing capability will also peak.

If your CPU utilization is running near 100% at all times, you have no spare capacity to handle a peak load. In applications that are latency-sensitive and require fast response times, high director CPU usage can reduce response times even though throughput and I/O processing capability might stay constant. Spikes in CPU utilization typically correlate to increased I/O load,

often typical of back-ups or batch processing jobs. These spikes are not necessarily a sign of a problem. However, monitor their height and duration, and their impact on response time.

 **NOTE:** The chart displays data only for the directors in the cluster to which you are currently connected. To simultaneously view statistics for another cluster, open a second browser session and connect to the second cluster.

Guidelines

Keep in mind the following guidelines when using the CPU Utilization chart:

- A completely idle system will be around 0-5% busy
- Less than 50% CPU utilization is safe
- CPU utilization of 50-80% warning
- CPU utilization of 80-90% danger
- CPU utilization of >90% critical

Corrective actions

- Consider migrating virtual volumes to less busy directors (moving them to different storage views).
- Consider adjusting the host failover driver settings to better load balance I/O across available directors.
- Stagger batch application start times causing troublesome CPU spikes (such as modify your back-up jobs schedule to start with some delay between jobs instead of starting all back-up jobs at once, if possible.)
- Look for signs of perhaps a rogue director firmware process or thread hogging CPU time.
- If normal baseline busy is constantly high (due to high I/O load), consider an engine expansion, or a VS1 to VS2 hardware upgrade, if possible.

Changing the view

To view the CPU utilization of a single director in your metro node system, select the director name from the **Director** drop-down.

Viewing the CPU Utilization chart

1. From the GUI main menu, click **Performance**.
2. In the Performance Dashboard, select the tab in which you want to display the CPU Utilization chart (or create a custom tab).
3. Click **+Add Content**.
4. Click the CPU Utilization chart icon.

Heap Usage chart

The **Heap Usage** chart shows a percentage of the heap memory used by the firmware on a director. In GeoSynchrony 5.x, the heap size is 6GB.

The greater the number of 4KB pages storing user data that the metro node director cache needs to keep track of, the greater the heap usage. There are safeguards in place so that the amount of information recorded about pages known by the cache will not exceed the heap size. If the amount of information recorded gets too large, the cache layer will remove the cached pages and knowledge about them from its cache.

During failure recovery, the amount of heap usage initially increases, and then decreases back to normal levels once the failure is handled. For this reason, it's important that a director's heap not get too full (< 85-90%) to account for this.

Guidelines

Keep in mind the following guidelines when using the **Heap Usage** chart:

- A value of around 40 to 50% is normal.


- Check if CPU utilization and front-end operations count is also high. The system might be simply at over-capacity, which may also be reflected in this category.
- Constantly rising heap usage is a potential sign of a director memory leak. Contact Dell EMC Support to identify these systems.

Corrective actions

- Check if CPU utilization and front-end operations count is also high. The system might be simply at over-capacity, which may also be reflected in this category.
- Constantly rising heap usage is a potential sign of a director memory leak. Contact Dell EMC Support to identify these systems.

Changing the view

To view the Heap usage of a single director in your metro node system, select the director name from the **Director** drop-down.

 **NOTE:** The chart displays data only for the directors in the cluster to which you are currently connected. To simultaneously view statistics for another cluster, open a second browser session and connect to the second cluster.

Viewing the Heap Usage chart

1. From the GUI main menu, click **Performance**.
2. In the Performance Dashboard, select the tab in which you want to display the chart (or create a custom tab).
3. Click **+Add Content**.
4. Click the **Heap Usage** chart icon.

Front-end Queue Depth chart

The **Front-end Queue Depth** chart provides the count of front-end operations per director. It describes the number of concurrent outstanding operations active in the system. This statistic is not directly related to front-end IOPS. IOPS could be easily be in the thousands range, whereas this field describes the concurrency of the operations and may be quite low, depending upon how many I/O's the application or host have outstanding to metro node over a given time.

Guidelines

Keep in mind the following guidelines when using the Front-end Queue Depth chart:

- Normal values are approximately 10-100 concurrent operations for a small to medium system.
- Anything greater than about 500 shows a lot of concurrent access in a system but is not alarming. Check the CPU utilization.
- > 1000 active could be indicative of a problem with metro node keeping up with host I/O demands. Check CPU utilization, front-end aborts, and be-aborts
- This value depends heavily upon the customer workload.

Corrective actions

- Monitor CPU % busy, front-end aborts, back-end aborts, and WAN health state (in metro node Metro systems).
- If other directors in a cluster are less busy, look to balance the workload across these directors.
- SCSI CDB trace information can identify if there are large dormant queues in the front-end. This can be bad for performance.
- For metro node Metro systems, WAN performance degradation could cause active operations to distributed-devices to not complete as quickly as they did previously. This may result in an I/O operations queue build up, causing this counter to increase. Check the WAN connectivity, health state, and performance.
- In VMware and certain host clustered environments, SCSI reservations could causing spikes in active operations, if there are a lot of reservation-intensive operations (vMotion, Storage vMotion, backup, or frequent file creations/deletions).

Changing the view

To view the Front-end Queue Depth of a single director in your metro node system, select the director name from the **Director** drop-down.

NOTE: The chart displays data only for the directors in the cluster to which you are currently connected. To simultaneously view Front-end Queue Depth for another cluster, open a second browser session and connect to the second cluster.

Viewing the Front-end Queue Depth chart

1. From the GUI main menu, click **Performance**.
2. In the Performance Dashboard, select the tab in which you want to display the Front-end Queue Depth chart (or create a custom tab).
3. Click **+Add Content**.
4. Click the Front-end Queue Depth chart icon.

Front-end Bandwidth chart

The **Front-End Bandwidth** chart on the Performance Dashboard displays the quantity of front-end reads and writes per second over time for directors on your metro node system. Generally bandwidth (measured in KB/s or MB/s) is associated with large block I/O (64KB or greater I/O requests).

NOTE: The chart displays data only for the cluster to which you are currently connected. To simultaneously view front end bandwidth charts for another cluster, open a second browser session and connect to the second cluster.

Guidelines

- Front-end performance should be compared to baseline numbers (native host to storage-array) when metro node performance issues arise. The underlying problem could be poor storage array performance. When you add metro node to your environment, know what your application throughput was beforehand.
- Front-end performance in metro node depends heavily upon the available back-end storage-array performance, and in Metro configurations, the WAN performance for distributed devices.
- Any running distributed rebuilds or data migrations might negatively affect available host throughput.
- Since metro node Local and Metro implement write through caching, naturally a small amount of write latency overhead (typically <1msec) is expected with metro node. This latency may affect applications that serialize their I/O and don't take advantage of multiple outstanding operations. These types of applications may see a throughput and IOPS drop with metro node in the data path.
- Understand that in a metro node Metro environment you incur extra WAN round-trip time on your write latency since writes need to be successfully written to each cluster's storage before the host is acknowledged. This extra latency could impact the throughput and IOPS of serialized-type applications.

Corrective actions

- Check CPU busy: If overly busy, metro node will be limited on the amount of bandwidth it can provide.
- Check back-end latency: If on average the back-end latency is large, or there are large spikes, there could be a poorly performing back-end fabric or an unhealthy, un-optimized, or over-loaded storage-array. Perform a back-end fabric analysis, and a performance analysis of all storage-arrays in question.
- Check front-end aborts: The presence of these indicate that metro node is taking too long to respond to the host. These might indicate problems with the front-end fabric, or slow SCSI reservations.
- Check back-end errors: If the metro node back-end is required to retry an operation because it is aborted, then this will add to the delay in completing the operation to the host.
- Check front-end queue depth: If this counter is large, this may explain larger than normal front-end latency. Follow front-end operations count corrective actions.
- Check metro node write delta time: If the time spent within metro node is more than usual, attempt to find out why. See corrective actions for write delta time.

- Check for bandwidth or IOPS over-provisioned metro node front-end ports. Be sure to balance hosts and LUNs across the available directors and front-end ports presented from metro node. Check the front-end fabric for saturation or over-capacity.
- Verify that front-end Fibre Channel ports, HBAs, and switch ports are configured to the proper port speeds.
- For your host multipathing software, configure them based on metro node best practices, and make sure the installed software versions are compatible with metro node (see the *Simple Support Matrix for metro node* document, available on Dell EMC Online Support.)

Corrective actions (metro node Metro only)

- For metro node Metro configurations, check the inter-cluster link health and maximum performance capabilities. From the metro node Management GUI, check the observed inter-cluster WAN bandwidth. If your application throughput appears low to you, and only seems to achieve something similar to what the WAN bandwidth reports, then chances are you are limited by the WAN. Therefore:
 - Ensure you have provisioned enough inter-cluster bandwidth for the desired application workload. Verify that your WAN configuration is supported by metro node (minimum supported bandwidth, supported inter-cluster latency, compatible WAN hardware and software).
 - If using Fibre Channel devices over dark fibre or DWDM, confirm that you have allocated enough buffer credits or configured the Fibre Channel WAN ports properly on your switches. Check for buffer credit starvation, c3 discards, and CRC errors. Some vendors may require extended fabric licenses to enable WAN features.
 - Validate your WAN performance before going live in production. Create multiple test distributed-devices and force them to rebuild. Observe the performance of the rebuilds.
 - When troubleshooting distributed device performance, if feasible, check local device performance. Export a test LUN from your storage array to metro node, then to the host, and then run a test I/O workload.
 - Check for any unexpected local or distributed rebuilds or data migrations. There will be some amount of performance impact to host application traffic that relies on the same virtual volumes and storage volumes. Tune the rebuild transfer-size setting using the CLI to limit rebuild/migration performance impact. Consider scheduling migrations during off-peak hours.

Changing the view

Use the following appropriate selection criteria to filter the data:

- **Director** — Allows you to select all directors or a specific director in the cluster.
- **Read** and **Write** check boxes — Allows you to select one or both check boxes to filter bandwidth for Reads and Writes.

Viewing the Front-end Bandwidth chart

1. From the GUI main menu, click **Performance**.
2. In the Performance Dashboard, select the tab in which you want to display the **Front-end Bandwidth** chart (or create a custom tab).
3. Click **+Add Content**.
4. Click the **Front-end Bandwidth** icon.


Front-end Latency chart

The **Front-end Latency** chart on the Performance Dashboard provides details of the front-end I/O latency statistics for your metro node system.

Front-end latency is defined as the amount of time an I/O spends within metro node. The reported metro node front-end latency should match closely to the host or application reported latency, unless there is significant added delay in the front-end fabric, HBA, multi-pathing software, or host operating system software.

It is important to distinguish how front-end latency write operations behave in the two metro node configurations:

- For metro node Local write operations, it includes the time spent protecting the disk blocks to one or more local storage-arrays.
- For metro node Metro write operations, it includes the time spent protecting the disk blocks to the storage-array at BOTH clusters. When writing to the remote cluster, the round-trip time on the WAN links adds to the front-end latency, depending upon the observed network delay between clusters.

 **NOTE:** The chart displays data only for the directors in the cluster to which you are currently connected. To simultaneously view front-end port latency for another cluster, open a second browser session and connect to the second cluster.

Guidelines

It is difficult to give absolute recommended values for front-end latencies since they are typically highly dependent upon back-end latencies.

- Metro node processing overhead on read misses and on write operations (in a metro node Local system) is roughly 1msec.
- Be aware of large spikes in latency, which might correlate to front-end or back-end errors, or to a large front-end queue depth.

Corrective actions

- Check CPU busy: If overly busy, the time for metro node to respond to I/O will increase.
- Check back-end latency: If on average the back-end latency is large, or there are large spikes, there could be a poorly performing back-end fabric or an unhealthy, un-optimized, or over-loaded storage array. Perform a back-end fabric analysis, and a performance analysis of all storage arrays in question.
- Check front-end aborts: The presence of these indicate that metro node is taking too long to respond to the host. These might indicate problems with the front-end fabric, or slow SCSI reservations. See front-end aborts corrective actions.
- Check back-end errors: If the metro node back-end is required to retry an operation because it is aborted, then this will add to the delay in completing the operation to the host. See back-end errors corrective actions.
- Check front-end queue depth: If this counter is large, this may explain larger than normal front-end latency. Follow front-end operations count corrective actions.
- Check metro node write delta time: If the time spent within metro node is more than usual, attempt to find out why. See corrective actions for write delta time.

The **Front-end Latency** chart allows you to:

- View front-end latency for a selected director or all directors
- Change the view to Read data only, Write data only, or both (default)
- Place the mouse over any point on the sent or received data indicators to view details for a specific director

Changing the view

Use the following appropriate selection criteria to filter the data:


- **Read** - Displays latency statistics pertaining to reads only.
- **Write** - Displays latency statistics pertaining to writes only.
- **Director** - Displays data for all directors or a specific director in the cluster.

Viewing the Front-end Latency chart

1. From the GUI main menu, click **Performance**.
2. In the Performance Dashboard, select the tab in which you want to display the Front-end Latency chart (or create a custom tab).
3. Click **+Add Content**.
4. Click the Front-end Latency chart icon.

Front-end Throughput chart

The **Front-End Throughput** chart on the Performance Dashboard displays the front-end I/Os per second over time for directors on your metro node system. Generally throughput (more commonly referred to as IOPS) is associated with small block I/O (512B to 16KB I/O requests.)

 **NOTE:** The chart displays data only for the ports in the cluster to which you are currently connected. To simultaneously view front-end throughput for another cluster, open a second browser session and connect to the second cluster.

Guidelines

- Front-end performance should be compared to baseline numbers (native host to storage array) when metro node performance issues arise. Poor throughput might be caused by poor storage array performance. Before adding metro node to your environment, know what your application throughput was beforehand.
- Front-end performance in metro node depends heavily upon the available back-end storage array performance, and in metro node Metro configurations, the WAN performance for distributed-devices.
- Any running distributed rebuilds or data migrations could negatively affect available host throughput.
- Since metro node Local and Metro implement write through caching, a small amount of write latency overhead (typically <1msec) is expected with metro node. This latency could affect applications that serialize their I/O and do not take advantage of multiple outstanding operations. These types of applications may see a throughput and IOPS drop with metro node in the data path.
- Understand that in a metro node Metro environment you will incur extra WAN round-trip time on your write latency since writes need to be successfully written to each cluster's storage before the host is acknowledged. Again, this extra latency could impact the throughput and IOPS of serialized-type applications.

Corrective actions

- Check CPU busy: If the CPU is overly busy, metro node becomes limited on the amount of throughput it can provide.
- Check back-end latency: If on average the back-end latency is large, or there are large spikes, there could be a poorly performing back-end fabric or an unhealthy, un-optimized, or over-loaded storage-array. Perform a back-end fabric analysis, and a performance analysis of all storage-arrays in question.
- Check front-end aborts: The presence of these indicate that metro node is taking too long to respond to the host. This might indicate problems with the front-end fabric, or slow SCSI reservations.
- Check back-end errors: If the metro node back-end is required to retry an operation because it is aborted, then this will add to the delay in completing the operation to the host.
- Check front-end queue depth: If this counter is large, this may explain larger than normal front-end latency. Follow front-end operations count corrective actions.
- Check metro node write delta time: If the time spent within metro node is more than usual, attempt to find out why. See corrective actions for write delta time.
- Verify the front-end average I/O Size, and confirm that you are sending small block I/O if you are trying to boost IOPS performance.
- Check for bandwidth/IOPS over-provisioned metro node front-end ports: Be sure to balance hosts and LUNs across the available directors and front-end ports presented from metro node. Check the front-end fabric for saturation/over-capacity.
- Verify front-end Fibre Channel ports, HBAs, and switch ports are configured to the proper port speeds.
- For host multipathing software, configure ports based on metro node best practices, and ensure the installed software versions are compatible with metro node (see the metro node Simple Support Matrix.)

Metro node Metro

- For metro node Metro configurations, check the inter-cluster link health and maximum performance capabilities. From the metro node Management GUI, check the observed inter-cluster WAN bandwidth. If your application throughput appears low, and only seems to achieve something similar to what the WAN bandwidth reports, then chances are you are limited by the WAN. Therefore:
 - Make sure you have provisioned enough inter-cluster bandwidth for the desired application workload. Verify that your WAN configuration is supported by metro node (minimum supported bandwidth, supported inter-cluster latency, compatible WAN hardware and software).
 - If using Fibre Channel devices over dark fibre or DWDM, confirm that you have allocated enough buffer credits or configured the Fibre Channel WAN ports properly on your switches. Check for buffer credit starvation, c3 discards, and CRC errors. Some vendors may require extended fabric licenses to enable WAN features.
 - Validate your WAN performance before going live in production. Create multiple test distributed-devices and force them to rebuild. Observe the performance of the rebuilds.
- When troubleshooting distributed-device performance, if feasible, check local device performance. Export a test LUN from your storage array to metro node then to the host, and run a test I/O workload.
- Check for any unexpected local or distributed rebuilds or data migrations. There will be some amount of performance impact to host application traffic that relies on the same virtual-volumes and storage volumes. Tune the rebuild transfer-size setting using the CLI to limit rebuild/migration performance impact. Consider scheduling migrations during off-peak hours.

Changing the chart view

Use the following appropriate selection criteria to filter the data:

- **Director** — Allows you to select all directors or a specific director in the cluster.
- **Read and Write check boxes** — Allows you to select one or both check boxes to filter throughput for Reads and Writes.


Viewing the Front-end Throughput chart

1. From the GUI main menu, click **Performance**.
2. In the Performance Dashboard, select the tab in which you want to display the **Front-end Throughput** chart (or create a custom tab).
3. Click **+Add Content**.
4. Click the Front-end Throughput chart icon.

Front-end Aborts chart

The **Front-end Aborts** chart displays the number of aborts per second over time for directors on your metro node system.

The presence of front-end aborts indicates the host gave up on I/Os submitted to metro node and aborted them, likely due to timeouts. This is bad under normal operations but acceptable during unusual circumstances such as metro node director failure recovery or non-disruptive-upgrades. These are only the symptom of poor host response time.

 **NOTE:** The chart displays data only for the cluster to which you are currently connected. To simultaneously view front-end aborts for another cluster, open a second browser session and connect to the second cluster.

Guidelines

A presence of these typically indicates host performance issues; however a small rate of these can be normal for some systems.

Corrective actions

- Monitor the CPU % busy, back-end errors, and WAN health state (in metro node Metro).
- In a metro node Metro, look at the WAN latency for irregularities.
- Look closely at the latency related categories (front-end read/write latency, and back-end read/write latency) for any high averages or large spikes. Try to correlate the spikes to the aborts.
- Look at metro node write delta time if larger than normal.
- Look at back-end errors, and if they happen to correlate to front-end aborts. If this is the case, examine the back-end fabric and storage array health.
- Examine the front-end fabric for changes, reported errors, properly negotiated speeds, and health state.

Changing the chart view

Use the following appropriate selection criteria to filter the data:

- **Director** - Allows you to select all directors or a specific director in the cluster.

Viewing the Front-end Aborts chart

1. From the GUI main menu, click **Performance**.
2. In the Performance Dashboard, select the tab in which you want to display the **Front-end Aborts** chart (or create a custom tab).
3. Click **+Add Content**.
4. Click the Front-end Aborts chart icon.

Write Latency Delta chart

This chart provides the delta between Front-end latency and Back-end Latency per director. This is a key metric for Local/Metro — the amount of overhead time metro node spends processing a write.

i **NOTE:** The chart displays data only for the directors in the cluster to which you are currently connected. To simultaneously view Write Latency Delta for another cluster, open a second browser session and connect to the second cluster.

Guidelines (metro node Local only)

- This is typically <1 msec
- 1-10 msec is concerning
- >10 msec should be investigated.
- Large spikes indicate metro node spent a long time adding processing time to the I/O — This is generally bad for performance.
- During disruptive events, this value may go as high as a few seconds, which could be due to I/Os timing out in the back-end.
- Note that in a heavily loaded system, this value climb as the directors become more busy, especially if the directors approach saturation (>95% director CPU busy).
- Large block writes (host I/O block sizes >128KB) add to extra metro node write processing time.

Guidelines (metro node Metro only)

- For distributed devices in write-through mode, this statistic includes the time required to write to the remote leg of the distributed Raid (WAN latency plus remote storage array write latency). The write is not acknowledged to the host without first writing to both legs of the distributed Raid, thus we only perform as fast as the slowest leg which in most cases will be the remote leg on the other cluster. This is why WAN performance is so critical. It is good to also check WAN round-trip latency.
- During disruptive events, this value may go as high as a few seconds, which could be due to I/Os timing out in the back-end, or a poorly performing WAN link.

Corrective actions

- Check CPU load: If CPU load is too high, this will mean less CPU cycles available to process data.
- Check back-end errors, and back-end latency: If the metro node back-end is aborting I/Os, must re-try them and this wait time is reflected in the metro node write delta time.
- Metro node Metro: Check WAN latency. The time spent writing to the remote leg of a distributed Raid is encapsulated in this write delta time.

Changing the view


To view the **Write Latency Delta** of a single director in your metro node system, select the director name from the **Director** drop-down.

Viewing the Write Latency Delta chart

1. From the GUI main menu, click **Performance**.
2. In the Performance Dashboard, select the tab in which you want to display the **Write Latency Delta** chart (or create a custom tab).
3. Click **+Add Content**.
4. Click the **Write Latency Delta** icon.

WAN Port Performance chart

The WAN Port Performance chart on the Performance Dashboard shows the WAN port performance by IP or Fibre Channel port for the cluster to which you are connected.

 **NOTE:** The chart displays data for the cluster to which you are connected. To simultaneously view the statistics for a remote cluster, open a second browser session and connect to the remote cluster's management console.

Guidelines

- The absolute value of the bytes (or bandwidth) and packet (or throughput) related categories might not indicate a problem. They will reflect the amount of traffic requested by other internal metro node firmware components, which is ultimately the host application requested load.
- If the observed maximums appear to be much lower than the available inter-cluster bandwidth or what you expect to see from your application, this is worth investigating.
- FC-COM specific: If Fibre Channel switches are used over dark fibre or DWDM WAN equipment, ensure that the WAN facing Fibre Channel ports have sufficient buffer credits allocated to the ports. A lack of buffer credits imposes an undesired limit on the maximum throughput on the WAN port.
- For Brocade switches:
 - An extended fabric license is required for each edge switch, and the WAN facing ports must be set to LS or LD mode. See the `portcfglongdistance` command.
 - Monitor the port's counters for non-zero values for `tim_txcrd_z` or time transmission credits are zero. This means the Fibre Channel port wanted to transmit a Fibre Channel packet, but did not have sufficient buffer credits to do so. Any non-zero value in this category implies performance issues on the WAN port.
 - If FCIP gateway devices are used between metro node clusters, ensure that the FCIP tunnel is configured properly.
- For Brocade FCIP switches:
 - Check for bandwidth rate limiting setting on the tunnel. See the `portshow fcip tunnel` command. Verify that the values for Minimum/Maximum Communication Rate are not causing a bottleneck.
 - Check for improper QoS settings on the tunnel. In the `portshow fcip tunnel` command output, check the values for QoS Percentages. Note that only if QoS has been set on the LAN facing Fibre Channel ports will QoS settings affect the fcip tunnel settings.

Corrective actions

- Check the ip-com and fc-com ports health on the directors.
- Check the WAN products for improperly configured ports, link errors, packet loss, QoS limitations, or large observed round-trip-times outside of supported metro node specs.
- If you suspect that you are not achieving optimal inter-cluster WAN throughput, verify the available inter-cluster bandwidth.
- Check the WAN latency using the `director ping` CLI command.
- Check the WAN Latency chart or the monitor category **com-cluster-io** for a history of WAN latency.
- Follow the WAN products' performance best practices.
- Find out from your WAN provider if anything has changed recently.
- WAN COM performance for remote read or remote write operations for distributed-devices and remote-devices are also limited to the storage array's performance. Make sure that the back-end performance on your remote cluster is not your bottleneck.
- IP-COM specific:
 - Check the network's maximum supported MTU (maximum transmission unit) size. Standard networks support 1500 byte MTUs. Jumbo frames can support up to about 9000 byte MTUs, which typically result in faster throughputs and consumes fewer CPU cycles. Jumbo frames however must be explicitly enabled on every device in the network path (`director tracepath` CLI command shows the number of hops.) WAN providers are typically reluctant to enable jumbo frames.
 - Check the WAN Message Size chart.
 - Check the port-group's socket buffer size (`socket-buf-size` in `/clusters/cluster-#/cluster-connectivity/option-sets/optionset-com-#/`). The default is set to 1MB. The optimal value for this is the network's delay-bandwidth product which is the latency or delay of the network multiplied by the available bandwidth, which is the amount of data required to be outstanding to fully utilize the network.
 - If this value is too large for a low bandwidth connection, there will be connection keepalive timeouts and possibly cluster departure events. Lower the `socket-buf-size` in this case.
 - If this value is too low for a high bandwidth connection, there won't be enough in-flight data outstanding on the WAN port to fully utilize its available capacity, and the observed performance will be much less than expected. Increase the `socket-buf-size` to the minimum delay-bandwidth product.

NOTE: The per ip-com-port statistics will not add up perfectly to the director-level statistics category. This is because of a very small amount (less than a MB/s typically) of IP port traffic sent over the director's management port which is not tracked in this chart, but is tracked on the director level UDT category.

Changing the view

Use the following appropriate selection criteria to filter the data:.

- **Director** — Displays data for all directors or a specific director in the cluster.
- **Port** — Displays data for all ports or for a specific port.
- **Bytes** — Displays the data in bytes.
- **Packets** — Displays the data in packets.
- **Sent** — Displays only the data sent over this WAN port.
- **Received** — Displays only the data received over this WAN port.

Viewing detailed data for a specific point in time

Place the mouse pointer on any point along the performance chart to display a popup box that shows detailed information about the port performance.

Viewing the WAN Port Performance chart

1. From the GUI main menu, click **Performance**.
2. In the Performance Dashboard, click **+Add Content**.
3. Click the **WAN Port Performance** chart icon.

WAN Latency chart

The **WAN Latency** chart provides a time-based view of the WAN Latency. The categories avg-lat/min-lat/max-lat each report values observed in the last 5 seconds or less. You can view, delete.

- *Average Latency* reflects the quality of service over a WAN COM port, and not just the WAN link. Factors influencing this are the wire latency, COM latency increases because of potential COM queue build-up, re-transmitted packets due to packet drops on the wire, or connection drops. Re-transmitted packets or connection drops can increase the maximum and average latency values. These values normalize once the packet drops are no longer occurring.
- *Minimum Latency* is the best indication of underlying true wire round-trip-time (RTT) latency. It is best to rely upon average latency as a true feel for the inter-cluster latency.

NOTE: The chart displays data only for the directors in the cluster to which you are currently connected. To simultaneously view **WAN Latency Delta** for another cluster, open a second browser session and connect to the second cluster.

Guidelines

- Know your expected inter-cluster round-trip time. This value should be similar to minimum latency, however, keep in mind that COM queuing effects or potential inter-cluster bandwidth saturation can result in average latency being higher than the minimum.
- It is not unusual for some observed maximum latencies to exceed the maximum supported round-trip-time values as defined by metro node product specifications (Metro <5 msec). So long as the averages are below these defined limits, the system is alright. However, watch for consistently high maximum values which may be an indication of WAN health issues.
- For Metro-Fibre Channel:
 - If Fibre Channel switches are used over dark fibre or DWDM WAN equipment, ensure that the WAN facing Fibre Channel ports have sufficient buffer credits allocated to the ports. A lack of buffer credits results in higher than expected WAN round-trip times, negatively affecting metro node's ability to communicate and send read or write requests between clusters.
- For Brocade switches:
 - An extended fabric license is required for each edge switch, and the WAN facing ports must be set to LS or LD mode. See the command `portcfglongdistance`.

- Monitor the port's counters for non-zero values for `tim_txcrd_z` or time transmission credits are zero. This means the FC port wanted to transmit a FC packet, but did not have sufficient buffer credits to so. Any non-zero value in this category implies performance issues on the WAN link.
- If FCIP gateway devices are used between metro node clusters, ensure that the FCIP tunnel is configured properly.
- For Brocade FCIP switches:
 - Check for bandwidth rate limiting setting on the tunnel. See the `portshow fcipunnel` command. Verify the values for Minimum/Maximum Communication Rate are not causing a bottleneck.
 - Check for improper QoS settings on the tunnel. From the `portshow fcipunnel` command output, check the values for QoS Percentages. Note that only if QoS has been set on the LAN facing Fibre Channel ports will QoS settings affect the fcipunnel settings.

Corrective actions

- Check the WAN devices for improperly configured ports, link errors, packet loss, QoS limitations, or large observed round-trip-times outside of supported metro node specs.
- Check the metro node WAN ports for issues.
- Follow the WAN products' performance best practices.
- If you suspect that you are not achieving optimal inter-cluster WAN throughput, verify the available inter-cluster bandwidth.
- Find out from your WAN provider if anything has changed recently.
- For remote read or remote write operations are limited to the remote cluster's storage-array's performance. Make sure that the back-end performance on your remote cluster is not your bottleneck.
- From the WAN endpoint devices, check the WAN link latency statistics for irregular round-trip times.

Corrective actions for Metro-IP


- Check the network's maximum supported MTU (maximum transmission unit) size. Standard networks support 1500 byte MTUs. Jumbo frames can support up to about 9000 byte MTUs, which typically result in faster throughputs and consumes fewer CPU cycles. Jumbo frames however must be explicitly enabled on every device in the network path
- Check the port-group's socket buffer size (`socket-buf-size` in `/clusters/cluster-#/cluster-connectivity/option-sets/optionset-com-#/`). The default is set to 1MB. The optimal value for this the network's delay-bandwidth product which is the latency or delay of the network multiplied by the available bandwidth, which is the amount of data required to be outstanding to fully utilize the network.
- If this value is too large for a low bandwidth connection, there will be connection keep alive timeouts and possibly cluster departure events. Lower the `socket-buf-size` in this case.
- If this value is too low for a high bandwidth connection, there won't be enough in-flight data outstanding on the WAN link to fully utilize it's available capacity, and the observed performance will be much less than expected. Increase the `socket-buf-size` to the minimum delay-bandwidth product.

Viewing the WAN Latency chart

1. From the GUI main menu, click **Performance**.
2. In the Performance Dashboard, select the tab in which you want to display the **WAN Latency** chart (or create a custom tab).
3. Click **+Add Content**.
4. Click the **WAN Latency** icon.

Rebuild Status dashboard

The **Rebuild Status** dashboard displays the status of any rebuilds or migration operations that are running on your metro node system. Rebuilds synchronize data from a source drive to a target drive. When differences arise between legs of a RAID-1 device (local or distributed), a rebuild updates the out-of-date leg.

 **NOTE:** The dashboard displays data only for the directors in the cluster to which you are currently connected. To simultaneously view Rebuild Status for another cluster, open a second browser session and connect to the second cluster.

Guidelines

- Local device rebuild rates are dictated by two factors: The available read performance of the source device, and the write performance of the target device.
- Distributed device rebuild rates are dictated by three factors: The available inter-cluster bandwidth, the read performance of the source device, and the write performance of the target device.
- Be aware of the potential performance impact of rebuilds on host I/O traffic. Tune the transfer-size of the rebuild accordingly. The default transfer-size is 128KB. Choose this value carefully depending upon your requirements:
- A smaller transfer size has less of an impact on host I/O, but results in longer rebuild times.
- A larger transfer size allows rebuilds to complete faster, however they cause more of an impact to host I/O and have the potential to starve out host I/O to a point of data unavailability if the host application is particularly sensitive.
- Dell EMC recommends that you set this to 128KB as a balance between minimal host impact and acceptable rebuild/migration performance.

Corrective actions

- Tune the RAID- local device's or distributed device's transfer-size appropriately for your environment and requirements.
- For distributed device rebuilds, verify the WAN inter-cluster health and available bandwidth.
- Verify the source and target storage array health and performance.
- The system can process up to 25 concurrent rebuilds. If this is too many and overwhelms one or more storage arrays or impacts host applications, scale back the number of concurrent rebuilds.

Filtering the Rebuild Status dashboard

You can filter the rebuild status display for:

- An individual cluster
- Distributed devices
- All rebuilds

This dashboard shows the following information about rebuilds. To see additional rebuild properties, click the rebuild name link to open the properties dialog box.

- Locality — Where the rebuild is taking place. The rebuild can be distributed between the storage on two clusters or it can take place at one of the clusters in a metro node configuration.
- Device — The device taking part in the rebuild.
- Rebuild Type — The type of rebuild. A rebuild can be a full rebuild or a logging rebuild. - A full rebuild copies the entire contents of the source to the target. - An incremental build uses a checksum differencing algorithm to transfer only those (chunks of) blocks that are different. - A comparison is used for thin devices to preserve their thinness. Comparison rebuilds do an additional read request to the target device to determine whether a write is necessary (read from source, read from target, compare data, only write if not equal). - A resync rewrites blocks that may have been affected by a director or inter-cluster link failure, guaranteeing that the mirror legs are identical. Applies only to distributed-devices.
- Rebuilt — Shows the amount of data that has been rebuilt.
- Total — Shows the total amount of data that will be rebuilt when the operation completes.
- % Done — Calculates the percentage of this device that is rebuilt/migrated.
- Throughput Shows the rate of rebuild/,migration that this operation is experiencing based on read and write rates.
- ETA — The estimated time remaining for this rebuild based on the amount of data that has been rebuilt and the throughput the rebuild operation is experiencing.

Finding a rebuild operation

Click  at the top of the screen to search for a specific rebuild job.

Sorting rebuild jobs

You can sort in the following ways:

- Click each column header to sort the data in the column, then click a second time to change the direction of the sort.

- Press and hold **Ctrl** and click a column header to create a sort order. A number appears in the column header to indicate the order of the sort.
- Click a sorted column header to remove the number and clear a sort order.
- Rearrange columns by clicking the column header, and dragging the column to the desired location.

Viewing the Rebuild Status dashboard

1. From the GUI main menu, click **Performance**.
2. Click **+** and select **Add Rebuild Dashboard**.

Virtual Volumes dashboard

The Virtual Volumes Performance dashboard

The **Virtual Volumes Performance Dashboard** shows performance metrics of the top 100 busiest volumes, ranked by total IOPS. By default, these volumes are displayed from the busiest IOPS (throughput) volumes to the least busy. You can sort the data in each column in ascending or descending order. Note that when you sort data in other columns, the data set is still the top 100 volumes in the system based on IOPS.

The Virtual Volumes list automatically refreshes every 60 seconds, displaying data from the last 60 second period. Select a volume in the list and then click **VIEW CHARTS** to display the Throughput, Bandwidth and Latency charts for the volume. See the [Virtual Volumes Throughput chart](#), [Virtual Volumes Latency chart](#) or [Virtual Volumes Bandwidth chart](#) for specific information for each chart.

You can view metrics only for the virtual volumes in the cluster you are logged into. To view metrics for virtual volumes on a second cluster, open another browser, connect to the second cluster, and then open the Virtual Volumes dashboard.

 **NOTE:** You cannot add virtual volume charts to other dashboards or add other charts to the Virtual Volumes Dashboard.

The Virtual Volumes Performance Dashboard shows the following performance statistics for each virtual volume listed:

- IOPS — Total count of read and write operations.
- Reads (KB/s) — Bandwidth for read operations.
- Writes (KB/s) — Bandwidth for write operations.
- Read Avg Latency (usec) — Average latency or response time for read operations.
- Write Avg Latency (usec) — Average latency or response time for write operations.

General guidelines

- For Latency statistics:
 - Satisfactory latency or response time depends heavily on the application's requirements.
 - It is difficult to give recommended values for front-end latency since it depends heavily on back-end latency.
 - In general, read or write latency values under 10msec are good, and greater than 100msec is usually cause for concern.
 - Different volumes will likely have different thresholds for what is acceptable. 10msec might be acceptable to one application but totally unacceptable to another.
- For Throughput and Bandwidth statistics:
 - There is recommendation for what is good or bad for IOPS and KB/s. It is typically what the application requests of the volume.
 - If values for these metrics are unsatisfactory, be aware of resource bottlenecks such as over-saturated front-end ports, or over-utilized metro node directors.
 - Identify performance-intensive applications such as nightly back-ups or data warehouse applications that might cause other latency-sensitive applications to suffer. Identify these busy volumes and adjust accordingly so they do not conflict. This includes actions such as moving them to their own front-end ports or directors (if possible), adjusting their maximum outstanding operations counts (queue depths), or staggering their start times to avoid busy times.

Corrective actions


- Unsatisfactory virtual volume performance might be the symptom of poor storage volume performance.

- For metro node Metro systems, poor virtual volume performance could be caused by poor inter-cluster WAN link performance and/or poor storage volume performance.
- Monitor other metro node performance statistics such as director CPU usage, front-end aborts, back-end aborts, storage volume latency, or WAN latency for correlations and possible causes of the poor performance.

Display the default charts for a virtual volume

1. Select the virtual volume in the list.
2. Click **View Charts** in the properties panel. By default, all charts display in a single view.

Finding a virtual volume in the list

1. Click  in the upper-right corner of the screen to display the **Search** text box.
2. In the **Search** text box, type the full or partial name of the volume and then press **Enter**. You can use the **Previous** and **Next** to move through the list of matches.

Viewing the Virtual Volumes dashboard

1. From the GUI main menu, click **Performance**.
2. Click **+** and select **Add Virtual Volumes Dashboard**.

Virtual Volume Throughput chart

The **Virtual Volume Throughput** chart provides a time-based view of the total throughput or IOPS for a virtual volume. Generally throughput, more commonly referred to as IOPS, is associated with small block I/O (512B to 16KB I/O) requests.

Guidelines

- The desired level of IOPS performance depends heavily on the host applications and their requested load. Therefore, it is not possible to provide a threshold of good or bad IOPS performance.
- Front-end performance in metro node depends heavily on the available back-end storage array performance, and in metro node Metro configurations, the WAN performance for distributed devices.
- Any running distributed rebuilds or data migrations might negatively affect available host throughput.
- Since metro node Local and Metro implement write through caching, a small amount of write latency overhead (typically <1msec) is expected. This latency may affect applications that serialize their I/O and do not take advantage of multiple outstanding operations. These types of applications may see a throughput and IOPS drop with metro node in the data path.
- In a metro node Metro environment you will incur extra WAN round-trip time on your write latency since writes need to be successfully written to both cluster's storage before the host is acknowledged. This extra latency may impact the throughput and IOPS of serialized-type applications.

Corrective actions

- Check CPU Utilization. If it is extremely busy, metro node will be limited in the amount of throughput it can provide.
- Check back-end latency. If on average the back-end latency is large, or there are large spikes, there could be a poorly performing back-end fabric or an unhealthy, unoptimized, or over-loaded storage array.
- Check front-end aborts. Their presence indicate that metro node is taking too long to respond to the host. These might indicate problems with the front-end fabric or slow SCSI reservations.
- Check back-end errors. If the metro node back-end is required to retry an operation because of an error, then this will add to the delay in completing the operation to the host.
- Check front-end operations count (queue depth). If this counter is large, this may explain larger than normal front-end latency.
- Perform a back-end fabric analysis and a performance analysis of the storage array(s) that hosts the underlying storage volume(s) for the virtual volume.
- Check for high metro node write delta time. Refer to the Corrective actions section in the [Write Latency Delta chart](#) topic.
- If you are trying to boost IOPS performance, verify the front-end average iosize, and confirm that you are sending small block I/O.

- Check for bandwidth/IOPS over-provisioned metro node front-end ports. Be sure to balance hosts and LUNs across the available directors and front-end ports presented from metro node. Check the front-end fabric for saturation or over-capacity.
- Verify that front-end FC ports, HBAs and switch ports are configured to the correct port speeds.
- Configure your host multipathing software based on metro node best practices, and ensure the installed software versions are compatible with metro node. For more information on compatibility, see the *Dell EMC Simple Support Matrix for metro node* document, available on Dell EMC Online Support and on the Solve Online.

For metro node Metro configurations

- Check the health of the inter-cluster link and maximum performance capabilities. From the GUI, check the inter-cluster WAN bandwidth. If your application throughput appears low and seems to only achieve something similar to what the WAN bandwidth reports, then you are probably limited by the WAN. Therefore:
 - Make sure you have provisioned enough inter-cluster bandwidth for the desired application workload. Verify that your WAN configuration is supported by metro node (minimum supported bandwidth, supported inter-cluster latency, compatible WAN hardware and software).
 - For Metro-FC, if the inter-cluster WAN is over a FC fabric, confirm that you have allocated enough buffer credits or that you have configured the FC WAN ports correctly on your switches. Check for buffer credit starvation, c3 discards, and CRC errors. Some vendors may require extended fabric licenses to enable WAN features.
 - Validate your WAN performance before going live in production. Create multiple test distributed devices and force them to rebuild. Observe the performance of the rebuilds.
- When troubleshooting distributed device performance, if feasible, check local device performance. Export a test LUN from your storage array to metro node, then to the host, and then run a test I/O workload.
- Check for any unexpected local or distributed rebuilds or data migrations. There will be some amount of performance impact to host application traffic that relies on the same virtual volumes and storage volumes. Tune the rebuild transfer-size setting to limit the performance impact of rebuild and migrations. Consider scheduling migrations during off-peak hours.

Changing the view

To view the throughput of a single director in your metro node system, select the director name from the **Director** drop-down.

Viewing the Virtual Volumes Throughput chart

1. From the GUI main menu, click **Performance**.
2. Click **+** and select **Add Virtual Volumes Dashboard**.

Virtual Volume Latency chart

The **Virtual Volume Latency** chart provides a time-based view of the IO Latency for a virtual volume broken down by read and write latency. Virtual volume latency is defined as the amount of time an I/O spends within metro node for a given virtual volume. The reported metro node front-end latency should match closely to the host or application reported volume latency unless there is significant added delay in the front-end fabric, HBA, multi-pathing software or host operating system software.

For metro node cache read miss operations, it includes the time spent retrieving the disk blocks from the storage array. Therefore, for non-metro node cached operations, front-end latency can perform only as fast as the back-end array. Contrast cache read misses to cache hit operations which will be fast. Read miss and hit latency is not reported separately in the read-latency metric, so it will be difficult to know the performance of each.

It is important to distinguish how front-end latency write operations behave in the following metro node configurations:

- For metro node Local write operations, it includes the time spent protecting the disk blocks to one or more local storage arrays.
- For metro node Metro write operations to distributed-devices, it includes the time spent protecting the disk blocks to the storage array at both clusters. When writing to the remote cluster, the round-trip time on the WAN links will add to the front-end latency, depending upon the network delay observed between clusters.

Guidelines

Keep the following guidelines in mind when using this chart:

- Satisfactory latency or response time is depends heavily on the application's requirements.
- It is not possible to give recommended values for front-end latency since it depends heavily on back-end latency.
- In general, read or write latency values under 10msec are good, and greater than 100msec is usually cause for concern. Different volumes may have different thresholds for what is acceptable.
- Metro node Local processing overhead on read misses and on write operations is roughly 1msec.
- Metro node Metro systems depend heavily on the performance of the inter-cluster WAN link.
- Be aware of large spikes in latency, which might correlate to front-end or back-end aborts, or to a large front-end queue depth.
- Understand what is normal latency for your environment, and then you will have a better idea of what is abnormal.

Corrective actions

- Check CPU Utilization chart. If it is extremely busy, the time for metro node to respond to I/O will increase.
- Check back-end latency. If on average the back-end latency is large, or there are large spikes, there could be a poorly performing back-end fabric or an unhealthy, non-optimized, or over-loaded storage array.
- Perform a back-end fabric analysis, and a performance analysis of the storage array(s) that hosts the underlying storage volume(s) for the virtual volume.
- Check front-end aborts. Their presence indicate that metro node is taking too long to respond to the host. These might indicate problems with the front-end fabric.
- Check back-end errors. If metro node back-end is required to retry an operation because it failed, then this will add to the delay in completing the operation to the host.
- Check front-end operations count (queue depth). If this counter is large, this may explain larger than normal front-end latency.
- Check for high metro node write delta time. Refer to the Corrective actions section in the [Write Latency Delta chart](#) topic.
- Understand the differences in average I/O Size of host requests. Large block requests (> 1MB) will take longer to complete, whereas small block requests (1KB-8KB) will be faster.

Changing the view

To view the latency of a single director in your metro node system, select the director name from the **Director** drop-down.

Virtual Volume Latency chart

1. From the GUI main menu, click **Performance**.
2. Click the **+** button, and then **Add Virtual Volumes Dashboard**.

Virtual Volume Bandwidth chart

The **Virtual Volume Bandwidth** chart provides a time-based view of the total bandwidth (or KB/s or MB/s) in reads and writes for a virtual-volume. Generally bandwidth (also referred to as KB/s or MB/s), is associated with large block I/O (64KB or greater I/O requests).

Guidelines

- The desired level of bandwidth performance depends heavily on the host applications and their requested load. Therefore, it is not possible to provide a threshold of good or bad bandwidth performance.
- Metro node front-end performance depends heavily on the available back-end storage array performance, and in metro node Metro configurations, the WAN performance for distributed devices.
- There is no absolute recommendation on what is good or bad for IOPS and KB/s. It is typically what the application requests of the volume.
- If values for these metrics are unsatisfactory, be aware of resource bottlenecks.
- Any running distributed rebuilds or data migrations might negatively affect available host bandwidth.
- Since the metro node Local and Metro implement write through caching, a small amount of write latency overhead (typically <1msec) is expected. This latency may affect applications that serialize their I/O and do not take advantage of multiple outstanding operations. These types of applications may see a throughput and IOPS drop with metro node in the data path.

- In a metro node Metro you will incur extra WAN round-trip time on your write latency since writes need to be successfully written to both cluster's storage before the host is acknowledged. This extra latency may impact the throughput and IOPS of serialized-type applications.

Corrective actions

- Check for bandwidth/IOPS over-provisioned metro node front-end ports. Be sure to balance hosts and LUNs across the available directors and front-end ports presented from metro node. Check the front-end fabric for saturation or over-capacity.
- Check CPU utilization. If unusually busy, metro node will be limited in the amount of bandwidth it can provide.
- Check back-end latency. If on average the back-end latency is large, or there are large spikes, there could be a poorly performing back-end fabric or an unhealthy, un-optimized, or over-loaded storage array. Perform a back-end fabric analysis and a performance analysis of all storage arrays in question.
- Check front-end aborts. Their presence indicate that metro node is taking too long to respond to the host. These might indicate problems with the front-end fabric or slow SCSI reservations.
- Check back-end errors. If the metro node back-end is required to retry an operation because of errors, then this will add to the delay in completing the operation to the host.
- Check front-end operations count (queue depth). If this counter is large, this may explain larger than normal front-end latency.
- Check for high metro node write delta time. Refer to the Corrective actions section in the [Write Latency Delta chart](#) topic.
- Check the front-end average iosize. For writes, iosizes larger than 128KB become serialized into 128KB requests. This can extend the time to complete large block transfers, or in extreme examples, cause the operation to timeout and fail.
- Verify that front-end Fibre Channel ports, HBAs and switch ports are configured to the correct port speeds.
- Configure your host multipathing software based on metro node best practices, and ensure the installed software versions are compatible with metro node. For more information on compatibility, refer to the *Simple Support Matrix for metro node* document, available on EMC Online Support and in the SolVe Desktop.

For metro node Metro configurations

- Check the health of the inter-cluster link and maximum performance capabilities. From the GUI, check the inter-cluster WAN bandwidth. If your application throughput appears low and is similar to what the WAN bandwidth reports, then you are probably limited by the WAN. In this case:
 - Make sure you have provisioned enough inter-cluster bandwidth for the desired application workload. Verify that your WAN configuration is supported by metro node (minimum supported bandwidth, supported inter-cluster latency, compatible WAN hardware and software).
 - For Metro-FC, if the inter-cluster WAN is over a FC fabric, confirm that you have allocated enough buffer credits or configured the FC WAN ports properly on your switches. Check for buffer credit starvation, c3 discards, and CRC errors. Some vendors may require extended fabric licenses to enable WAN features.
 - Validate your WAN performance before going live in production. Create multiple test distributed devices and force them to rebuild. Observe the performance of the rebuilds.
- When troubleshooting distributed device performance, if feasible, check local device performance. Export a test LUN from your storage array to metro node, then to the host, and run a test I/O workload.
- Check for any unexpected local or distributed rebuilds or data migrations. There will be some performance impact to host application traffic that relies on the same virtual volumes and storage volumes. Tune the rebuild transfer-size setting to limit the performance impact of rebuild and migrations. Consider scheduling migrations during off-peak hours.

Changing the view

To view the bandwidth of a single director in your metro node system, select the director name from the **Director** drop-down.

Viewing the Virtual Volume Bandwidth chart

1. From the GUI main menu, click **Performance**.
2. Click **+** and select **Add Virtual Volumes Dashboard**.

Front End Ports dashboard


The Front End Ports dashboard

The **Front End Ports** dashboard shows performance metrics for all metro node front-end ports, sorted by the busiest total IOPS ports to the least busiest ports. You can sort each column in ascending or descending order.

Note that historical data is not available for ports. The dashboard automatically refreshes every five seconds, displaying data from the last five second period.

For each front-end port, the dashboard shows the following performance statistics:

- Director—Displays data for all directors or a specific director in the cluster.
- Port—Displays data for all ports or for a specific port.
- IOPs (Total) — Total read and write operations per second.
- Read (IOPS) — Read operations per second.
- Write (IOPS) — Write operations per second.
- Queue Depth (count) — Number of outstanding operations.
- Reads (KB/s) — Bandwidth for read operations.
- Writes (KB/s) — Bandwidth for write operations.
- Read Avg Latency (usec) — Average latency or response time for read operations.
- Write Avg Latency (usec) — Average latency or response time for write operations.

 **NOTE:** Metro node back-end ports metrics are not currently available.

Guidelines


- It is best practice to evenly balance all front-end ports whenever possible.
- For Queue Depth:
 - Queue Depth is sometimes referred to as number of outstanding operations. Do not confuse this with operations per second (IOPS).
 - Excessively high queue depth values on a port tends to result in high (poor) host response time.
 - Maximum per port queue depth is around 1100 operations.
 - Front-end ports with consistently high per-port queue depths with high latency should be addressed. Spread the workload across additional front-end ports.
 - Per port queue depths depends heavily upon application workloads and HBA adapter settings.
 - Follow the Dell EMC recommended values for host and HBA queue depths. Refer to the *Host Connectivity Guides*, available on Dell EMC Online Support.
- For Latency statistics:
 - Satisfactory latency or response time depends heavily on the application's requirements.
 - It is difficult to give absolute recommended values for front-end port latency since it depends largely on back-end latency.
 - Read or write latency values under 10msec are good, and greater than 100msec is usually cause for concern.
 - High latency ports should have host workload moved or shared across other lower latency ports, if the port's workload is the problem instead of another underlying problem.
- For throughput (IOPS) and Bandwidth (KB/s) statistics:
 - Monitor per port MB/s maximum usage to avoid bandwidth saturation (for example, one 8Gbps FC port's maximum available bandwidth is typically 800 MB/s or less.)
 - If values for these metrics are unsatisfactory, be aware of resource bottlenecks such as over-saturated front-end ports or host initiators, or over-utilized metro node directors.
 - Identify performance-intensive applications such as nightly backups, or data warehouse applications that might cause other latency-sensitive applications to suffer.
 - Check the **Virtual Volumes** dashboard to identify busy volumes.

Corrective actions

- Whenever possible, use all available front-end ports.

- Unsatisfactory front-end port performance might be a symptom of poor back-end performance (metro node to storage array data path.)
- For metro node Metro systems, poor front-end port performance could result from either poor inter-cluster WAN link performance, and/or poor remote cluster storage volume performance.
- Monitor other metro node performance statistics such as director CPU usage, front-end aborts, back-end aborts, storage volume latency, or WAN latency for correlations and possible causes of the poor performance.

Finding a director or port in the list

1. Click  in the upper-right corner of the screen to open the **Search** text box.
2. In the **Search** text box, type the name of the director or port, and then press **Enter**. The matching item is highlighted. You can click **Previous** and **Next** to move through the list of matches if applicable.

Viewing the Front End Ports dashboard


1. From the GUI main menu, click **Performance**.
2. Click **+** and then select **Add Front End Ports Dashboard**.

System Health

The **System Health** dashboard provides a visual representation of the results of the CLI `health-check` command. The `health-check` command scans the major components of the system and outputs any errors and warnings. The GUI displays the output of this command in various charts using colored status indicators that allow you to quickly view the cluster health, component health, and health of other components such as metro node Witness status (if installed). Use this dashboard to validate the system health before or after any significant operations such as non-disruptive upgrades or reconfiguration of the system.

The **System Health** dashboard shows the following charts:

- **Cluster Health** — Displays the health status and other information for the cluster. The Transition/Health Indications section at the bottom of the cluster health chart lists general errors and warnings for the cluster.
- **Component Health** — Shows the status of the storage objects and other components in each cluster.
- **Metro node Product Information** — Shows the type of metro node configuration, hardware, WAN connectivity (if applicable), and version. Click the **View Version Details** link under the product information to view more detailed software version information about the metro node system components.
- **Metro node Witness** (if applicable)—Shows the status and other details about metro node Witness. To see additional metro node Witness details, click the **View Status Details** link at the bottom of the metro node Witness chart.

The last refresh time displays at the upper-right corner of the screen, however, you can manually refresh the data by clicking the Refresh icon  in the upper-right corner of the screen.

Colored status indicators

Each chart uses color in the status indicators to indicate the status of the corresponding components. The number on the status indicator represents the number of components in that state. For example, if a cluster has four directors; two may report a warning status (yellow) and two may be functioning (green).

Click anywhere in the status indicator to go to the corresponding component view to manage the component. For the **Directors** and **Metadata** status indicators, additional information displays in a separate popup window when you click the status indicator. In the popup window, click a director to view its properties. The **System Health** dashboard uses the following colors to indicate the health of a component:

- Green—OK, the component is functioning.
- Yellow—Warning, the component is changing state or needs user attention.
- Red—Error, an error has occurred that requires user action.
- Grey—Unknown status, metro node cannot determine the status of the component.

Corrective actions

For information about corrective actions for any errors or warnings reported, see the Troubleshooting section of metro node in the SolVe Desktop, available on Dell EMC Online Support.

Opening the System Health dashboard

From the GUI main menu, select **Dashboard > System Health**.

Provisioning storage

This section describes how to work with various storage objects to provision storage.

Topics:

- [Guide](#)
- [Provisioning from storage volumes](#)
- [Provision Job properties](#)
- [Distributed storage](#)
- [Storage arrays](#)
- [Storage volumes](#)
- [Devices](#)
- [Extents](#)
- [Distributed devices](#)
- [Virtual volumes](#)
- [Consistency group](#)

Guide

For the steps to begin provisioning storage, see [Guide](#). In the Guide, each step is a link that opens the UI screen that is required to perform the step. To view the Guide, click **Provision Storage**, and then expand the **GUIDE** at the upper-right corner of the screen. For more information about a step, click **MORE** for that step to open the associated help topic.

Distributed storage

Distributed storage sees storage objects that are created by using storage from both clusters. To view distributed storage, click **Provision storage**, and then select **Distributed Storage** from the **Provision on** drop-down, and then select a required view from the **View By** drop-down.

Provisioning from storage volumes

Provision from Storage Volumes wizard allows you to provision a virtual volume directly from a storage volume or preserve data on an existing storage volume that you want to expose to hosts. The wizard simplifies the provisioning process by automatically claiming storage (if not already claimed) and creating all of the underlying storage objects (extents and devices), and then creating a local or distributed virtual volume that is the total capacity of the selected storage volume. When provisioning from storage volumes, you can create only one virtual volume at a time. Each virtual volume that you create maps to a storage volume on the array. You can provision from storage volumes using integrated or non-integrated storage arrays. [Provisioning Overview](#) provides more information about integrated and non-integrated storage.

Creating a new virtual volume is as simple as selecting or creating a consistency group for the volume, selecting mirroring options (optional), and then selecting an available storage volume from an array on the cluster on which you want to create the virtual volume. You do not need to claim the storage volume before you begin to provision, and you can select claimed or unclaimed storage volumes from the selected array.

The **Provision from Storage Volumes** wizard consists of the following steps:

1. Select an existing consistency group for the volume or create a consistency group.
2. Select mirroring options (optional) and synchronization options (if applicable) and provide a name for the virtual volume.
3. Select a storage volume to use to create the virtual volume.
4. Expose the virtual volume to hosts by selecting a storage view (optional).
5. Review your selections, and submit the provisioning request.
6. View the results of the operation.


Creating thin virtual volumes

To create a thin virtual volume, the storage volume must be thin capable. If you do not select a thin capable storage volume, the resulting volume is thick.

Launching the wizard

You can launch the Provision from Storage Volumes wizard from several different locations that are listed below. Depending on where you launch the wizard, you are given the option to create a specific type of consistency group by default.

- **Provision Storage** menu (Local or distributed)
- **Distributed Virtual Volumes** view (Distributed only)
- **Virtual Volumes** view on a specific cluster (Local only)
- **Storage Volumes** view on a specific cluster (Local only)

Complete each step of the wizard. For additional information about any of the steps, click  in the upper-right corner of the screen.

Provision Job properties

Provision Job Name	The optional name given to the provisioning job. If you did not provide a name for the job, the base volume name is used.
Target Cluster	The cluster on which the job was initiated.
Status	The status of the job.
Status Details	An explanation of the job status.
Created	The date and time the job was created.
Last Updated	The last date and time the job status was automatically updated.
Description	A description of the volumes being created.
Volume Options	The options selected for the volumes.
Consistency Group	The name of the consistency group for the volumes (if any).
Total Capacity	The total capacity of the volumes being created.
Thin	Indicates if the volumes are thin.
Resources affected at <cluster>	Storage View — The storage view selected for the volumes (if any). Storage Array — The storage array from which the volumes were provisioned. Storage Pool — The storage pool from which the volumes were created. Storage Group — The storage group automatically assigned or selected for the virtual volumes (if applicable).

Distributed storage

Distributed storage refers to storage objects (consistency groups, virtual volumes, devices) that are created by using storage from both clusters. Select the **Distributed Storage** from the **Provision on** drop-down in **Provision Storage** view to conveniently access the distributed storage objects. From this section you can view a grouping of all distributed consistency groups, distributed virtual volumes, and distributed devices.

Storage arrays

Supported storage arrays



Metro node supports a variety of storage systems and third-party arrays. For an up-to-date list of supported arrays and other components, see the *Simple Support Matrix for metro node* document, available on Dell EMC Online Support and on the SolVe Online.

Using the Storage Arrays view

Metro node automatically discovers storage arrays that are connected to the back-end ports. All arrays connected to each director in the cluster are listed in the **Storage Arrays** view. In some cases you may need to [rediscover](#) the storage array, if for example:

- The array was not discovered correctly (you do not see the expected number of logical units on the array when you [view logical units](#) for the array).
- You made changes to the storage exported in a previously discovered array.

Arrays are sorted by name, in ascending alphabetical order. The arrow in the column indicates the direction of the sort;

ascending () or descending (). You can [sort](#) columns of data, [select](#) one or more arrays in the list, and [show or hide properties](#) on the screen.

This screen shows the following information about storage arrays:



Column	Description
Storage Array	The name of the discovered storage array.
Connectivity Status	Indicates if the directors have at least one path to both of the array's controllers. Storage array connectivity status provides more information.
Logical Unit Count	The number of LUNs exported from the array.
Auto Failover	Indicates whether metro node will trespass the LUNs after a controller failure. This attribute is reported from the array, and displays only for active/passive and explicit ALUA arrays.

Storage Arrays Properties Panel

This screen also displays the properties of the storage arrays in a separate panel at the extreme right side of the screen.

Tasks in this Screen

In the **Storage Arrays** view, you can also do the following:

- Search for an extent — Click the  icon corresponding to the element and type the full or the partial name in the field that appears. Press **Enter**.
- Refresh screen data — Click the Refresh icon .
- Create virtual volumes — Select a storage array, and then click **CREATE VIRTUAL VOLUMES** to start the Create Virtual Volumes wizard.
- Claim storage — Select a storage array, and then click **CLAIM STORAGE** to start the Claim Storage wizard.
- Set thin rebuilds — Select a storage array, select **MORE > Set Thin Rebuild**.
- Show logical units for an array — Select a storage array and select **MORE > Show Logical Units**. If you do not see the expected number of LUNs, [rediscover](#) the array.
- Rediscover an array — Select a storage array, and select **MORE > Rediscover Array** to rediscover the array that might not be displayed correctly.
- Forget storage — Select an array, and select **MORE > Show Logical Units**. On the **Show Logical Units** screen, select the LUN, and click **Forget Unreachable LUNs**.

Opening the Storage Arrays view

From the GUI main menu, do one of the following:

1. In a metro node Local:
 - a. Select **Provision Storage**.
 - b. In the **View by** drop-down, select **Storage Arrays**.
2. In a metro node Metro:
 - a. Select **Provision Storage**.
 - b. In the **Provision on** drop-down, select the cluster of the storage array.
 - c. In the **View by** drop-down, select **Storage Arrays**.

Rediscovering a storage array

About this task

Rediscover a storage array if any changes you made to the array are not reflected in the GUI. Specifically, you may need to perform this operation if you:

- Add a new array and the number of expected storage volumes do not display when you view the logical units for the array
- Modify an existing array, for example, export a new LUN and it is not visible in the GUI

Depending on how many storage volumes are in the storage array, it can take several minutes to complete.

Steps

1. From the GUI main menu, do one of the following:
 - a. In a metro node Local:
 - i. Select **Provision Storage**.
 - ii. In the **View by** drop-down, select **Storage Arrays**.
 - b. In a metro node Metro:
 - i. Select **Provision Storage**.
 - ii. In the **Provision on** drop-down, select the cluster of the storage array.
 - iii. In the **View by** drop-down, select **Storage Arrays**.
2. In the **Storage Arrays** view, select an array, and then click **MORE > Rediscover Array**.
A message appears to confirm the operation.
3. To forget any unreachable LUNs while rediscovering this array, click **Forget unreachable LUNs**.
4. Click **Yes**.
5. Show logical units for the array again to make sure that all expected storage volumes are visible.

Viewing logical units for an array

Steps

1. From the GUI main menu, do one of the following:
 - a. In a metro node Local:
 - i. Select **Provision Storage**.
 - ii. In the **View by** drop-down, select **Storage Arrays**.
 - b. In a metro node Metro:
 - i. Select **Provision Storage**.
 - ii. In the **Provision on** drop-down, select the cluster of the storage array.
 - iii. In the **View by** drop-down, select **Storage Arrays**.
2. In the **Storage Arrays** view, select an array, and then select **MORE > Show Logical Units**.
The **Show Logical Units** screen lists all LUNs discovered on the array. The total number of LUNs appear in the upper-left corner of the screen. If you do not see the expected number of LUNs, [rediscover](#) or [forget](#) any unreachable arrays.
 - To view the properties of a LUN, click the LUN name link to open the properties dialog box.
 - The **ALUA Support** column indicates if ALUA is supported on the array. If ALUA is not supported on the array, **None** displays for the ALUA-specific properties. For more information about ALUA, see the array vendor's documentation.

3. Click **Close**.

Forgetting a storage array

About this task

Forget a storage array if it has been deleted from use by metro node and it is no longer used for I/O. You can forget a LUN from the **Show Logical Units** screen, or from the **Rediscover Storage Array** dialog box while rediscovering a storage array.

Steps

1. From the GUI main menu, do one of the following:
 - a. In a metro node Local:
 - i. Select **Provision Storage**.
 - ii. In the **View by** drop-down, select **Storage Arrays**.
 - b. In a metro node Metro:
 - i. Select **Provision Storage**.
 - ii. In the **Provision on** drop-down, select the cluster of the storage array.
 - iii. In the **View by** drop-down, select **Storage Arrays**.
2. In the **Storage Arrays** view, select a storage array, and then select **MORE>Show Logical Units**.
3. In the **Show Logical Units** screen, click **Forget Unreachable LUNs**.
A confirmation dialog box is displayed.
4. Click **Yes**.
5. View the results of the operation, and then click **Close**.

Storage array connectivity status

The following table describes the array connectivity states.

Connectivity state	Definition
OK	The array is functioning normally; both directors can see each LUN exported from the array.
Unknown	Metro node is unable to report on the connectivity status of the array.
Error	One or more directors cannot see a LUN exported from the array.
Degraded	A LUN exported from the array does not have two paths to each director.

Storage volumes

About storage volumes

A storage volume is a LUN exported from an array. When metro node discovers an array, the **Storage Volumes** view shows all LUNs exported from the array. You must claim, and optionally name these storage volumes before you can use them in a metro node cluster. Once claimed, you can divide a storage volume into multiple extents (up to 128), or create a single extent that uses the entire capacity of the storage volume.

To claim storage volumes, use the Claim Storage wizard, which assigns a meaningful name to the storage volume. Meaningful names help to associate a storage volume with a specific storage array and LUN on that array, and helps during troubleshooting and performance analysis.

For thinly-provisioned storage volumes, you can request a thin rebuild for local or distributed mirrors. [Setting a storage volume rebuild type](#) provides more information.

The Storage Volumes view

The **Storage Volumes** view lists all storage volumes in the selected cluster. By default, storage volumes are sorted in ascending alphabetical order, then by name. The arrow in the default sort column indicates the direction of the sort; ascending (↑) or descending (↓). You can sort columns of data, select one or more storage volumes in the list, and show or hide properties links on the screen.

In the main panel of the screen, you can perform these tasks:

- Create virtual volumes
- Provision virtual volumes
- Set the rebuild type
- Claim all the storage on the supporting array
- Unclaim storage
- Create extents
- Show ITLs

This screen shows the following information about storage volumes:

Column	Description
Storage Volume	The name of the storage volume.
Capacity	The size of the storage volume.
Health	The overall health of the storage volume. Storage volume status provides more information.
Status	Indicates how the storage volume is functioning in the cluster. Storage volume status provides more information.
Use	Indicates whether the storage volume is Unclaimed, Claimed, Used, Unusable, or contains metadata (used as a meta-volume).
Provision Type	Indicates how the storage volume was provisioned. Possible value: <ul style="list-style-type: none">• Legacy — The storage volume was provisioned from non-integrated storage arrays.
Thin Rebuild	Indicates if a thin rebuild will be used for the storage volume.
Thin Capable	Indicates if the storage volume is thin-capable.
Vendor	The name of the supporting array's manufacturer.
Largest Free Chunk	The total number of bytes in the largest range of blocks available on the storage volume. (Displays when selected in the View filter in the upper-right corner of the screen. Viewing the largest free chunk on a storage volume provides more information.)
Storage Array	The storage array on which this storage volume was created. To view the storage array properties, click the storage array name link.



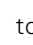

Storage Volume Properties Panel

This screen also displays the properties of the storage volumes in a separate panel at the extreme right side of the screen. Using the **Storage Volume Properties** panel, you can:

- View the details of the storage volume
- Turn on thin rebuild on the storage volume
- View the storage volume map
- Rename the storage volume
- View the details of the storage array from where the storage volume has been created
- Rename the storage array

Tasks in this Screen

You can also perform the following tasks in the **Storage Volumes** screen:

- Search for a storage volume — Click the  icon corresponding to the element and type the full or the partial name in the field that appears. Press **Enter**.
- Refresh screen data — Click the Refresh icon .
- Provision from storage volumes — Click **PROVISION** to open the Provision from Storage Volumes wizard.
- Create local or distributed virtual volumes — Click the **CREATE VIRTUAL VOLUMES** drop-down and select **Local** or **Distributed** to start the Create Virtual Volumes from Claimed Storage wizard.
- Show ITLs for a selected storage volume — Select the storage volume, and then click **MORE > Show ITLs**.
- Set rebuild type for a selected storage volume — Select a storage volume, and then click **MORE > Set Rebuild Type**.
- Claim all storage from a supporting array — Select a storage volume, and then click **MORE > Claim All Storage on Supporting Array** to open the Claim Storage wizard.
- Unclaim selected storage volumes — Select a storage volume, and then click **MORE > Unclaim Storage**.
- Create extents — Select the storage volume, and then click **MORE > Create Extents** to open the Create Extents wizard.
- Turn on Thin Rebuild — Select the storage volume, and then click **TURN ON THIN REBUILD** in the **Storage Volume Properties** panel.
- View Map — Select the storage volume, and then click **VIEW MAP** in the **Storage Volume Properties** panel.
- Rename a storage volume — Select the storage volume, and then click **Rename** under the **STORAGE VOLUME** tab in the **Storage Volume Properties** panel. Click  to save the changes.
- Rename a storage array — Select the storage volume, and then click **Rename** under the **STORAGE ARRAY** tab in the **Storage Volume Properties** panel. Click  to save the changes.

Opening the Storage Volumes view

From the GUI main menu, do one of the following:

1. In a metro node Local:
 - a. Click **Provision Storage**.
 - b. In the **View by** drop-down, click **Storage Volumes**.
2. In a metro node Metro:
 - a. Click **Provision Storage**.
 - b. In the **Provision on** drop-down, select the cluster of the storage volume.
 - c. In the **View by** drop-down, select **Storage Volumes**.

Claiming storage volumes

About this task

Use the Claim Storage wizard to claim storage. You must claim the storage volumes before you can use them in a cluster (except for the metadata volume, which is created from an unclaimed storage volume). Only after claiming a storage volume, can you use it to create extents, devices, and then virtual volumes. Before you claim storage volumes, view the discovered arrays to ensure that they were discovered correctly, and you have the expected number of unclaimed storage volumes. To encapsulate storage volumes already in use, see the Encapsulation procedures in the SolVe Desktop, available on Dell EMC Online Support.

You can claim storage in the following ways:

- Claim and name all storage volumes from an array.
- Claim and name one or more storage volumes at a time.
- Claim storage volumes from the same array using different storage types (an identifier that indicates the category or type of data that are stored on the storage volume).

Some storage arrays require a Name-Mapping file to claim storage. When you select an array to claim storage, metro node automatically detects whether the storage volumes require a Name-Mapping file or not. If you select an array that requires a name-mapping file, the wizard informs you, and allows you to upload the file. If you know that the array requires a name-mapping file, create the file before you begin to claim storage.

If the array does not require a Name-Mapping file, you have the option of creating and using this file to customize the names of the claimed storage volumes. For instructions to create a Name Mapping file and other configuration information, see the Configure storage arrays for metro node procedure in the SolVe Desktop, available on Dell EMC Online Support. To provision 3PAR LUNs for metro node, see the 3PAR documentation.

Steps

1. From the GUI main menu, do one of the following:
 - In metro node local
 - a. Click **Provision Storage**.
 - b. In the **View by** drop-down, click **Storage Volumes**.
 - In metro node Metro
 - a. Click **Provision Storage**.
 - b. In the **Provision on** drop-down, select the cluster of the storage volume.
 - c. In the **View by** drop-down, select **Storage Volumes**.
2. In the **Storage Arrays** view, click **Claim Storage**.
The Claim Storage wizard is displayed.
3. Follow the instructions on the wizard screens to proceed. For additional help using the wizard, click **Help**.

The Claim Storage wizard

Use the Claim Storage wizard to claim storage an array. To encapsulate storage volumes already in use, see the Encapsulation procedures in the SolVe Desktop, available on Dell EMC Online Support.

Naming storage volumes

By default, storage volumes are identified by their SCSI Vital Product Data (VPD), for example: VPD83T3:600601604a950d00fe9bc44d11b3da11. If the array requires a Name Mapping file, use this file to map the default storage volume names to more user-friendly names that allow you to easily identify the storage volume.

For arrays that do not require a Name Mapping file, the storage volumes are named using one of the following formats:

- `<array_name><last_four_digits_of_array_serial_number>_<device_number>`
- `<array_name><last_four_digits_of_array_serial_number>_<LUN_number>`

You can rename the storage volumes after claiming.

Selecting a storage type

Select a storage type to assign different categories of data to different types of storage media as a means of reducing total storage cost. Categories of data can be based on levels of protection needed, performance requirements, frequency of use, and other considerations. As an example, you could designate all storage designated as "High" to be used for highly classified or critical information, which would be stored on more expensive, high quality storage arrays. Data designated as "Medium" might be seldom-used data, which could be stored on a less expensive storage array, and so forth. By default, the storage type is included in the storage volume name.

Selecting a thin-rebuild


For thinly-provisioned storage volumes, you can select if the storage volume will use thin rebuilds. [Setting a thin rebuild for storage volumes](#) provides more information.

Setting a storage volume rebuild type

About this task

For XtremIO, VNX and VMAX arrays, metro node automatically detects if a storage volume is thin capable. For all other arrays, metro node does not automatically detect if a storage volume is thinly-provisioned on the array. Therefore, if a storage volume is thinly-provisioned, or thin capable, choose a thin rebuild. When you select a thin rebuild, metro node preserves the thinness of the target device by copying only non-zero data.

You can specify the type of rebuild for a storage volume when you claim the storage (from the **Storage Arrays** view), or you can do so later by selecting claimed storage volumes (**Storage Volumes** view). Note that the claimed storage volumes may be in use. The change takes effect immediately.

 **NOTE:** For thin capable storage volumes, metro node will automatically select a thin rebuild even if you did not select this option in the wizard during the claiming process.


Steps

1. From the GUI main menu, do one of the following:
 - In a metro node Local:
 - a. Select **Provision Storage**.
 - b. In the **View by** drop-down, select **Storage Volumes**.
 - In a metro node Metro:
 - a. Select **Provision Storage**.
 - b. In the **Provision on** drop-down, select the cluster of the storage volume.
 - c. In the **View by** drop-down, select **Storage Volumes**.
2. Select the storage volume, and then select **MORE > Set Rebuild Type**.
The **Set Claimed Volumes Rebuild Type** dialog box opens.
3. Select **Yes** under **Use Thin Rebuilds**.
4. Click **OK**.

Unclaiming a storage volume

About this task

Storage volumes are either discovered on the array, or created from storage pools. Regardless of how a storage volume was created, when you unclaim it, the status changes to Unclaimed, making it available for claiming again. However, if the storage volume was created from a storage pool, you have the option of deleting the storage volume from the array and returning the allocated storage capacity to the pool on the array. You can select multiple storage volumes to unclaim, and the selection can contain a mix of storage volumes created with and without integrated services.

 **NOTE:** You cannot unclaim a storage volume that is currently in use.

Steps

1. From the GUI main menu, do one of the following:
 - In a metro node Local:
 - a. Select **Provision Storage**.
 - b. In the **View by** drop-down, select **Storage Volumes**.
 - In a metro node Metro:
 - a. Select **Provision Storage**.
 - b. In the **Provision on** drop-down, select the cluster of the storage volume.
 - c. In the **View by** drop-down, select **Storage Volumes**.
2. Select one or more claimed storage volumes, and then select **MORE > Unclaim Storage**.
A confirmation dialog box is displayed.
 - If the confirmation box displays a check box, this means that your selection contains at least one storage volume created from a storage pool. Select the check box to delete the storage volume on the array and return the storage capacity of those storage volumes to the pool from which they were created.
 - If the confirmation box does not contain a check box, click **Yes**.

3. Review the results of the operation, and then click **Close**.


Viewing ITLs for a storage volume

About this task

The **Show ITLs** dialog box shows the selected storage volume's back-end path (metro node back-end port, storage target port, and LUN). For each storage volume there should be at least two initiator-target-LUN (ITLs) for each director, indicating that each director has two paths to the storage volume. The two paths should include two back-end ports on each director, each connected and zoned to a separate port on the array. This ensures that there is high availability and no single point of failure.

In this dialog box, you can also search for a specific initiator, target or LUN, and [sort](#) the columns of data. This information displayed is read-only.

Steps


1. From the GUI main menu, do one of the following:
 - In a metro node Local:
 - a. Select **Provision Storage**.
 - b. In the **View by** drop-down, select **Storage Volumes**.
 - In a metro node Metro:
 - a. Select **Provision Storage**.
 - b. In the **Provision on** drop-down, select the cluster of the storage volume.
 - c. In the **View by** drop-down, select **Storage Volumes**.
2. Select the storage volume, and then select **MORE > Show ITLs**.
The **Show ITLs** dialog box is displayed and lists all ITLs for the storage volume.
3. To search an initiator, target or LUN, click the  icon corresponding to the element and type the name in the field that appears.
4. Click **Close**.

Renaming a storage volume

About this task

Rename a storage volume at any time in the **Storage Volume** properties dialog box.

Steps

1. From the GUI main menu, do one of the following:
 - a. Click **Provision Storage**.
 - b. In the **View by** drop-down, click **Storage Volumes**.
2. In the **Storage Volume** view, click the storage volume name link.
3. In the properties panel, click the **Rename** under **STORAGE VOLUME** tab , and the type a new name for the storage volume.
4. Click .

Storage Volume status

The following tables define a store volumes Health, Operational and I/O states.

Health state	Definition
OK	The storage volume is functioning normally.

Health state	Definition
Degraded	The storage volume may be out-of-date and/or the underlying disk is not performing optimally.
Unknown	Metro node cannot determine the storage volume's Health state, or the state is invalid.
Non-recoverable error	The storage volume may be out-of-date compared to its mirror (applies only to a storage volume that is part of a RAID-1 Metadata Volume), and/or metro node cannot determine the Health state.
Critical failure	Metro node has marked the storage volume as hardware-dead.
Isolated	The storage volume is not performing optimally, which caused metro node to prevent IO to the mirror.

Operational state	Definition
OK	The storage volume is functioning normally.
Degraded	The storage volume may be out-of-date and/or the underlying disk is not performing optimally.
Unknown	Metro node cannot determine the storage volume's Operational state, or the state is invalid.
Error	Metro node has marked the storage volume as hardware-dead.
Starting	The storage volume is not yet ready.
Lost communication	The storage volume is unreachable.

I/O state	Definition
Alive	I/O is proceeding normally to the storage volume.
Dead	Metro node has marked the storage volume as dead; I/O cannot proceed on the storage volume. This can happen when a certain number of I/Os to the storage volume fails.
Unreachable	Metro node has lost communication with the storage volume.

Devices

About devices

Devices combine extents or other devices into one large device with specific RAID techniques such as mirroring or striping. Devices can only be created from extents or other devices. A device's storage capacity is not available until you create a virtual volume on the device and export that virtual volume to a host. You can create only one virtual volume per device. Creating a virtual volume on a device indicates the device configuration is complete.

Note that devices use storage from one cluster only. To create a device from storage in both clusters in a metro node Metro, refer to [Creating distributed devices](#).

There are two ways to create devices:

- Simple — a device configured by using one component (an extent).
- Complex — a device that has more than one component, combined by using a specific RAID type. The components can be extents or other devices (both simple and complex).

To create a device using different device types, use the CLI. The *CLI Guide for metro node* provides more information.



Figure 1. Simple device configuration




Figure 2. Complex device configuration

Using the Devices view

The **Devices** view shows all local devices in the selected cluster. Use this view to quickly see the status of a device, search for a device, or view the underlying components of a device. Devices are sorted by name in ascending alphabetical order. You can [sort](#) columns of data, [select](#) one or more objects in the list, and [show/hide properties](#) on the screen.



This screen shows the following information about devices. To view additional device properties, click the device name link.

Column	Description
Device	The name of the top-level device.
Geometry	The underlying RAID structure of the device.
Capacity	The size of the device.
Health	The overall health of the device. Device status provides more information.

Column	Description
Status	Indicates how the device is functioning in the cluster. Device status provides more information.
Virtual Volume	The name of the virtual volume built on the device.

Additional tasks

You can also perform the following tasks from the **Devices** view:

- View the components of a device — Select the device. The underlying components display in the components panel.
- View the properties of an underlying extent — Click the extent name under the Components tab in the panel on the right.
- Create a device — Click **CREATE** to open the Create Devices wizard.
- View a device map — Select the device, and click the **VIEW MAP** under the **DEVICE** tab in the Device Properties panel.
- Create virtual volumes on top of a device — Select a device, and then select **MORE > Create Virtual Volumes**.
- Add a local mirror — Click **Add Local Mirror** to open the Add Local Mirrors wizard.
- Add a remote mirror — Click **Add Remote Mirror** to open the Add Remote Mirror wizard.
- Search for a device — Type the device name in the **Find** text box at the top of the screen. If this box is not visible, click  in the upper-right corner of the screen.
- Delete a device — Select the device, and then click **Delete**.
- Refresh screen data — Click  in the upper-right corner of the screen.

Opening the Devices view

From the GUI main menu, do one of the following:

1. In a metro node Local:
 - a. Select **Provision Storage**.
 - b. In the **View by** drop-down, select **Devices**.
2. In a metro node Metro:
 - a. Select **Provision Storage**.
 - b. In the **Provision on** drop-down, select the cluster of the devices.
 - c. In the **View by** drop-down, select **Devices**.

The Create Devices wizard

Use the Create Devices wizard to create the following types of devices:

- RAID-1

Additionally, you can use the **1:1 Mapping of extents to devices** option in the wizard to create a device for each selected extent. While creating these devices, you have the option of creating a virtual volume on the device. Supported RAID types provides more information about the RAID types. For each type of RAID, you can select a maximum of 64 extents.

To create a device using different device types, use the CLI. The *CLI Guide for metro node* provides more information.

RAID-1 devices

When creating a RAID-1 device, you must select a source and target extent in order to synchronize the data between extents. RAID-1 creates a device that mirrors I/O from the source to the target extent. For RAID-1 devices, you must select extents from different storage volumes. Once you select one extent from a storage volume, all other extents from that storage volume are removed from your selection options. You can select any number of source and target extents to create a RAID-1 device. When you view the device map, the extents display as children of the device.

1:1 Mapping of Extents to Devices

One to one mapping creates a device for each selected extent. The device is the total capacity of each selected extent. Note that 1:1 mapping creates a RAID-0 device with stripe depth of 1.

Virtual volume names

If a virtual volume was created on the device, the virtual volume is given the device name, followed by `_vol`. For example, if the device name is "Mydevice1", the virtual volume created on the device will be named `Mydevice1_vol`. You can rename the device or virtual volume at any time in the properties dialog box.

Device names

To avoid entering a different name for each device you are creating, you can enter a base name to be used for all the devices. Each device is given this name followed by an underscore and an incremented number. For example:


- `base name_1`
- `base name_2`

You can also change the name of individual devices if desired.

The Add Local/Remote Mirror wizards

You can mirror a device locally or remotely using the Add Local Mirror wizard or the Add Remote Mirror wizard. When you add a local mirror, a device in the local cluster is used to create a RAID-1 device. If the local device you are mirroring is already a RAID-1, the mirror is added as another leg.

For a remote mirror, a device in the remote cluster is used to create a distributed device. You can mirror a device with or without a volume on top. When creating a remote mirror, the mirror must not be in a consistency group.

 **NOTE:** For a mirror virtual volume to be thin-enabled, the supporting devices must be thin-capable.


Synchronizing data

For a remote mirror, you can indicate whether or not to perform an initial data synchronization between the devices if the device does not have a virtual volume. Note that this is an initial synchronization. If you choose not to perform an initial synchronization, future writes are still be mirrored. If the device has a virtual volume, the data is automatically synchronized. If the device contains data you want to preserve, you should synchronize. Note that any data on the target device will be overwritten.

Selecting a consistency group or detach rule for distributed devices

Because a distributed device is created when you mirror a device remotely, you must select a consistency group or a detach rule for the distributed device. If you add the distributed device to a consistency group, the rule set used for the group will be used for the distributed device.

You have the option of adding the volumes to a consistency group or selecting a rule set for the distributed devices. If you do not add the distributed device to a consistency group, you must select from the predefined detach rules, or you can use the CLI to create a detach rule which will automatically appear in the **Rule Set** drop-down selection box. For more information on creating and managing detach rules, refer to the *CLI Guide for metro node*.

 **NOTE:** You can only apply user-created rule sets to distributed devices that are not in a consistency group.

Selecting target devices to use as mirrors

When creating a mirror, you can select target devices automatically or manually. If you click **AUTO-GENERATE MIRROR MAPPINGS** in either wizard, metro node selects a target device of the same size. If no devices are available, the appropriate message displays. While you can manually select a target device of larger capacity, this is not recommended because you cannot use the extra capacity. You can also save mirror mappings to a file on your system.

Viewing the status of IO to a device

About this task

The Service status attribute in the **Device Properties** dialog box indicates the status of I/O to a device. [Device status](#) provides a definition of the service status.

Steps

1. From the GUI main menu, do one of the following:
 - a. In a metro node Local:
 - i. Select **Provision Storage**.
 - ii. In the **View by** drop-down, select **Devices**.
 - b. In a metro node Metro:
 - i. Select **Provision Storage**.
 - ii. In the **Provision on** drop-down, select the cluster of the devices.
 - iii. In the **View by** drop-down, select **Devices**.
2. In the **Devices** view, click the device name.
The **Device Properties** dialog box appears. You can see the Service Status in the dialog box.

Creating a device

Steps


1. From the GUI main menu, do one of the following:
 - a. In a metro node Local:
 - i. Select **Provision Storage**.
 - ii. In the **View by** drop-down, select **Devices**.
 - b. In a metro node Metro:
 - i. Select **Provision Storage**.
 - ii. In the **Provision on** drop-down, select the cluster of the devices.
 - iii. In the **View by** drop-down, select **Devices**.
2. In the **Devices** view, click **CREATE**.
The Create Devices wizard is displayed.
3. Follow the instructions in the wizard. For additional help on any of the screens, click **Help**.

Renaming a device

About this task

You can rename a device at any time.

Steps

1. From the GUI main menu, do one of the following:
 - a. In a metro node Local:
 - i. Select **Provision Storage**.
 - ii. In the **View by** drop-down, select **Devices**.
 - b. In a metro node Metro:
 - i. Select **Provision Storage**.
 - ii. In the **Provision on** drop-down, select the cluster of the devices.
 - iii. In the **View by** drop-down, select **Devices**.
2. In the **Devices** view, click the device name link to open the device properties panel.
3. Click the **Rename** under the **DEVICE** tab in the Device Properties panel, and then type a new name in the **Device Name** field.
4. Click .

Deleting a device

About this task

You may delete one or more top-level devices that do not have virtual volumes. If the device's children are other devices, the child devices become top-level devices. If the device's children are extents, the status of those extents change from `Used` to `Claimed`.


Steps

1. From the GUI main menu, do one of the following:
 - a. In a metro node Local:
 - i. Select **Provision Storage**.
 - ii. In the **View by** drop-down, select **Devices**.
 - b. In a metro node Metro:
 - i. Select **Provision Storage**.
 - ii. In the **Provision on** drop-down, select the cluster of the devices.
 - iii. In the **View by** drop-down, select **Devices**.
2. In the **Devices** view, click **Delete**.
3. Review the devices to delete, and then click **OK**.
4. Review the results, and then click **Close**. You can sort the results.
If the device had child devices, the dialog box also displays the child devices have now become top-level devices.

Mirroring a device

About this task

You can mirror a device locally or remotely. When you add a local mirror, a device in the local cluster is used to create a RAID-1 device. If the local device you are mirroring is already a RAID-1, the mirror is added as another leg. For a remote mirror, a device in the remote cluster is used to create a distributed device. You can mirror a device with or without a volume on top.

 **NOTE:** When creating a remote mirror, the mirror must not be in a consistency group.

Steps

1. From the GUI main menu, do one of the following:
 - a. In a metro node Local:
 - i. Select **Provision Storage**.
 - ii. In the **View by** drop-down, select **Devices**.
 - b. In a metro node Metro:
 - i. Select **Provision Storage**.
 - ii. In the **Provision on** drop-down, select the cluster of the devices.
 - iii. In the **View by** drop-down, select **Devices**.
2. In the **Devices** view:
 - To add a local mirror, select **MORE > Add Local Mirror**. The Add Local Mirror wizard opens.
 - To add a remote mirror, select **MORE > Add Remote Mirror**. The Add Remote Mirror wizard opens.
3. Follow the instructions on the wizard screens. For additional help on any wizard screen, click **Help**.

Device status

The following tables define the Health, Operational, and Service status for devices.

Health status	Description
OK	The device is functioning normally.
Degraded	The device may be out-of-date compared to its mirror. (This state applies only to devices that are part of a RAID-1 device.)

Health status	Description
Unknown	Metro node cannot determine the device's Health state, or the state is invalid.
Non-recoverable error	The device may be out-of-date compared to its mirror (applies only to devices that are part of a RAID-1 device), and/or metro node cannot determine the Health state.
Minor Failure	One or more children of the device is out-of-date, but will rebuild.
Major Failure	One or more children of the device is out-of-date and will never rebuild, possibly because they are dead or unavailable.
Critical Failure	All children of this device are dead or unavailable.
Isolated	One or more storage volumes under this mirror is not performing optimally. Metro node has automatically marked the mirrors as isolated to prevent I/O to them.

Operational status	Description
OK	The device is functioning normally.
Degraded	The device may be out-of-date compared to its mirror. (This applies only to devices that are part of a RAID-1 device.)
Unknown	Metro node cannot determine the device's Operational state, or the state is invalid.
Stressed	One or more children of the device is out-of-date and will never rebuild.
Error	One or more children of the device is hardware-dead.
Starting	The device is not yet ready.

Service status	Description
Running	The device is functioning normally.
Suspended	The device is accepting new I/O; pending I/O requests are frozen. This state could be due to unreachable, dead, or corrupt storage.

Device component properties

Component Name The name of the device's component. You can change the name from this dialog box.

Geometry The underlying RAID structure of the device.

Health The overall health of the component.

Health Indications A list of reasons for the component's unhealthy state.

Operational Status Indicates how the component is functioning in the cluster. See Operational states.

Service Status Indicates whether I/O can proceed on the device. See Service states.

Device properties

Device Name The name of the top-level device. You can change the name in the properties dialog box. The name must be unique across clusters.

Geometry The underlying RAID structure of the device.

Virtual Volume The name of the virtual volume created on the device.

Locality Indicates where the device is visible: Local or Remote.

Cluster The name of the cluster to which the device belongs.

Block Count	The number of blocks in the underlying extent.
Block Size	The size of the data blocks used internally when managing I/O. The block size determines the granularity of storage volume I/O and the rebuild rate of distributed devices.
Capacity	The size of the device.
Thin Capable	Indicates if the underlying storage volume is thinly-provisioned.
Storage Array Family	<p>The family of arrays that the supporting array belongs to. Possible values:</p> <ul style="list-style-type: none"> • XTREMIO — The array belongs to the XtremIO family of arrays. • CLARiion — The array belongs to the VNX or Unity family of arrays. • SYMMETRIX — The array belongs to the VMAX or Symmetrix family of arrays. • Mixed — The supporting arrays belong to different array families. For example, a RAID-1 device with one leg on a VNX array and the other leg on a VMAX array. • Other — metro node does not recognize the array or does not support creating thin volumes on the array. • - (dash) The Storage Array Family property is not populated. There may be a connectivity problem.
Stripe Depth	For RAID-0, the depth of the stripe in bytes.
Health	The overall health of the device. See Health states .
Health Indications	The reasons for the device's unhealthy state. If the Health Indications include rebuilding, you can view the rebuild status for the device by clicking the Go to Rebuild Status view link.
Operational Status	Indicates how the device is functioning in the cluster. See Operational states .
Service Status	Indicates the status of I/O to the device. See Service states .

Distributed device properties

Distributed Device Name	The name of the device. You can change the name in this dialog box.
Transfer Capacity	The maximum capacity at which a distributed device can mirror data on the underlying devices. The transfer capacity is always the same as the capacity of the smallest device, since the system cannot mirror more data than the smallest device can handle.
Virtual Volume	The name of the virtual volume created on the device. By default, the virtual volume inherits the name of the device with <code>_vol</code> appended at the end of the name.
Thin Capable	Indicates if the underlying storage volume is thinly-provisioned.
Storage Array Family	<p>The family of arrays that the supporting storage array belongs to. Possible values:</p> <ul style="list-style-type: none"> • XTREMIO — The array belongs to the XtremIO family of arrays. • CLARiion — The array belongs to the VNX or Unity family of arrays. • SYMMETRIX — The array belongs to the VMAX or Symmetrix family of arrays. • Mixed — The supporting arrays belong to different array families. For example, you might have a RAID-1 device with one leg on a VNX array and the other leg on a VMAX array. • Other — metro node does not recognize the array or does not support creating thin volumes on the array. • - (dash) — The Storage Array Family property is not populated. There may be a connectivity problem.
Rule Set	The name of the rule set applied to the device. You can change the rule set in this dialog box.
Health	The overall health of the device. See Health states .
Health Indications	Lists the reasons for the device's unhealthy state. If the Health Indications of the distributed device indicate that it is rebuilding, click the Go to Rebuild Status view link to monitor the rebuild status of the distributed device.

Operational Status	Indicates how the device is functioning. See Operational states .
Service Status	Indicates whether I/O can proceed on the device. See Service states .

Add capacity to virtual volumes

Use this dialog box to expand a virtual volume using concatenation.

Add Capacity to Virtual Volume	The selected virtual volume that you want to expand.
Current Capacity	The size of the selected virtual volume.
Additional Capacity Needed	Enter the amount of space to add to this virtual volume. Metro node will look for volumes with at least that available capacity. To enter the value using a different unit, select the correct unit from the drop-down list. When you enter the amount of space needed, you can click Show Available to list the storage volumes with a minimum capacity equal to the additional capacity needed. If you change the capacity needed, click Show Available again to update Available Storage Volumes list.
Total New Capacity	The new size of the virtual volume based on the additional capacity needed.
Available storage volumes	The list of available storage volumes to choose from based on the additional capacity needed. The name of the storage volume and the largest free chunk available in the storage volume are displayed.

Extent properties

Extent Name	The name of the extent. You can change the name of the extent from this dialog box. By default, extents are named after the underlying storage volume with the prefix "extent_".
Use	Indicates whether the extent is currently in use. Possible values: Claimed or Used.
Used By	The name of the device created from the extent.
Storage Volume	The name of the underlying storage volume.
Block Count	The number of blocks in the underlying extent.
Block Offset	The starting block number for the extent. If not specified, extents are created using sequential block numbers on a storage volume. To specify the starting block number, use the CLI.
Block Size	The size of the data blocks used internally when managing I/O. The block size determines the granularity of disk I/O and the rebuild rate of distributed devices.
Capacity	The size of the extent.
Thin Capable	Indicates if the underlying storage volume is thinly-provisioned.
Storage Array Family	The family of arrays that the supporting array belongs to. Possible values: <ul style="list-style-type: none"> • XTREMIO — The array belongs to the XtremIO family of arrays. • CLARiion — The array belongs to the VNX or Unity family of arrays. • SYMMETRIX — The array belongs to the VMAX or Symmetrix family of arrays. • Other — metro node does not recognize the array or does not support creating thin volumes on the array. • - (dash) — The Storage Array Family property is not populated. There may be a connectivity problem.
Health	Indicates the overall health of the extent. See Health states .
Health Indications	Displays the reasons for the extent's unhealthy state.
I/O Status	The status of I/O to the underlying storage volume. See I/O states .

Operational Status

Indicates how the extent is functioning in the cluster. See [Operational states](#).

Extents

About extents

An extent is a range of blocks of a storage volume. You can only create a full-size extent that spans the entire capacity of the storage volume. Extents are used to create devices, and then virtual volumes on top of those devices.

Extent names

Extents are automatically named after the underlying storage volume using the "extent_" prefix. For example, an extent that is created from a storage volume that is named *Storagevolume1* will be named *Extent_Storagevolume1_1*. You can change the name at any time in the **Extent Properties** dialog box.



Figure 3. Extent configurations

Using the Extents view

The **Extents** view shows all extents in a selected cluster. By default, extents are sorted by Use, showing all claimed extents first. The arrow in the default sort column indicates the direction of the sort; ascending (↑) or descending (↓). You can [sort](#) columns of data, [select](#) one or more objects in the list, and [show or hide](#) properties links on the screen.

In the main panel of the screen, you can perform these tasks:

- Create new extents
- Create devices
- Delete extents

This screen also shows the following information about extents:

Column	Description
Extent	The name of the extent. By default, the name includes the underlying storage volume name.
Capacity	The size of the extent.
Use	Indicates whether the extent is currently in use. Possible values: Claimed or Used.
Health	The overall health of the extent. Extent status provides more information.
Status	Indicates how the extent is functioning in the cluster. Extent status provides more information.

Column	Description
Storage Volume	The name of the underlying storage volume.



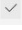

Extent Properties Panel

This screen displays the properties of the extents in a separate panel at the extreme right side of the screen. Using the **Extents Properties** panel, you can:

- View the details of the extent
- View the extent map
- Rename the extent
- View the details of the storage volume from where the extent has been created
- Turn on thin rebuild on the storage volume
- View the storage volume map
- Rename the storage volume

Tasks in this Screen

You can also perform the following tasks from this screen:

- Search for an extent — Click the  icon corresponding to the element and type the full or the partial name in the field that appears. Press **Enter**.
- Refresh screen data — Click the Refresh icon .
- Create an extent - Click **CREATE** to open the Create Extents wizard.
- Create devices from extents — Click **CREATE DEVICES** to open the Create Devices using the Extents wizard.
- View the Map of the Extent — Select the extent, and then click **VIEW MAP** under the **EXTENT** tab in the Extent Properties panel.
- Delete an extent — Select a claimed extent, and then click **DELETE**. The **Use** column indicates if an extent is in use.
- Rename an extent — Click the extent, and then click **Rename** under the **EXTENT** tab in the **Extent Properties** panel. Click  to save the changes.
- View the storage volume map — Select the extent, and then click **VIEW MAP** under the **STORAGE VOLUME** tab in the **Extent Properties** panel.
- Rename a storage volume — Click the extent, and then click **Rename** under the **STORAGE VOLUME** tab in the **Extent Properties** panel. Click  to save the changes.

Opening the Extents view

From the GUI main menu, do one of the following:

1. In a metro node Local:
 - a. Select **Provision Storage**.
 - b. In the **View by** drop-down, select **Extents**.
2. In a metro node Metro:
 - a. Select **Provision Storage**.
 - b. In the **Provision on** drop-down, select the cluster of the extent.
 - c. In the **View by** drop-down, select **Extents**.

The Create Extents wizard

Use the Create Extents wizard to create extents from selected storage volumes. By default, metro node creates one extent equal to the size of the selected storage volume. However, you can create an extent that is less than the capacity of the storage volume, and use the unused capacity later. You can create a maximum of 128 extents per storage volume.

Best practices

- If the storage volume is larger than the wanted virtual volume, create an extent the size of the wanted virtual volume. Do not create smaller extents, and then use different RAID types to concatenate or stripe the extents.
- If the storage volume is smaller than the wanted virtual volume, create a single extent per storage volume, and then use devices to concatenate or stripe these extents into a larger device.

When a storage volume fails to create an extent

The **View Results** dialog box shows any storage volumes that failed to complete the operation. It may happen for any of the following reasons:

- The storage volume may have become unreachable shortly before or during the operation.
- The storage volume may have started to return read or write failures on other extents that have already been exported as virtual volumes.

If it occurs, investigate the status of the storage volume in the **Storage volumes** view and proceed accordingly.

Creating an extent

About this task

Create extents from storage volumes. You can create an extent equal to the size of the selected storage volume, or less than the capacity of the storage volume, and use the unused capacity later. You can create a maximum of 128 extents per storage volume.

Steps

1. From the GUI main menu, do one of the following:
 - a. In a metro node Local:
 - i. Select **Provision Storage**.
 - ii. In the **View by** drop-down, select **Extents**.
 - b. In a metro node Metro:
 - i. Select **Provision Storage**.
 - ii. In the **Provision on** drop-down, select the cluster of the extent.
 - iii. In the **View by** drop-down, select **Extents**.
2. In the **Extents** view, click **CREATE**.
The Create Extents wizard is displayed.
3. Complete the wizard. For additional help on any of the wizard screens, click **Help**.


Renaming an extent

About this task

You can rename an extent at any time.

Steps

1. From the GUI main menu, do one of the following:
 - In a metro node Local:
 - a. Select **Provision Storage**.
 - b. In the **View by** drop-down, select **Extents**.
 - In a metro node Metro:
 - a. Select **Provision Storage**.
 - b. In the **Provision on** drop-down, select the cluster of the extent.
 - c. In the **View by** drop-down, select **Extents**.
2. In the **Extents** view, click the extent name link.
3. In the **Extent Properties** panel, click **Rename** and type a new name for the extent in the **Extent Name** field.

- Click  to save the new name.

Deleting an extent

About this task

You can delete one or more unused extents.

Steps

- From the GUI main menu, do one of the following:
 - In a metro node Local:
 - Select **Provision Storage**.
 - In the **View by** drop-down, select **Extents**.
 - In a metro node Metro:
 - Select **Provision Storage**.
 - In the **Provision on** drop-down, select the cluster of the extent.
 - In the **View by** drop-down, select **Extents**.
- In the **Extents** view, select the extents to delete, and then click **Delete**.
- Review the extents to delete, and then click **OK**.
- Review the results of the Delete operation, and then click **Close**. You can sort the results in the dialog box.

Extent status

The following tables define an extent's Health, Operational, and I/O status.

Health status	Description
OK	The extent is functioning normally.
Degraded	The extent may be out-of-date compared to its mirror. (This state applies only to extents that are part of a RAID-1 device.)
Unknown	Metro node cannot determine the extent's Health state, or the state is invalid.
Non-recoverable error	The extent may be out-of-date compared to its mirror (applies only to extents that are part of a RAID-1 device), and/or the Health state cannot be determined.
Isolated	One or more storage volumes under this mirror is not performing optimally. metro node has automatically marked the mirrors as isolated to prevent I/O to them.

Operational status	Description
OK	The extent is functioning normally.
Degraded	The extent may be out-of-date compared to its mirror. (This state applies only to extents that are part of a RAID-1 device.)
Unknown	Metro node cannot determine the extent's Operational state, or the state is invalid.
Starting	The extent is not yet ready.

I/O status	Description
Alive	I/O is proceeding normally on the extent.
Dead	The underlying storage volume is marked as hardware-dead.
Unreachable	The underlying storage volume is unreachable.

Distributed devices

About distributed devices

Distributed devices are configured by using storage from both clusters, and are used only in a metro node Metro system. A distributed device's components must be other devices, and those devices must be created from storage in both clusters. You can also create a distributed device by mirroring a device remotely.

Devices can be components of other devices. If a device is not a component of another device, it is referred to as a top-level device. Distributed devices can be created only from top-level devices.

To create a distributed device, a Logging volume must exist in the system. Logging volumes are used to keep track of blocks written during an inter-cluster link failure or when one leg of a distributed RAID-1 becomes unreachable and then recovers. After a link failure is restored or an unreachable leg recovers, metro node uses the information in logging volumes to synchronize the mirrors by sending only changed blocks across the link.




Figure 4. Distributed device configurations

The Distributed Devices view

The **Distributed Devices** view shows all distributed devices in the system. Use this view to quickly see the status of a distributed device, the components of a distributed device, or search for a distributed device. By default, distributed devices are sorted ascending alphabetical order. The arrow in the default sort column indicates the direction of the sort; ascending (↑) or descending (↓). You can [sort](#) columns of data, [select](#) one or more objects in the list, and [show/hide properties](#) on the screen.

This view shows the following information about distributed devices. To see additional properties, click the distributed device name link to open the properties dialog box.

Column	Description
Device	The name of the distributed device. You can change the name in the properties dialog box.
Geometry	Indicates the underlying RAID structure of the distributed device.
Capacity	The size of the distributed device.
Health	The overall health of the distributed device. Distributed device status provides more information.
Status	Indicates how the distributed device is functioning. Distributed device status provides more information.
Virtual Volume	The name of the virtual volume created on the distributed device.
Rule Set	The name of the Rule Set applied to the distributed device. You can change the Rule Set in the Distributed Device Properties dialog box.

Additional tasks

You can perform these additional tasks from this screen:

- View properties of supporting devices — Click a device name in the **Component of Selected device** panel to open its properties dialog box.
- Change the rule set for a distributed device — Click the distributed device name to display the properties dialog box. In the properties panel, click **Change** to select a new rule set, and then click OK.
- Search for a distributed device — Type the device name in the **Search** text box and press **Enter**.

Opening the Distributed Devices view

From the GUI main menu, do the following:

1. In a metro node Local:
 - a. Select **Provision Storage**.
 - b. In the **View by** drop-down, select **Devices**.
2. In a metro node Metro:
 - a. Select **Provision Storage**.
 - b. In the **Provision on** drop-down, select the **Distributed Storage**.
 - c. In the **View By** drop-down, select **Devices**.

The Create Distributed Device from Claimed Storage Volumes wizard

Use the Create Distributed Device from Claimed Storage wizard to create a distributed device. You can also create a distributed device by adding a remote mirror to the device. [Mirroring a device](#) provides for more information.

Creating thin distributed virtual volumes

To create thin distributed virtual volumes, the supporting arrays must belong to the same array family and the storage volumes must also be thin capable.

Consistency groups and rule sets

If you add a distributed device to a consistency group, a volume is automatically created on the device. If you do not add the device to a consistency group, creating a virtual volume on the device is optional. The virtual volume inherits the distributed device name, with `_vol` appended to the end of the name.

Distributed devices in a consistency group use the rule set applied to the consistency group. If the distributed device will not be in a consistency group, you must select a rule set for the distributed device. You can change the rule set in the distributed device properties dialog box at any time. Also, you can use the CLI to create your own detach rule for the distributed device (as long as the device is not in a consistency group). For more information about creating and managing detach rules, see the *CLI Guide for metro node*.

Data synchronization

You can indicate whether to perform an initial data synchronization from the source device to the target device or not. It is an initial data synchronization. If you choose not to perform an initial synchronization, future writes will still be mirrored. If the source device contains data you want to preserve, then synchronize the data. If the device does not contain data or needs formatting, then you do not need to synchronize.

If the device has a virtual volume, it is automatically synchronized. Any data on the target device will be overwritten.

Logging volumes

Logging volumes are created during initial system setup and are required in each cluster to keep track of any blocks written during a loss of connectivity between clusters. After an inter-cluster link is restored, the logging volume is used to synchronize distributed devices by sending only changed blocks over the inter-cluster link. The *CLI Guide for metro node* provides more information about logging volumes.


To create a distributed device, a logging volume must exist in the system. If no logging volume exists in the system, an error message appears to inform you that you cannot create a distributed device without a logging volume. In the GUI, the logging volume is automatically selected for the distributed device. To specify a logging volume to be used for a distributed device, use the CLI.

Distributed device rule sets

Rule sets are predefined rules that determine which cluster continues I/O during an inter-cluster link failure or cluster failure. In these situations, until communication is restored, most I/O workloads require specific sets of virtual volumes to resume on one cluster and remain suspended on the other cluster.

In a metro node Metro configuration, there is a Management Console on the management server in each cluster. You can create distributed devices by using the GUI or CLI on either management server. During communication problems, the default rule set makes the cluster that is used to create the distributed device the winner cluster. I/O proceeds on the winner cluster and is suspended on the loser cluster. The following table describes how I/O proceeds when you use the default rule sets:

Rule set	Cluster-1	Cluster-2
Cluster-1 detaches	Services I/O	Suspends I/O
Cluster-2 detaches	Suspends I/O	Services I/O


 **NOTE:** Rule sets are applied per distributed device.

You can change a distributed device's default rule set in the **Distributed Device Properties** dialog box if the distributed device is not in a consistency group. The *CLI Guide for metro node* provides detailed information about creating and managing rule sets.

Changing the rule set for a distributed device

About this task

By default, the GUI applies the cluster-1-detaches rule set to all distributed devices, however, you can change this rule set after creating the device as described in the following procedure. When a cluster detaches, I/O continues on the detached (winning) cluster. Distributed device rule sets provides more information about rule sets.

 **NOTE:** You cannot use this option to change the rule set of a distributed device that is in a consistency group.

Steps

1. From the GUI, do the following:
 - a. In a metro node Local:
 - i. Select **Provision Storage**.
 - ii. In the **View by** drop-down, select **Devices**.
 - b. In a metro node Metro:
 - i. Select **Provision Storage**.
 - ii. In the **Provision on** drop-down, select the **Distributed Storage**.
 - iii. In the **View By** drop-down, select **Devices**.

The **Distributed Devices** screen is displayed.
2. To view the distributed device properties of a device at the right panel of the screen, click the distributed device name link.
3. In Device tab, click **Change**, and then select a new rule set from the **Rule Set** drop-down.
4. Click **OK** to save your changes.

Creating a distributed device

About this task

A distributed device uses storage from both clusters in a metro node Metro. To create a distributed device, use the Create Distributed Devices from Claimed Storage wizard. The wizard uses claimed storage to allow you to easily create both legs of a distributed device in a single operation.

You can also create a distributed device by adding a remote mirror to a device. [Mirroring a device](#) provides more information.

Steps

1. From the GUI main menu, do the following:
 - a. In a metro node Local:
 - i. Select **Provision Storage**.
 - ii. In the **View by** drop-down, select **Devices**.
 - b. In a metro node Metro:
 - i. Select **Provision Storage**.
 - ii. In the **Provision on** drop-down, select the **Distributed Storage**.
 - iii. In the **View By** drop-down, select **Devices**.
2. Click **CREATE**.

The Create Distributed Devices is displayed.
3. Complete the wizard. For additional help using the wizard, click **Help**.

Deleting a distributed device

About this task

The Delete operation deletes a distributed device and the unexported virtual volume on the device. It leaves the underlying device structure intact, and returns the top-level devices to the appropriate cluster. The Delete operation will fail if you try to delete a distributed device that has an exported virtual volume.

Steps

1. In the GUI main menu, do the following:
 - a. In a metro node Local:
 - i. Select **Provision Storage**.
 - ii. In the **View by** drop-down, select **Devices**.
 - b. In a metro node Metro:
 - i. Select **Provision Storage**.
 - ii. In the **Provision on** drop-down, select the **Distributed Storage**.
 - iii. In the **View By** drop-down, select **Devices**.


2. In the **Distributed Devices** screen, select the distributed devices to delete, and then click **DELETE**.
A confirmation dialog box is displayed.
3. Review the distributed devices to delete and the clusters to which the underlying devices will return upon completing the operation. If satisfied, click **Yes**.
The results dialog box is displayed.
4. Review the results of the operation and proceed accordingly.
5. Click **Close** to continue.
The underlying devices are now visible in the **Devices** screen under the appropriate clusters.

Renaming a distributed device

About this task

When you rename a distributed device, the name of the virtual volume on the device does not change. To rename the virtual volume, refer to [Renaming a virtual volume](#).

Steps

1. In the GUI main menu, do the following:
 - a. In a metro node Local:
 - i. Select **Provision Storage**.
 - ii. In the **View by** drop-down, select **Devices**.
 - b. In a metro node Metro:
 - i. Select **Provision Storage**.
 - ii. In the **Provision on** drop-down, select the **Distributed Storage**.
 - iii. In the **View By** drop-down, select **Devices**.
2. In the **Distributed Devices** view, click the distributed device name link in the **Device** column.
3. In the **Distributed Device Properties** panel, click **Rename** and type a new name in the **Distributed Device Name** field.
4. Click .

Distributed Device status

The following tables define the Health, Operational, and Service status for distributed devices.

Health status	Description
OK	The distributed device is functioning normally.
Unknown	Metro node cannot determine the distributed device's health state, or the state is invalid.
Non-recoverable error	Metro node cannot determine the distributed device's health state.
Minor Failure	Either one or more children of the distributed device is out-of-date and will rebuild, or the Logging volume for the distributed device is unhealthy.
Major Failure	One or more children of the distributed device is out-of-date and will never rebuild, possibly because they are dead or unavailable.

Operational status	Description
OK	The distributed device is functioning normally.
Degraded	The distributed device may have one or more out-of-date children that will eventually rebuild.
Unknown	Metro node cannot determine the distributed device's Operational state, or the state is invalid.
Stressed	One or more children of the distributed device is out-of-date and will never rebuild.
Error	One or more components of the distributed device is hardware-dead.

Operational status	Description
Starting	The distributed device is not yet ready.

Service status	Description
Running	The distributed device is functioning normally.
Winner–running	This cluster detached the distributed device while the other cluster was unreachable, and is now sending I/O to the device.
Cluster–unreachable	Metro node cannot reach the cluster; the status is unknown.
Potential–conflict	This cluster has detached the distributed device in order to resume I/O while the other cluster was unreachable. This creates a potential conflict when the clusters can communicate again.
Need–resume	The other cluster detached the distributed device while it was unreachable. The distributed device needs to be manually resumed for I/O to resume at this cluster.
Need–winner	All clusters are reachable again, but both clusters had detached this distributed device and resumed I/O. You must pick a winner cluster whose data will overwrite the other cluster's data for this distributed device.
Suspended	The distributed device is not accepting new I/O; pending I/O requests are frozen.

Virtual volumes

About virtual volumes

A virtual volume is created on a device or a distributed device, and is presented to a host through a storage view. Virtual volumes are created on top-level devices only, and always use the full capacity of the device or distributed device.

Exporting volumes to hosts

A virtual volume is not visible to a host until you add it to a storage view. Exporting a virtual volume consists of the following steps:

1. Registering an initiator
2. Creating a storage view
3. Adding initiators to the storage view
4. Adding ports to the storage view
5. Adding virtual volumes to the storage view

Volume visibility

You can enable or disable remote access for a virtual volume. Remote access allows a virtual volume in one cluster to be visible in the other cluster. After remote access is enabled, the virtual volume will display in both clusters with global visibility. Global visibility indicates that the virtual volume is visible in both clusters. Local visibility indicates that it is visible only at the cluster under which you are viewing the virtual volume.

Volume names

Virtual volume names are assigned automatically based on the device name. For example, if a device is named Testdevice, by default, the volume created on that device will be named Testdevice_vol. You can rename the virtual volume at any time in the **Virtual Volume Properties** dialog box.

Expanding volumes

You can non-disruptively increase the capacity of a virtual volume by selecting a storage volume with available capacity and specifying whether to use all or a portion of the capacity. metro node automatically creates an extent based on the specified capacity. To expand a volume using existing extents or devices, use the CLI. The *CLI Guide for metro node* provides more information.

Consistency groups

Consistency groups allow you to group volumes together and apply a set of properties to all the volumes in the group. Consistency groups ensure consistent behavior among the volumes in the group, and eliminates setting specific values to each virtual volume individually. About consistency groups provides detailed information about consistency groups and how they are used.

Monitoring volume performance


The Virtual Volumes Dashboard shows performance metrics of the top 100 busiest volumes, ranked by total IOPS. Use this data to monitor virtual volume performance or investigate poor volume performance. Refer to Virtual Volumes Dashboard for more information.

Configuration limits for volumes

The *Release Notes for metro node* for this release provide information on configuration limits when creating virtual volumes.

The Virtual Volumes view

The **Virtual Volumes** view shows all the virtual volumes in the selected cluster. Use this view to quickly see the status of a virtual volume, or the supporting device for a virtual volume. By default, virtual volumes are sorted by name in ascending alphabetical order. The arrow in the default sort column indicates the direction of the sort; ascending (↑) or descending (↓). You can sort columns of data, select one or more objects in the list, and show or hide on the screen.

 **NOTE:** You cannot sort on the **Visibility** column.

This screen also shows the following information about virtual volumes. To see additional properties, click the virtual volume name link to open the properties dialog box.

Name	Definition
Virtual Volume	The name of the virtual volume.
Capacity	The size of the virtual volume.
Health	The overall health of the virtual volume. Virtual volume status provides more information.
Operational Status	Indicates how the virtual volume is functioning in the cluster. Virtual volume status provides more information.
Service Status	Indicates how the virtual volume is functioning in the cluster. Virtual volume status provides more information.
Locality	The location of the supporting device. Possible values: Local, Remote, or Distributed.
Visibility	Indicates where a virtual volume is visible: Global (both clusters) or Local (only at this cluster).
Thin Enabled	Indicates if the virtual volume is created as a thin volume. Possible values: <ul style="list-style-type: none">Enabled — The volume was created as a thin volume.Disabled — The volume was not created as a thin volume although the underlying storage is thin-capable.Unavailable — The underlying storage is not thin-capable.


Name	Definition
Expandable by	The amount of space that is available for volume expansion. A volume that has an Expandable by of 0 can still be expanded using concatenation. A volume that has an Expandable by of > 0 can be expanded using storage volume expansion.
Expansion Method	The expansion method available for this volume. Possible values: <ul style="list-style-type: none"> • <code>concatenation</code> — The volume can be expanded using Concatenation or RAID-C expansion. • <code>storage-volume</code> — The volume can be expanded to the Expandable capacity using storage volume expansion. • <code>not-supported</code> — The volume does not support expansion.
Expansion Status	The expansion status of the volume. The volume can be: <ul style="list-style-type: none"> • <code>-</code> (dash) This volume can be expanded. • <code>failed</code> — The last volume expansion on this volume failed. <code>unknown</code> — The volume expansion status is unknown. • <code>in-progress</code> — The volume cannot be expanded because it has a volume expansion in progress.
Supporting Device	Shows the volume's supporting device. Click a device name to see its status and properties.

Virtual volume views and filters

To reduce the number of columns displayed on the page, you can select a specific view and filter the information in that view.

Additional tasks

You can also perform the following tasks from the **Virtual Volumes** view:

- View Virtual Volume properties — The screen displays the properties of the virtual volumes in a separate panel at the right side of the screen.
- View Supporting Device properties—Select the virtual volume, and then click the **Supporting Device** tab under the Virtual Volumes Properties panel.
- View Consistency Groups properties- Select the virtual volume, and then click the **Consistency Group** tab under the Virtual Volumes Properties panel.
- View a volume map —Select virtual volume, and then click the **VIEW MAP** under the **VIRTUAL VOLUMES** tab in the Virtual Volume Properties panel.
- Provision from pools —Select Virtual Volumes, and then select **CREATE > Provision from Pools**.
- Provision from storage volumes — Select Virtual Volumes, and then select **CREATE > Provision from Storage Volumes**.
- Rename virtual volumes — Select the virtual volume, and then click **Rename** under the **VIRTUAL VOLUME** tab in the Virtual Volume Properties panel. Click ☐ to save the changes.
- Search for virtual volumes — Type the virtual volume name in the **Search** box and press **Enter**.
- Refresh screen data — Click the Refresh icon .
- Create Virtual Volumes — Click **Provision from Pools**, **Provision from Storage Volumes**, or **Create from Devices**.
- Add a local or remote mirror — Select a virtual volume, and then select the **MODIFY>Add Local Mirror** or **Add Remote Mirror**.
- Expand volumes — Select an expandable virtual volume, and then select **MODIFY>Expand** to launch the Expand Volume wizard and add chunks of free storage volumes to the virtual volume.
- Enable or disable remote access — Select a virtual volume, and then select **MODIFY> Enable Remote Access** or **MODIFY>Disable Remote access**. This menu option allows you to enable remote access to this volume from another cluster in the metro node configuration. (metro node Metro only)
- Export— To export the Virtual Volumes to Storage view, click **MODIFY>EXPORT**.
- Delete volumes — Select an unexported volume, and then select **DELETE>Delete**.
- Tear down virtual volumes — Select an unexported virtual volume, and then select **DELETE>Tear Down** to delete all the associated devices and extents.

Opening the Virtual Volumes view

From the GUI main menu, do one of the following:

In metro node local

1. Click **Provision Storage**.
2. In the **View by** drop-down, click **Virtual Volumes**.

In metro node Metro

1. Click **Provision Storage**.
2. In the **Provision on** drop-down, select the cluster of the storage volume.
3. In the **View by** drop-down, select **Virtual Volumes**.

The Distributed Virtual Volumes view

The **Distributed Virtual Volumes** view lists all distributed virtual volumes in the system. Use this view to quickly see the status of a distributed virtual volume. By default, virtual volumes are sorted by name in ascending alphabetical order. The arrow in the default sort column indicates the direction of the sort; ascending (↑) or descending (↓). You can [sort](#) columns of data, [select](#) one or more objects in the list, and [show or hide](#) on the screen.

In the main panel of the screen, you can perform these tasks:

- Create a distributed virtual volumes
- Modify an existing distributed virtual volume
- Delete distributed virtual volumes

This screen also shows the following information about distributed virtual volumes:

Name	Definition
Virtual Volume	The name of the volume.
Capacity	The size of the volume.
Health	The overall health of the volume. Virtual volume status provides more information.
Operational Status	Indicates how the virtual volume is functioning. Virtual volume status provides more information.
Service Status	Indicates how the virtual volume is functioning. Virtual volume status provides more information.
Thin Enabled	Indicates if the virtual volume is created as a thin volume. Possible values: <ul style="list-style-type: none">• <code>Enabled</code> — The volume was created as a thin volume.• <code>Disabled</code> — The volume was not created as a thin volume although the underlying storage is thin-capable• <code>Unavailable</code> — The underlying storage is not thin-capable.
Expandable by	The amount of space that is available for volume expansion. A volume that has an Expandable by of 0 can still be expanded using concatenation. A volume that has an Expandable by of > 0 can be expanded using storage volume expansion.
Expansion Method	The expansion method available for this volume. Possible values: <ul style="list-style-type: none">• <code>concatenation</code> — The volume can be expanded using Concatenation or RAID-C expansion.• <code>storage-volume</code> — The volume can be expanded to the Expandable capacity using storage volume expansion.• <code>not-supported</code> — The volume does not support expansion.
Expansion Status	The expansion status of the volume. The volume can be: <ul style="list-style-type: none">• <code>-</code> (dash) This volume can be expanded.• <code>failed</code> — The last volume expansion on this volume failed. <code>unknown</code> — The volume expansion status is unknown.

Name	Definition
	<ul style="list-style-type: none"> <code>in-progress</code> — The volume cannot be expanded because it has a volume expansion in progress.
Consistency Group	The name of the consistency group to which the volume belongs.
Supporting Device	Shows the volume's supporting device. Click a device name to see its status and properties.






Distributed Virtual Volume Properties Panel

This screen also displays the properties of the distributed virtual volume in a separate panel at the extreme right side of the screen. Using the **Distributed Virtual Volume Properties** panel, you can:

- View the additional details of the distributed virtual volume
- Rename the distributed virtual volume
- View the consistency groups
- Rename the consistency groups and modify their detach rules
- View the details of the supporting devices
- Rename the supporting devices

Tasks in this screen

You can also perform the following tasks from the **Distributed Virtual Volumes** view:

- Search for a virtual volume — Click the  icon corresponding to the element and type the name in the field that appears.
- Refresh screen data — Click the Refresh icon .
- Create virtual volumes — Click **Provision from Pools**, **Provision from Storage Volumes**, or **Create from Distributed Devices** from **CREATE** drop-down to start the Create Virtual Volumes wizard.
- Modify virtual volumes — Click **Expand** or **Export** from **MODIFY** drop-down to modify the properties of the virtual volume.
- Delete virtual volumes — Click **Delete** or **Delete Device** from **DELETE** drop-down to delete an unexported virtual volume.
- Rename virtual volumes — Select a virtual volume, and then click **Rename** that corresponds to the virtual volume name under the **VIRTUAL VOLUME** tab in the **Virtual Volume Properties** panel. Click  to save the changes.
- Rename consistency groups — Click a distributed virtual volume and click **Rename** that corresponds to the consistency group name under the **CONSISTENCY GROUP** tab in the **Virtual Volume Properties** panel. Click  to save the changes.
- Modify the detach rules of the consistency groups — Click a distributed virtual volume and select an option from the **Detach Rules** drop-down under the **CONSISTENCY GROUP** tab in the **Virtual Volume Properties** panel.
- Rename supporting devices — Select a virtual volume, and then click **Rename** that corresponds to the supporting device name under the **SUPPORTING DEVICE** tab. Click  to save the changes.

Opening the Distributed Virtual Volumes view

1. From the GUI main menu, select **Provision Storage**.
2. In the **Provision on** drop-down, select **Distributed**.
3. In the **View by** drop-down, select **Virtual Volumes**.


Creating a virtual volume

There are several ways to create local or distributed virtual volumes in the GUI:

- Provision from storage pools — Provisioning from Pools provides detailed information about integrated storage and the provisioning process. To provision from storage pools, select **CREATE > Provision from Pools** from Provision Storage view.
- Provision from storage volumes — Uses integrated or non-integrated storage to provision a single virtual volume the size of a selected storage volume, or preserve data on an existing storage volume that you want to expose to hosts. Provisioning

from storage volumes provides detailed information. To provision from storage volumes, select **CREATE > Provision from Storage Volumes** from Provision Storage view.


- Create a virtual volume on top of a device (top-level) — Use this option if you did not create a virtual volume at the time that you created the device. This option allows you to add a virtual volume later. The volume is given the device name, followed by _vol. To create a virtual volume on top of a device, select **MORE>Create Virtual Volumes** from the **Devices** view.
- Use the **Create Virtual Volumes** wizard — Automatically creates a virtual volume using claimed or unclaimed storage volumes, depending on where you launch the wizard. Launching the wizard from the **Storage Arrays** view gives you the option of using claimed or unclaimed storage from the selected storage array, eliminating the separate step required to claim storage first. Launching from the **Storage Volumes** view uses claimed storage only. If the wizard fails to create the virtual volume, you can roll back the sub-components (extents and devices) created during the operation, or leave the sub-components and manually create the virtual volume on the device. To open the Create Virtual Volumes wizard, click **Create Virtual Volumes** from the **Storage Volumes** view or the **Storage Arrays** view.

 **NOTE:** To create thin volumes, the supporting arrays must belong to the same family of arrays and the storage volumes must also be thin-capable.

About virtual volume expansion

See the following method of volume expansion:

- Storage volume expansion — Allows online, non-disruptive expansion of the virtual volume taking the full "current capacity" of the storage volume and making that "configured capacity" of the virtual volume. [Expanding a volume using storage volumes](#) provides more information.

 **NOTE:** Volumes cannot be expanded during a rebuild.

Expanding a virtual volume using storage volumes

About this task

Storage volume expansion allows online, non-disruptive of the virtual volume, using the full current capacity of a storage volume and making that the configured capacity of the virtual volume. [About virtual volume expansion](#) provides more information.

Steps

1. From the GUI main menu, do one of the following:
 - In metro node local
 - a. Click **Provision Storage**.
 - b. In the **View by** drop-down, click **Virtual Volumes**.
 - In metro node Metro
 - a. Click **Provision Storage**.
 - b. In the **Provision on** drop-down, select the cluster of the storage volume.
 - c. In the **View by** drop-down, select **Virtual Volumes**.
2. In the Virtual volumes list, select one or more volumes with the Expansion Method of `storage-volume` and a capacity greater than 0.
3. Select **MODIFY>Expand**.
4. **Expand Virtual Volumes Confirmation** dialog box is displayed. Click **Yes** to expand the Virtual Volumes. **Expand Virtual Volumes Result** is displayed.

Enabling or disabling remote access for a volume

About this task

Remote access allows you to make a virtual volume on a device in one cluster visible in the other cluster. Once you make a virtual volume remotely accessible, you can include the volume in storage views in either cluster. The virtual volume's **Visibility** attribute is Global in the **Virtual Volumes** view of both clusters. Note that you cannot disable remote access on a virtual volume that is currently exported. To disable remote access, you must first remove the virtual volume from the storage view.

NOTE: Remote virtual volumes suspend I/O during inter-cluster link outages. As a result, the availability of the data on remote virtual volumes is directly related to the reliability of the inter-cluster link.

Steps

1. From the GUI main menu, do one of the following:
 - In metro node local
 - a. Click **Provision Storage**.
 - b. In the **View by** drop-down, click **Virtual Volumes**.
 - In metro node Metro
 - a. Click **Provision Storage**.
 - b. In the **Provision on** drop-down, select the cluster of the storage volume.
 - c. In the **View by** drop-down, select **Virtual Volumes**.
2. In the **Virtual Volumes** view, select the virtual volume, and then click **Enable Remote Access** or **Disable Remote Access**.
3. Review the virtual volumes, and then click **YES**.
4. Review the results of the operation, and then click **CLOSE**.

Manually assigning LUN numbers to volumes

By default, metro node automatically assigns LUN numbers to virtual volumes (0 - 16383) when creating a storage view or adding volumes to a view. However, you can manually assign LUN numbers. [Adding virtual volumes to a storage view](#) provides more information.

NOTE: You cannot change the LUN number of a virtual volume already in a storage view. To change the LUN number of a virtual volume already in a view, remove the virtual volume from the storage view, and then add it again using the desired LUN number.

Deleting a volume

About this task

Deleting a virtual volume deletes the virtual volume and leaves the underlying devices and extents intact. However, the storage volume status remains claimed. After the delete operation, the data on the volume is no longer accessible. Before you delete a virtual volume, be sure that you no longer need the data, or the data has been copied.

Before deleting a volume, be aware of the following:

- If the volume to delete is on a distributed device, delete the distributed device, which will automatically delete the virtual volume. [Deleting a distributed device](#) provides more information.
- If the volume to delete is a remote volume, delete the volume from the cluster that owns the underlying device (not the other cluster where the volume is visible).

NOTE: You can delete only unexported virtual volumes. To delete an exported virtual volume, you must first remove the volume from the storage view.

Steps

1. From the GUI main menu, do one of the following:
 - In metro node local
 - a. Click **Provision Storage**.
 - b. In the **View by** drop-down, click **Virtual Volumes**.
 - In metro node Metro
 - a. Click **Provision Storage**.
 - b. In the **Provision on** drop-down, select the cluster of the storage volume.
 - c. In the **View by** drop-down, select **Virtual Volumes**.
2. In the **Virtual Volumes** view, select the volume to delete and then click **DELETE > Delete**.


3. In the confirmation dialog box, review the volumes to delete, and if satisfied, click **Yes**.

Renaming a volume

About this task

Rename a virtual volume at any time. Renaming a virtual volume does not impact host IO to the volume.

Steps


1. From the GUI main menu, do one of the following:
 - In metro node local
 - a. Click **Provision Storage**.
 - b. In the **View by** drop-down, click **Virtual Volumes**.
 - In metro node Metro
 - a. Click **Provision Storage**.
 - b. In the **Provision on** drop-down, select the cluster of the storage volume.
 - c. In the **View by** drop-down, select **Virtual Volumes**.
2. In the **Virtual Volumes** view, select the volume name to see the properties at the right side of the screen.
3. In the **Virtual Volumes properties** dialog box, click **Rename** in Virtual Volume tab and type a new name in the **Name** field. The name must be unique within the cluster.
4. Click .

Tearing down a volume

About this task

Tearing down a virtual volume deletes the virtual volume and the underlying devices and extents. This operation makes the underlying storage volume capacity used by the virtual volume available for use again. Note that the underlying storage volume status remains `claimed`.

If the virtual volume you are tearing down is in a storage view, remove the volume from the storage view before performing this operation.

 **NOTE:** You cannot tear down a virtual volume on a distributed device. Deleting a virtual volume provides more information.

Steps

1. From the GUI main menu, do one of the following:
 - In metro node local
 - a. Click **Provision Storage**.
 - b. In the **View by** drop-down, click **Virtual Volumes**.
 - In metro node Metro
 - a. Click **Provision Storage**.
 - b. In the **Provision on** drop-down, select the cluster of the storage volume.
 - c. In the **View by** drop-down, select **Virtual Volumes**.
2. In the **Virtual Volumes** view, select an unexported volume to tear down.
3. Click the **Delete** drop-down, and then click **Tear Down**.
4. Review the virtual volumes to tear down, and if satisfied, click **Yes**.
5. Click **Close**.

Virtual Volume status

The following tables define a virtual volume's Health, Operational, and Service states.

Health status	Definition
OK	The virtual volume is functioning normally.
Unknown	Metro node cannot determine the virtual volume's Health state, or the state is invalid.
Non-recoverable error	Metro node cannot determine the virtual volume's Health state.
Minor Failure	One or more of the virtual volume's underlying devices is out-of-date, but will rebuild.
Major Failure	One or more of the virtual volume's underlying devices is out-of-date, but will never rebuild.

Operational status	Definition
OK	The virtual volume is functioning normally.
Degraded	The virtual volume may have one or more out-of-date devices that will eventually rebuild.
Unknown	Metro node cannot determine the virtual volume's Operational state, or the state is invalid.
Stressed	One or more of the virtual volume's underlying devices is out-of-date and will never rebuild.
Error	One or more of the virtual volume's underlying devices is hardware-dead.
Starting	The virtual volume is not yet ready.

Service status	Definition
Running	The virtual volume is accepting I/O.
Inactive	The virtual volume is part of an inactive or unhealthy storage view that the hosts cannot see.
Unexported	The virtual volume is not exported to a host.
suspended	I/O is suspended to the virtual volume.
Cluster unreachable	The virtual volume's cluster is unreachable.
Need resume	Applies to virtual volumes on distributed devices only. The other cluster detached the distributed device while it was unreachable. The distributed device needs to be manually resumed for I/O to resume at this cluster.
Disconnected	One or more legs on the underlying device have been disconnected.

Pool properties

Pool Name	The name of the pool available on the array.
Storage Array	The name of the managed array on which the Pool was created.
Supported Pool Types	The type of storage supported by the pool. Possible values are: Thin and/or Thick.
Policy	The name of the policy associated with the storage pool.
Free Capacity	The amount of free capacity available from the pool.
Used Capacity	The amount of capacity already used from the pool.
Total Capacity	The total capacity of the pool.

Virtual volume properties

Virtual Volume Name	The name of the virtual volume. You can change the name in this field.
Supporting Device	The device on which the virtual volume was created.
Consistency Group	Identifies the metro node consistency group to which this virtual volume belongs.
Locality	Designates where the virtual volume is located. Values can be cluster 1, cluster 2 or distributed.
Visibility	Indicates where a virtual volume is visible: Global (both clusters) or Local (only at this cluster).
Expansion Method	<p>The expansion method available for this volume. Possible values:</p> <p><code>concatenation</code> — The volume can be expanded using Concatenation or RAID-C expansion.</p> <p><code>storage-volume</code> — The volume can be expanded to the Expandable capacity using storage volume expansion.</p> <p><code>not-supported</code> — The volume does not support expansion.</p>
Expandable By	The amount of space that is available for volume expansion. A volume that has an Expandable by of 0 can still be expanded using concatenation. A volume that has an Expandable by of > 0 can be expanded using storage volume expansion.
Expansion Status	<p>The expansion status of the volume. The volume can be:</p> <ul style="list-style-type: none">– (dash) — This volume can be expanded.<code>failed</code> — The last volume expansion on this volume failed.<code>unknown</code> — The volume expansion status is unknown.<code>in-progress</code> — The volume cannot be expanded because it has a volume expansion in progress.
Block Count	The number of blocks in the underlying extent.
Block Size	The size of a single block, in kilobytes.
Capacity	The size of the virtual volume.
Thin Enabled	<p>Indicates if the virtual volume is created as a thin volume. Possible values:</p> <ul style="list-style-type: none">• <code>Enabled</code> — The volume was created as a thin volume.• <code>Disabled</code> — The volume was not created as a thin volume although the underlying storage is thin-capable.• <code>Unavailable</code> — The underlying storage is not thin-capable. <p>You can change this property for XtremIO, VMAX and VNX virtual volumes only in the CLI. Refer to the CLI Guide for information.</p>
Storage Array Family	<p>The family of arrays that the supporting array belongs to. Possible values:</p> <ul style="list-style-type: none">• <code>XTREMIO</code> — The array belongs to the XtremIO family of arrays.• <code>CLARiion</code> — The array belongs to the VNX or Unity family of arrays.• <code>SYMMETRIX</code> — The array belongs to the VMAX or Symmetrix family of arrays.• <code>Mixed</code> — The supporting arrays belong to different array families. For example, you might have a RAID-1 device with one leg on a VNX array and the other leg on a VMAX array.• <code>Other</code> — Metro node does not recognize the array or does not support creating thin volumes on the array.• – (dash) — The Storage Array Family property is not populated. There may be a connectivity problem.
Health	The overall status of the virtual volume. See Virtual volume states.
Health Indications	The reasons for the virtual volumes unhealthy state. See Virtual volume status.
Operational Status	Indicates how the virtual volume is functioning in the cluster. See Virtual volume status.
Service Status	Indicates how the virtual volume is functioning in the cluster. See Virtual volume status.

Show ITLs dialog box

This dialog box shows a read-only list of Initiator-Target-LUN combinations that shows the storage volume's back-end path (metro node back-end port, storage port, and LUN). you can also search for a specific initiator, target or LUN by entering the number in the **Find** text box and pressing **Enter**.

The following table describes the fields shown in the dialog box:

Field/Column	Description
Storage Volume Name	The name of the selected storage volume.
System ID	The array's ID number for the storage volume.
Initiator	The Port WWN of the metro node backend port.
Target	The Port WWN of the storage array target port.
LUN	The Logical Unit Number for the storage volume.

Logical unit properties

Name	The VPD identifier for the logical unit.
Active AAO Controller	The active (or ALUA optimized) controller for the logical unit.
Active AAO Visibility	Lists all directors that have a path to the active (or ALUA optimized) controller for this logical unit.
Passive AAN Controller	The passive (or ALUA non-optimized) controller for the logical unit.
Passive AAN Visibility	Lists all directors that have a path to the passive (or ALUA non-optimized) controller for this logical unit.
LUNs	The LUN number on the array.
Storage Volume	The storage volume to which this logical unit belongs.
ALUA Support	Indicates if Asymmetrical Logical Unit Access (ALUA) is supported on the array. Depending on the type of controller exposed to the LUN, one of these values will display. Possible values:

Implicit — The logical unit can change its ALUA access state internally within the array without any external input requesting a change.

Explicit — The logical unit can change its ALUA access state in response to a request from an initiator (metro node back end software).

Implicit explicit — The logical unit has both implicit and explicit capabilities.

None — ALUA is not supported on the array.

Connectivity Status	The connectivity status of the logical unit.
----------------------------	--

Visibility	Indicates how the logical unit is visible to the controllers. Possible values:
-------------------	--

Active optimized — All commands and task management functions can be sent on this path; this path has the highest performance (ALUA arrays only).

Active non-optimized — All commands and task management functions can be sent on this path, but the path has a lower performance than an active optimized path (ALUA arrays only).

Active — The controller actively processing I/O requests.

Both — Both controllers process I/O and provide standby capability for each other. One controller is actively processing I/O while the other is idle in standby mode waiting to process I/O should the active primary controller fail or be taken offline.

None — ALUA is not supported on the array.

ALUA Support field values

Value	Definition
Implicit	The logical unit can change its ALUA access state internally within the array without any external input requesting a change.
Explicit	The logical unit can change its ALUA access state in response to a request from an initiator (metro node back-end software).
Implicit explicit	The logical unit has both implicit and explicit capabilities.
None	ALUA is not supported on the array.

Visibility field values

Value	Definition
Active optimized	All commands and task management functions can be sent on this path; this path has the highest performance (ALUA arrays only).
Active non-optimized	All commands and task management functions can be sent on this path, but the path has a lower performance than an active optimized path (ALUA arrays only).
Active	The controller is actively processing I/O requests.
Passive	The controller is in standby mode waiting to process I/O should the active primary controller fail or be taken offline.
Both	Both controllers process I/O and provide standby capability for each other. One controller is actively processing I/O while the other is idle in standby mode waiting to process I/O should the active primary controller fail or be taken offline.
None	ALUA is not supported on the array.

Extent or Device mobility job properties

Mobility Job Name	The name of the mobility job.
Source	The name of the source extent or source device.
Target	The name of the target extent or target device.
Start Time	The date and time the job started. Note that N/A may display in this field if you are connected to one cluster and view the properties of a job that was started on another cluster.
Transfer Speed	The maximum number of bytes of data transferred at a time from the source to the target. You can change the transfer speed in the properties dialog box while the job is in the queue or in progress. Mobility job transfer speed provides more information on selecting the transfer speed.
Status	The status of the extent or device mobility job. Mobility job status provides more information.

Metro node port properties

Enabled Port	Indicates whether or not the port is enabled. Use the checkbox to enable or disable the port.
Port Name	The ID number for the metro node port.
Port Name	Name of the port.
Director ID	The ID of the director on which the port is located.
Port WWN	The port's World Wide Name (WWN). The WWN is a unique network identifier for the port.

Port Node WWN	The port's node World Wide Name (WWN). The WWN is a unique network identifier for the port's node.
Export Status	Indicates whether or not this port is used to export storage. See metro node port status .

Storage array properties

Storage Array Name	The name of the storage array manufacturer.
Storage Array Family	<p>The family of arrays that the storage array belongs to. Possible values:</p> <ul style="list-style-type: none"> • XTREMIO — The array belongs to the XtremIO family of arrays. • CLARiiON — The array belongs to the VNX or Unity family of arrays. • SYMMETRIX — The array belongs to the VMAX or Symmetrix family of arrays. • Other — Metro node does not recognize the array or does not support creating thin volumes on the array. • – (dash) — The Storage Array Family property is not populated. There may be a connectivity problem.
Controllers	A list of controllers in the array. This list is empty if the array is not supported.
Ports	A list of all port WWNs discovered for the array.
Logical Unit Count	The number of LUNs exported from the array.
Connectivity Status	Indicates if all LUNs exported from the array are visible to both directors. See Storage array connectivity status .

Storage view properties

Storage view Name	The name of the storage view. You can change the name in this field.
Operational Status	Indicates how the storage view is functioning in the cluster. See Storage view status .
Active	Indicates whether or not the storage view is active.

Storage volume properties


Storage Volume Name	The name of the storage volume. You can change the name in this field.
Use	Indicates whether the storage volume is Claimed, Unclaimed, or Used.
Used By	The name of the extent created from the storage volume.
Storage Volume Type	Describes the type of data stored on the storage volume: Data protected, Metadata or Normal
Integrated Services	The name of the supporting storage array.
Thin Rebuild	Indicates if a thin rebuild was selected for the storage volume.
System ID	The internal system ID for the storage volume.
Block Count	The number of blocks in the storage volume.
Block Size	The size of a single block in kilobytes.
Capacity	The size of the storage volume.
Thin Capable	Indicates if the storage volume is thinly-provisioned.

Storage Array Family	<p>The family of arrays that the supporting array belongs to. Possible values:</p> <ul style="list-style-type: none"> • XTREMIO — The array belongs to the XtremIO family of arrays. • CLARiion — The array belongs to the VNX or Unity family of arrays. • SYMMETRIX — The array belongs to the VMAX or Symmetrix family of arrays. • Other — Metro node does not recognize the array, or does not support creating thin volumes on the array. • - (dash) — The Storage Array Family property is not populated. There may be a connectivity problem.
Largest Free Chunk	The total number of bytes in the largest range of blocks available for use on the storage volume.
Health	The overall health of the storage volume. See Health states .
Health Indications	The overall health of the storage volume. See Health states .
I/O Status	The status of I/O to the storage volume. See I/O states .
Operational Status	Indicates how the storage volume is functioning in the cluster. See Operational states .

Create Virtual Volumes dialog box

Use this dialog box to create a virtual volume from an existing device. The Available Devices list shows all available top-level devices. Select a device from the list and then click **Add** to add the device to the Selected Devices list.

To create thin virtual volumes, select the **Create thin virtual volumes** checkbox and make sure that all selected devices are thin-capable.

 **NOTE:** Selecting a mix of devices (thin capable and not thin capable) will result in creating thin virtual volumes on thin devices and thick virtual volumes on thick devices.

Consistency group

About consistency groups

Consistency groups allow a set of volumes to be grouped and have the same set of properties. Grouping volumes ensures consistency of application data spanning multiple volumes, and provides an application consistent copy of the data that can be used to restart the application at any metro node cluster in the event of a cluster failure or inter-cluster link failure.

See the *Administration Guide for metro node* for a detailed description.

Types of consistency groups



Consistency groups can consist of local, global, or distributed volumes. The following table describes the attributes that define each type of consistency group.


Consistency Group Type	Storage at Clusters	Visibility
Local (Local volumes only)	cluster-1 or cluster-2	cluster-1 or cluster-2
Global (Local volumes that are visible at both clusters)	cluster-1 or cluster-2	cluster-1 and cluster-2
Distributed (Distributed volumes only)	cluster-1 and cluster-2	cluster-1 and cluster-2

Applying detach rules

When you create a consistency group that contains distributed or global virtual volumes, you must select a detach rule to determine if and where I/O continues during an inter-cluster link failure. Note that detach rules are applied per consistency group. [Creating a consistency group](#) provides detailed information on the detach rules available for consistency groups and their effect on I/O during a link failure.

Using the Consistency Groups view

The **Consistency Groups** view shows all the consistency groups in the selected cluster. Use this view to quickly view information about a consistency group. By default, consistency groups are sorted by name in ascending alphabetical order. The arrow in the default sort column indicates the direction of the sort; ascending () or descending (). You can [sort](#) columns of data, [select](#) one or more consistency groups in the list, and [show or hide properties](#) in the view.

 **NOTE:** The default view shows both distributed and non-distributed consistency groups. To display only distributed consistency groups, see the [View distributed consistency groups](#).

Default view

The screen shows the following information about the Consistency Groups:

Column	Description
Consistency Group	The name of the consistency group.
Operational Status	The overall status of the consistency group. If the consistency group is distributed and status at both clusters is not the same, the status that is not OK displays. Click the status link for detailed information about the status. If the consistency group is distributed, the GUI displays the status at each cluster. Consistency group status provides more information.
Detach Rule	The detach rule that applies to the consistency group.
Volumes	Indicates the number of volumes in the consistency group.
Visibility	Lists the clusters at which the consistency group is visible.
Storage at Clusters	Indicates if the volumes in the group use storage from one cluster, or both clusters.
Selected Consistency Group panel	Shows the virtual volumes in the selected consistency group, and allows you to add or remove virtual volumes from the consistency group.

Additional tasks

You can also perform the following tasks from the **Consistency Groups** view:

- View consistency group properties — Click the consistency group name to view the properties at the right side of the screen in the properties panel.
- View properties of virtual volumes in a consistency group — Click the virtual volume name in the **Components of Selected Consistency Group** panel on the right.
- Create a consistency group — Click **CREATE**.
- Delete a consistency group — Select a consistency group, and then click **DELETE**.
- Rename a consistency group — Click the consistency group name link to open its properties panel. Click **Rename** to change the Consistency Group Name..
- Search for a consistency group — Type the complete or partial group name in the **Search** text box, and then press **Enter**.
- Add or remove virtual volumes from a consistency group — Select a consistency group in the list, go to **VIRTUAL VOLUMES** tab in the properties panel, and then click **Add** or **Remove**.

Opening the Consistency Groups view

From the GUI main menu, do one of the following:

1. In a metro node Local:
 - a. Select **Provision Storage**.
 - b. In the **View by** drop-down, select **Consistency Groups**.
2. In a metro node Metro:
 - a. Select **Provision Storage**.
 - b. In the **Provision on** drop-down, select the cluster of the consistency groups.
 - c. In the **View by** drop-down, select **Consistency Groups**.

Distributed Consistency Groups view

The Distributed Consistency Groups screen shows all distributed consistency groups. Note that distributed consistency groups are used in metro node Metro configurations only. Use this view to quickly view information about a distributed consistency group. By default, groups are sorted by name in ascending alphabetical order. The arrow in the default sort column indicates the direction of the sort; ascending (



) or descending (



). You can [sort](#) columns of data, [select](#) one or more consistency groups in the list, and [show or hide properties](#) in the view.

The screen shows the following information about the distributed consistency groups:

Column	Description
Consistency Group	The name of the consistency group.
Operational Status	The overall status of the consistency group. If the consistency group is distributed and status at both clusters is not the same, the status that is not OK displays. Click the status link for detailed information about the status. If the consistency group is distributed, the GUI displays the status at each cluster. Consistency group status provides more information.
Detach Rule	The detach rule that applies to the consistency group.
Volumes	Indicates the number of volumes in the consistency group.

Additional tasks

You can do the following tasks from the **Distributed Consistency Groups** view:

- View distributed consistency group properties — Click the consistency group name to view the properties at the right side of the screen in the properties panel.
- View properties of virtual volumes in a the group — Click the virtual volume name in the **Components of Selected Consistency Group** panel on the right.
- Create a distributed consistency group — Click **CREATE**.
- Delete a group — Select a consistency group, and then click **DELETE**.
- Rename a consistency group — Click the consistency group name link to open its properties panel. Click **Rename** to change the Consistency Group Name..
- Search for a group — Type the complete or partial group name in the **Search** text box, and then press **Enter**.
- Add or remove virtual volumes from the group — Select a consistency group in the list, go to **VITRUAL VOLUMES** tab in the properties panel, and then click **Add** or **Remove**.

Opening the Distributed Consistency Groups view

From the GUI main menu, do one of the following:

1. In a metro node Local:

- a. Select **Provision Storage**.
 - b. In the **View by** drop-down, select **Consistency Groups**.
2. In a metro node Metro:
 - a. Select **Provision Storage**.
 - b. In the **Provision on** drop-down, select the **Distributed Storage**.
 - c. In the **View by** drop-down, select **Consistency Groups**.

Create Consistency Group wizard

Use the Create Consistency Group wizard to create a consistency group and add volumes to the group. You can create the group and add the volumes later. Adding or removing a volume from a consistency group provides more information. Note the following:

- You can create a maximum of 1,024 groups.
- There is a maximum of 1,000 volumes in a group.
- A volume can belong to only one consistency group at a time.

Configuring consistency groups

Consistency groups can consist of local, global, or distributed volumes. The following table describes the attributes that define each type of consistency group.

Consistency Group	Storage at Clusters	Visibility
Local (Local volumes only)	cluster-1 or cluster-2	cluster-1 or cluster-2
Global (Local volumes that are visible at both clusters)	cluster-1 or cluster-2	cluster-1 and cluster-2
Distributed (Distributed volumes only)	cluster-1 and cluster-2	cluster-1 and cluster-2

Choosing detach rules

When you add distributed volumes or volumes with global visibility to a consistency group, you must select a detach rule for the group. The rule you select is applied to all volumes in the group and determines I/O behavior during an inter-cluster link failure. The following table describes the predefined detach rules used for consistency groups, and the effect on I/O to volumes in the group during a failure. You can change the detach rule at any time in the consistency group properties dialog box.

Detach Rule	Description	Cluster-1	Cluster-2
Winner: cluster-x delay_time. Where delay_time = 5 seconds (default)	Cluster-1 detaches cluster-2 if a link failure lasts more than delay_time seconds. Note that the cluster that detaches becomes the winner cluster and services I/O; the cluster that gets detached becomes the loser cluster and suspends I/O.	Serves I/O	Suspends I/O
	Cluster-2 detaches cluster-1 if a link failure lasts more than delay_time seconds. Note that the cluster that detaches becomes the winner cluster and services I/O; the cluster that gets detached becomes the loser cluster and suspends I/O.	Suspends I/O	Serves I/O
No Automatic Winner	Suspends I/O on both clusters. This detach rule is generally used if you have created external scripts or use other programs to determine which cluster detaches and continues I/O during a link failure.	Suspends I/O	Suspends I/O

Types of consistency groups

Consistency groups can consist of local, global, or distributed volumes. The following table describes the attributes that define each type of consistency group.

Consistency Group Type	Storage at Clusters	Visibility
Local (Local volumes only)	cluster-1 or cluster-2	cluster-1 or cluster-2
Global (Local volumes that are visible at both clusters)	cluster-1 or cluster-2	cluster-1 and cluster-2
Distributed (Distributed volumes only)	cluster-1 and cluster-2	cluster-1 and cluster-2

Creating a consistency group

About this task

Use the Create Consistency Group wizard to create a consistency group and add volumes to the group. You can create the group and add the volumes later.

Steps

- From the GUI main menu, do one of the following:
 - In a metro node Local:
 - Select **Provision Storage**.
 - In the **View by** drop-down, select **Consistency Groups**.
 - In a metro node Metro:
 - Select **Provision Storage**.
 - In the **Provision on** drop-down, select the cluster of the consistency groups.
 - In the **View by** drop-down, select **Consistency Groups**.
- In the **Consistency Groups** view, click **CREATE**.
The Create Consistency Group wizard is displayed.
- Complete the wizard. For additional help using the wizard, click **Help**.

Adding a volume to a consistency group

About this task

Add a volume from a consistency group at any time. If the volume is exported in a storage view, it will remain in the storage view.

Be aware of the following when you add volumes to consistency groups:

- When you add a volume, the volume's visibility changes to the group's visibility. [Enabling or disabling remote access](#) provides information on changing a volume's visibility.
- When you add a distributed volume, the group's rule set is applied to the volume.

Steps

- From the GUI main menu, do one of the following:
 - In a metro node Local:
 - Select **Provision Storage**.
 - In the **View by** drop-down, select **Consistency Groups**.
 - In a metro node Metro:
 - Select **Provision Storage**.
 - In the **Provision on** drop-down, select the cluster of the consistency groups.
 - In the **View by** drop-down, select **Consistency Groups**.
- In the **Consistency Groups** view, select the consistency group.
The virtual volumes in the consistency group (if any) displays in the properties panel on the right.
- To add the virtual volumes, Click **Add**.
The **Add Virtual Volumes** dialog box is displayed. The virtual volumes already in the consistency group are listed in the Selected Virtual Volumes list.
- Select one or more virtual volumes from the Available Virtual Volumes list, and then click **Add** to add them to the Selected Virtual Volumes list.

5. Click **OK**.
The results dialog box opens.
6. Review the results, and then click **Close**.

Removing a volume from a consistency group

About this task

Remove a volume from a consistency group at any time. If the volume is exported in a storage view, it will remain in the storage view.

Be aware of the following when you remove volumes from consistency groups:

- When you remove a volume, it retains the consistency group's visibility. [Enabling or disabling remote access](#) provides information on changing a volume's visibility.
- When you remove a distributed volume, it retains the rule set applied to the consistency group.

Steps

1. From the GUI main menu, do one of the following:
 - a. In a metro node Local:
 - i. Select **Provision Storage**.
 - ii. In the **View by** drop-down, select **Consistency Groups**.
 - b. In a metro node Metro:
 - i. Select **Provision Storage**.
 - ii. In the **Provision on** drop-down, select the cluster of the consistency groups.
 - iii. In the **View by** drop-down, select **Consistency Groups**.
2. In the **Consistency Groups** view, select the consistency group.
The virtual volumes in the consistency group (if any) displays in the properties panel on the right.
3. Select Virtual Volume, and then click **Remove**.
A confirmation dialog box appears.
4. Review the volume to remove, and then click **Remove**.
The results dialog box opens.
5. Review the results, and then click **Close**.

Deleting a consistency group

About this task

Deleting a consistency group deletes the group, not the volumes in the group. If the volumes in the group are currently in a storage view, they will remain in the storage view.

Steps

1. From the GUI main menu, do one of the following:
 - a. In a metro node Local:
 - i. Select **Provision Storage**.
 - ii. In the **View by** drop-down, select **Consistency Groups**.
 - b. In a metro node Metro:
 - i. Select **Provision Storage**.
 - ii. In the **Provision on** drop-down, select the cluster of the consistency groups.
 - iii. In the **View by** drop-down, select **Consistency Groups**.
2. In the **Consistency Groups** view, select one or more consistency groups, and then click **Delete**.
3. Review the consistency group to delete, and then click **OK**.
4. Review the results of the operation, and then click **Close**.

Consistency Group status

The Operational status column in the **Consistency Groups** view shows the overall status of the consistency group at both clusters. If the status is the same at both clusters, that status displays in the Operational status column. If the status is not the same at both clusters, the less than optimal status will display.

The following table lists and defines the consistency group operational states.

Operational status	Definition
OK	The consistency group is servicing I/O normally at both clusters.
Suspended	I/O is suspended on the volumes in the group at one or both clusters.
Degraded	I/O is continuing to the volumes, but there are problems at one or both clusters.
Unknown	The status is unknown, likely because management connectivity is lost at one or both clusters.

The following table lists and describes the status details that may appear when you click one of the status links in the previous table.

Status detail	Description
requires-resolve-conflicting-detach	After the inter-cluster link is restored, two clusters have discovered that they have detached one another and resumed I/O independently. The clusters are continuing to service I/O on their independent versions of the data. You must issue the consistency-group resolve-conflicting-detach command in the CLI to make the view of data consistent again at the clusters.
rebuilding-across-clusters	One or more distributed member volumes is being rebuilt.
rebuilding-within-cluster	One or more local rebuilds is in progress at this cluster.
requires-resume-after-data-loss-failure	There have been at least two concurrent failures, and data has been lost. This can happen when, for example, a director fails shortly after the inter-cluster link fails, or when two directors fail at almost the same time. You must issue the consistency-group resume-after data-loss command in the CLI to select a winning cluster and allow I/O to resume.
requires-resume-after-rollback	A cluster has detached its peer cluster and rolled back the view of data, but is waiting for you to issue the consistency-group resume-after-rollback command in the CLI before resuming I/O.
cluster-departure	Not all the visible clusters are in communication.
requires-resume-at-loser	After the inter-cluster link is restored, a cluster that suspended I/O during the outage has discovered that its peer was declared the winner and resumed I/O. You must issue the consistency-group resume-at-loser command to make the view of the data consistent with the winner cluster, and to resume I/O at the loser cluster.
unhealthy-devices	I/O has stopped in this consistency group because one or more volumes is unhealthy and cannot perform I/O.
will-rollback-on-link-down	If a link failure occurs, the winner cluster will roll back the view of the data in order to resume I/O. This status detail appears when the static detach rule configured in detach-rule makes one of the clusters in active-clusters a loser during a link failure.

Consistency group properties

Consistency Group Name	The name of the consistency group. You can change the name in the properties dialog box.
Detach Rule (metro node Metro only)	<p>The detach rule that applies to the consistency group. It designates which cluster detaches if the clusters lose connectivity. Possible values:</p> <p>no-automatic-winner — The consistency group does not select a winning cluster.</p> <p>winner <cluster-name> <delay> — The cluster specified by <cluster-name> is declared the winner after the inter-cluster link outage lasts more than the number of seconds specified by delay. If a consistency group has a detach-rule configured, the rule applies to all volumes in the consistency group, and overrides any rule-sets applied to individual volumes.</p>
Visibility	Lists the clusters at which the consistency group is visible.
Storage At Clusters	Indicates if the volumes in the consistency group use storage from one cluster, or both clusters.

Step 1: Select or create a consistency group for the virtual volume

Select or create a consistency group for the virtual volume.

To view the topology of a consistency group, select the group in the list. In the map that displays to the right, the Storage icon



and blue shading around a cluster indicate that the selected consistency group uses storage from that cluster.

- To add the virtual volume to an existing consistency group, select **Add storage to existing consistency group** and then select the consistency group in the list.
- To filter the list of consistency groups by name, type the full or partial name of the consistency group in the **Filter** text box. To filter by storage location, click the drop-down box to the right of the **Filter** text box, and then select the storage location.
- To create a new consistency group for the virtual volume, select **Create new consistency group**.


Step 1: Create a consistency group



Complete the following fields to create a consistency group for the virtual volume.

Consistency Group Name — Type a name for the consistency group.

Consistency Group Type — Select the type of consistency group to create: Local or Distributed.

- For local consistency groups, select the cluster on which to create the consistency group.
- For Distributed consistency groups, select a detach rule. The detach rule determines if IO continues and at which cluster if an inter-cluster link failure occurs.

 **NOTE:** Detach rules are applied per consistency group.

Detach rule	Cluster-1	Cluster-2
No automatic winner	Suspends I/O  NOTE: This detach rule is generally used if you have created external scripts or use other programs to determine which cluster detaches and continues I/O during a link failure.	Suspends I/O  NOTE: This detach rule is generally used if you have created external scripts or use other programs to determine which cluster detaches and continues I/O during a link failure.
Winner: <cluster-1> (5 seconds)	Serves I/O within 5 seconds	Suspends I/O
Winner: <cluster-2> (5 seconds)	Suspends I/O	Serves I/O within 5 seconds


Step 2: Select volume options

Select mirroring on the cluster where the virtual volumes will be created (optional), the quantity, size and base name for the virtual volumes, and then enter a name for the provisioning job (optional).

High Availability — Select this check box to enable mirroring at the cluster indicated. The option to mirror at a cluster is disabled if you have less than two arrays in the cluster.

- For a local consistency group, select the check box to enable mirroring at the local cluster.
- For a distributed consistency group, select the appropriate check box to enable mirroring at one cluster or both clusters, or leave both check boxes clear to disable mirroring.

Number of Volumes — The number of volumes to create. You can create a maximum of 10 volumes at a time. Each volume will have the same RAID geometry.

 **NOTE:** For Unity arrays, you can provision only one volume at a time.

Capacity Per Volume — The size of each volume. Select the size in MB, GB or TB. The minimum size is 150MB. Each volume will have the same capacity.

Total Capacity — The total capacity of all the virtual volumes to be created.

Volume Base Name — Type a base name for the virtual volumes. A number will be added to the base name for each volume created. For example, using the base name Test, volumes will be named Test_1, Test_2, Test_3, and so forth. You can reuse the base name to provision additional volumes. Metro node will continue to increment the number at the end of the name, beginning with the highest existing numeric suffix.

Job Name — Select the **Specify a job name** check box, and then type a name for the provisioning job to be created. This field is optional. If you do not specify a job name, the volume base name will be used. If the job name you enter is already in use, a numeric suffix will be added to the job name to make it unique, for example, *jobname_n*. Where n is the first available number that makes the name unique. Or, you can choose to use a different name. You can see currently used job names in the **Provision Jobs Status** page under **Monitoring**. The job name must meet the following criteria:

- Maximum of 64 characters
- Must begin with a letter
- Must include only numbers, letters, underscores, or dashes

To discard all selections made on the previous screen and restart the wizard from Step 1, click **Start Over** in the lower-right corner of the screen.

To go back to a specific step, click **Back** until you reach that step.

Step 3: Select a storage pool

Select a storage array, and then select a storage pool and storage group on the array for the volumes to be created.

V3GR arrays only: Selecting a policy/workload

For V3GR arrays, policies define a level of performance (based on an average response time) required by the storage workload. Refer to the array documentation for more information.

Note the following:


- If you select a policy, storage pools and storage groups with the same policy are available for selection.
- If you select **None**, pools with any type of policy, and workloads with no policy are available for selection.
- If you select **Optimized**, only system-optimized pools and storage groups with optimized policies are available for selection.
- Selecting a workload is optional, however, if you select a policy and a workload, the storage pools and storage groups with the same Policy+Workload are listed for selection.

Selecting a storage pool

The Storage Pool section shows the following information about the storage pools on the selected array:


- **Name** — The name of the pool on the array.
- **Policy** — The type of policy associated with the storage pool.

- **Type** — Indicates whether the pool is capable of provisioning thick volumes, thin volumes, or both. If you select a thin pool, a thin rebuild will be automatically set for the storage volume that is created during the provisioning process.
- **Free** — The amount of storage available for provisioning.
- **Total** — The total amount of storage available from the pool.

 **NOTE:** To create a thin volume, the selected pools must be thin capable and the supporting arrays must belong to the same array family.

Selecting a storage group

By default, metro node will automatically select a storage group for the newly created volumes, based on compatible policies, and the lowest amount of volumes in the group. However, if you want to select the storage group, click the **Storage Group** drop-down box, and then select `Select from list` to show the available storage groups.

 **NOTE:** For V3GR arrays only, the **Storage Group** drop-down box changes to a selection list because a storage group is required. For all other arrays, the storage group is optional.

The Storage Group section shows the following information about available storage groups:

- **Name** — The name of a storage group on the selected array.
- **Policy** — The type of policy associated with the storage pool (if any). metro node lists only storage groups with policies that match that of the selected pool, and storage groups with no associated policies.
- **Parent** — The name of the parent storage group (if applicable).
- **Masking views** — The name of the masking view associated with the storage group (if applicable).
- **# Volumes** — The number of volumes currently in the storage group.

To discard all selections made on previous screens and restart the wizard from Step 1, click **Start Over** in the lower-right corner of the screen.

To go back to a specific step, click **Back** until you reach that step.

Step 3: Select a pool for each mirror on the second cluster

Select a storage pool on the second cluster to provide storage for the virtual volumes to be created.

V3GR arrays only: Selecting a policy or workload

For V3GR arrays, policies define a level of performance (based on an average response time) required by the storage workload. Refer to the array documentation for more information.


Note the following:

- When you select a policy, storage pools and storage groups with the same policy are available for selection.
- If you select `None`, pools with any type of policy, and workloads with no policy are listed for selected for selection.
- If you select `Optimized`, only system-optimized pools and storage groups with optimized policies are available for selection.
- Selecting a workload is optional, however, if you select a policy and a workload, the storage pools and storage groups with the same Policy+Workload are listed for selection.

Selecting a storage pool


The Storage Pool section shows the following information about the storage pools on the selected array:

- **Name** — The name of the pool on the array.
- **Policy** — The type of policy associated with the storage pool.
- **Type** — Indicates whether the pool is capable of provisioning thick volumes, thin volumes, or both. If you select a thin pool, a thin rebuild will be automatically set for the storage volume that is created during the provisioning process.
- **Free** — The amount of storage available for provisioning from the pool.
- **Total** — The total amount of storage available from the pool.

 **NOTE:** To create a thin volume, the selected pools must be thin capable and the supporting arrays must belong to the same family of arrays.

Selecting a storage group

By default, metro node will automatically select a storage group for the newly created volumes, based on compatible policies, and the lowest amount of volumes in the group. However, if you want to select a storage group, click the **Storage Group** drop-down box, and then select `Select from list` to show the available storage groups to choose from.

 **NOTE:** For V3GR arrays only, the **Storage Group** drop-down box changes to a selection list because a storage group is required. For all other arrays, the storage group is optional.

The Storage Group section shows the following information about available storage groups:

- **Name** — The name of a storage group on the selected array.
- **Policy** — The type of policy associated with the storage pool (if any). Metro node lists only storage groups with policies that match that of the selected pool, and storage groups with no associated policies.
- **Parent** — The name of the parent storage group (if applicable).
- **Masking views** — The name of the masking view associated with the storage group (if applicable).
- **# Volumes** — The number of volumes currently in the storage group.

To discard all selections made on previous screens and restart the wizard from Step 1, click **Start Over** in the lower-right corner of the screen.

To go back to a specific step, click **Back** until you reach that step.

Step 3: Select a pool for each mirror in the cluster

Select an array to provide storage for each mirror in the cluster, and then select a pool from the list of available pools on the array. To ensure high availability, the array selected for one mirror is automatically filtered from the list of arrays available for other mirror.

V3GR arrays only: Selecting a policy or workload

For V3GR arrays, policies define a level of performance (based on an average response time) required by the storage workload. Refer to your array documentation for more information.


Note the following:

- When you select a policy, storage pools and storage groups with the same policy are available for selection.
- If you select `None`, pools with any type of policy, and workloads with no policy are listed for selected for selection.
- If you select `Optimized`, only system-optimized pools and storage groups with optimized policies are available for selection.
- Selecting a workload is optional, however, if you select a policy and a workload, the storage pools and storage groups with the same Policy+Workload are listed for selection.

Selecting a storage pool


The Storage Pool section shows the following information about the storage pools on the selected array:

- **Name** — The name of the pool on the array.
- **Policy** — The type of policy associated with the storage pool.
- **Type** — Indicates whether the pool is capable of provisioning thick volumes, thin volumes, or both. If you select a thin pool, a thin rebuild will be automatically set for the storage volume that is created during the provisioning process.
- **Free** — The amount of storage available for provisioning from the pool.
- **Total** — The total amount of storage available from the pool.

 **NOTE:** To create thin volumes, the selected pools must be thin capable and the supporting arrays must belong to the same array family.

Selecting a storage group

By default, metro node will automatically select a storage group for the newly created volumes, based on compatible policies, and the lowest amount of volumes in the group. However, if you want to select a storage group, click the **Storage Group** drop-down box, and then select `Select from list` to show the available storage groups to choose from.

 **NOTE:** For V3GR arrays only, the **Storage Group** drop-down box changes to a selection list because a storage group is required. For all other arrays, the storage group is optional.

The Storage Group section shows the following information about available storage groups:


- **Name** — The name of a storage group on the selected array.
- **Policy** — The type of policy associated with the storage pool (if any). Metro node lists only storage groups with policies that match that of the selected pool, and storage groups with no associated policies.
- **Parent** — The name of the parent storage group (if applicable).
- **Masking views** — The name of the masking view associated with the storage group (if applicable).
- **# Volumes** — The number of volumes currently in the storage group.

To discard all selections made on previous screens and restart the wizard from Step 1, click **Start Over** in the lower-right corner of the screen.

To go back to a specific step, click **Back** until you reach that step.

Step3: Create thin virtual volumes

The pools you have selected support creating thin virtual volumes. To create thin volumes, Select the **Create thin virtual volumes** check box, otherwise, click **Next** to continue without creating thin volumes.

 **NOTE:** To create thin virtual volumes, the selected storage arrays must belong to the same array family, and pools must support creating thin volumes. For Unity and XtremIO arrays, the storage volumes are always thin.

Select a storage view for the virtual volume(s) (optional)

Select a storage view for the virtual volume(s), or select **None** to continue and select a storage view later.

- If you select an active storage view, the virtual volume is added to that storage view and automatically becomes visible to hosts with access to the storage view. In some cases a rescan of the host may be required. If the storage view is inactive, go to the **Storage Views** screen and add the components necessary to activate the storage view. [Creating a storage view](#) provides more information.
- If you select **None**, when you are ready to expose the virtual volume to hosts, you must manually add the virtual volume to a storage view. [Adding or removing virtual volumes from a storage view](#) provides more information.

To discard all selections made on previous screens and restart the wizard from the Step 1, click **Start Over** in the lower-right corner of the screen.

To go back to a specific step, click **Back** until you reach that step.

Step 5: Review your selections

This screen shows a summary of your selections in the wizard and a topology map of the volume to be created.

Review your selections, and if satisfied, click **Provision Storage** in the lower-right corner of the screen to submit the provisioning request.

To make changes to a specific step, click **Back** until you reach that step.

To discard all selections made on previous screens and restart the wizard from the Step 1, click **Start Over** in the lower-right corner of the screen.

Step 6: View results

This screen shows the results of the provisioning request.

- To manage the provisioning job, click the **View job Status** link to open the **Provisioning Jobs Status** screen.
- To provision more storage, click **Provision More** in the lower-right corner of the screen to go back to Step 1 in the wizard, otherwise, click **Done** to close the wizard.

Step 2: Select volume options

Select high availability options, synchronization options, and then enter a name for the volume. If a cluster has only one array, the ability to mirror is disabled for that cluster.

- **Source Cluster** — Select the cluster where you want to create the virtual volume.
- **High Availability** — Select a check box to enable mirroring at the cluster indicated. You can select a single cluster, both clusters, or leave both check boxes unselected to disable mirroring at both clusters. When you select the check box to enable mirroring at a cluster, `Local Mirroring` appears above the cluster in the map.
- **Synchronization Options** — Select the appropriate option to indicate whether or not you want to synchronize data from the source storage volume to the target storage volume.
- **Volume Name** — Type a name for the virtual volume.

To discard all selections made on the previous screen and restart the wizard from Step 1, click **Start Over** in the lower-right corner of the screen.

To go back to a specific step, click **Back** until you reach that step.

Step 2: Select volume options

Select high availability options and a name for the virtual volume. The mirroring option is enabled only for the cluster where the virtual volume will be created. If the cluster has only one array, the option to mirror is disabled.

Source Cluster — The cluster where the virtual volume will be created.

High Availability — Select the check box to mirror the virtual volume at the cluster indicated.

Volume Name — Type a name for the virtual volume.


To discard all selections made on the previous screen and restart the wizard from the Step 1, click **Start Over** in the lower-right corner of the screen.

To go back to a specific step, click **Back** until you reach that step.

Step 3: Select a storage volume to create the virtual volume

Select one storage volume to use for the virtual volume. The virtual volume will use the entire capacity of the storage volume. You can select claimed or unclaimed storage volumes.

Select an array in the drop-down box, and then select a storage volume from the list of available storage volumes on the array.

 **NOTE:** To create a thin volume, make sure that the storage volume is thin capable.


To discard all selections made on previous screens and restart the wizard from Step 1, click **Start Over** in the lower-right corner of the screen.

To go back to a specific step, click **Back** until you reach that step.

Step 3: Select a source and target storage volume

Select a source and a target storage volume on the cluster where the virtual volume will be created. The source storage is used to create the virtual volume and the target storage is used to mirror the virtual volume. The virtual volume to be created will use the entire capacity of the source storage volume. To ensure high availability, the array selected to provide the source storage is automatically filtered from the list of arrays available for the target storage.

Select only one source and one target storage volume. The storage volume can be claimed or unclaimed.

 **NOTE:** To create a thin volume, the selected source and target storage volumes must come from the same array family, and must also be thin capable.

Source Storage — Select an array in the drop-down box, and then select from the list of available storage volumes on the array.

Target Storage — Select an array in the drop-down box, and then select a storage volume from the list of available storage volumes on the array.

To discard all selections made on previous screens and restart the wizard from Step 1, click **Start Over** in the lower-right corner of the screen.

To go back to a specific step, click **Back** until you reach that step.

Step 3: Create thin volumes

The storage volumes you selected support creating thin virtual volumes. To create thin virtual volumes, click the **Create thin virtual volumes** check box, or click **Next** to continue without creating thin volumes.

Step 3: Select a target storage volume on the remote cluster

Select a target storage volume on the remote cluster. Data written to the source storage volume will be copied onto this storage volume.

Note the following:

- You can select claimed or unclaimed storage volumes.
- The size the target volume must be greater than or equal to the size of the source volume.


Select an array in the drop-down box, and then select from the list of available storage volumes on the array.

To discard all selections made on previous screens and restart the wizard from the Step 1, click **Start Over** in the lower-right corner of the screen.

To go back to a specific step, click **Back** until you reach that step.

Step 3: Select target storage on the remote cluster

When mirroring on both clusters, select target storage volumes on the remote cluster. The data on the source storage volume will be mirrored to the target storage volumes on the remote cluster. To ensure high availability, the array selected for one target is automatically filtered from the list of arrays available for the other target.

 **NOTE:** The size of the target volume must be greater than or equal to the size of the source volume. You can select claimed or unclaimed storage.

To discard all selections made on previous screens and restart the wizard from the Step 1, click **Start Over** in the lower-right corner of the screen.

To go back to a specific step, click **Back** until you reach that step.

Step 6: View results

This screen shows the results of the provisioning request.

To provision more storage, click **Provision More** in the lower-right corner of the screen to go back to Step 1 in the wizard.

To close the wizard, click **Done**.

Show Logical Units

Name (VPD ID)	The VPD identifier for the logical unit.
Storage Volume	The storage volume to which this logical unit belongs.
LUNs	The LUN number on the array.
Connectivity	The connectivity status of the logical unit.
Visibility	Indicates how the logical unit is visible to the controllers. See possible values .
ALUA Support	Indicates if Asymmetrical Logical Unit Access (ALUA) is supported on the array. Depending on the type of controller exposed to the LUN, one of these values will display.

Exporting storage

This section describes how to export storage.

Topics:

- [Initiators and metro node ports](#)
- [Storage views](#)

Initiators and metro node ports

About Initiators and metro node ports

Metro node ports

The virtual volumes that are created on a device are not visible to hosts until you export them. Virtual volumes are exported to a host through front-end ports on the metro node directors and HBA ports on the host/server. For failover purposes, two or more front-end ports can be used to export the same volumes. Typically, a storage view has two ports assigned to it from two different metro node directors to provide maximum redundancy. When volumes are added to a view, they are exported on all metro node ports in the view using the same LUN numbers.

Initiators

For an initiator to see the virtual volumes in a storage view, it must be registered and included in the storage view's registered initiator list. The initiator must also be able to communicate with the front-end ports over Fibre Channel connections—either through a direct connection or through a fabric.

A volume is exported to an initiator as a LUN on one or more front-end port WWNs. Typically, initiators are grouped so that all initiators in a group share the same view of the exported storage (they can see the same volumes by the same LUN numbers on the same WWNs).

Using the Ports view

The **Ports** view lists all metro node front-end ports used in a cluster. By default, ports are sorted by name, in ascending order. The arrow in the default sort column indicates the direction of the sort; ascending (



) or descending (



). You can [sort](#) columns of data, [select](#) one or more objects in the list, and [show or hide](#) on the screen.

This screen also shows the following information about ports.

Column	Description
Port	Display the details of port.
Port name	The name of the metro node port.
Director	Indicates the name of director which is linked.
Enabled	Indicates whether the port is enabled or not.
Export Status	Indicates how the port is functioning. Port status provides more information.

Column	Description
Port WWN	A unique network identifier for the port.



Port Properties Panel

This screen displays the properties of the ports in a separate panel at the extreme right side of the screen. Using the **Port Properties** panel, you can:

- View the details of the port
- Enable or disable a port
- View the details of the discovered initiators of the port

Tasks in this Screen

You can perform the following tasks from the **Ports** view:

- Search for a port — Click the  icon corresponding to the element and type the full or the partial name in the field that appears. Press **Enter**.
- Refresh screen data — Click the Refresh icon .
- Enable a port — Select a disabled port. In the **Port Properties** panel, click **ENABLE PORT**.
- Disable a port — Select an enabled port. In the **Port Properties** panel, click **DISABLE PORT**.
- Show the discovered initiators of the port — In the **Port Properties** panel, click the **INITIATORS** tab.

Opening the Ports view

From the GUI main menu, do one of the following:

1. In a metro node Local:
 - a. Select **Provision Storage**.
 - b. In the **View by** drop-down, select **Ports**.
2. In a metro node Metro:
 - a. Select **Provision Storage**.
 - b. In the **Provision on** drop-down, select the cluster of the port.
 - c. In the **View by** drop-down, select **Ports**.

Using the Initiators view

The **Initiators** view lists all initiators (discovered and registered) in a cluster. These initiators are discovered from the fabric name server and are not necessarily logged into the target ports. By default, initiators are sorted by name, in ascending alphabetical order. The arrow in the default sort column indicates the direction of the sort; ascending (

↑

) or descending (

↓

). You can [sort](#) columns of data, [select](#) one or more objects in the list, and [show or hide](#) on the screen.

In the main panel of the screen, you can perform these tasks:

- Register a new initiator manually.
- Register an initiator.
- Unregister a registered initiator.
- Rediscover an initiator that does not display properly.

This screen shows the following information about initiators:

Column	Description
Initiator	The name that is given to a registered initiator, or the default name discovered by metro node. All unregistered initiator names are prefixed

Column	Description
	with "UNREGISTERED-". You can rename the initiator during the registration process.
Host Type	The type of host associated with the initiator.
Node WWN	A unique network identifier for the HBA's interface card. All ports on the card share this node World Wide Name (WWN).
Port WWN	A unique network identifier for the port.




Initiator Properties Panel

This screen displays the properties of the storage volumes in a separate panel at the extreme right side of the screen. Using the **Initiator Properties** panel, you can:

- View the details of the initiator.
- Rename the initiator.
- View the details of the target ports.

Tasks in this Screen

You can also perform the following tasks from the **Initiators** view:

- Search for an extent—Click the  icon corresponding to the element and type the name in the field that appears.
- Refresh screen data—Click the Refresh icon .
- Register a new initiator manually - Click **REGISTER INITIATOR MANUALLY** and type the following details in the pop-up screen:
 - New initiator name—The name of the new initiator.
 - Host Type—The type of the host.
 - Port WWN—The WWN of the initiator port.
 - Node WWN—The WWN of the initiator node.
- Register an initiator—Select the initiator, and click **MORE > Register**.
- Unregister an initiator—Select a registered initiator and click **MORE > Unregister**.
- Rediscover an initiator—Select an initiator, and click **MORE > Rediscover initiator**.
- Rename an initiator—Click the initiator, and then click **Rename** under the **INITIATOR** tab in the **Initiator Properties** pane. Click  to save the changes.
- View the status and properties of the initiator and the discovered target ports—Select the initiator to display its properties and the target ports in the **Initiator Properties** panel. To view the properties of the target port, click the port name under the **TARGET PORTS** tab.

Opening the Initiators view

From the GUI main menu, do one of the following:

1. In a metro node Local:
 - a. Select **Provision Storage**.
 - b. In the **View by** drop-down, select **Initiators**.
2. In a metro node Metro:
 - a. Select **Provision Storage**.
 - b. In the **Provision on** drop-down, select the cluster of the port.
 - c. In the **View by** drop-down, select **Initiators**.

Enabling or disabling a port

Steps

1. From the GUI main menu, do one of the following:
 - In a metro node Local:
 - a. Select **Provision Storage**.
 - b. In the **View by** drop-down, select **Ports**.
 - In a metro node Metro:
 - a. Select **Provision Storage**.
 - b. In the **Provision on** drop-down, select the cluster of the port.
 - c. In the **View by** drop-down, select **Ports**.
2. In the **Ports** view, select the port.
3. In the **Port Properties** pane, click **ENABLE PORT** to enable the port, or **DISABLE PORT** to disable the port. The change is reflected in the **Enabled** column in the **Ports** view.

Registering an initiator

About this task

When initiators are connected directly or through a Fibre Channel fabric, metro node automatically discovers them and populates the Initiators screen. Once discovered, you must register the initiators with metro node before they can be added to a storage view and access storage. Registering an initiator gives a meaningful name to the Port World Wide Name (WWN) (typically the server's DNS name) to allow you to identify the host. All discovered initiator names are prefixed with UNREGISTERED-(example, UNREGISTERED-0x200000cc05bb004d) . If metro node did not automatically discover an initiator, you must register the initiator manually.

Ensure that you specify the correct host type in the Host Type column as this attribute cannot be changed in the **Initiator Properties** dialog box after the registration is complete. To change the host type after registration, you must use the CLI. The *CLI Guide for metro node* provides more information. [Supported host types](#) lists the host types that are supported with metro node.



NOTE: You must register each port of a host's HBA/HBC as a separate initiator.

Steps

1. From the GUI main menu, do one of the following:
 - a. In a metro node Local:
 - i. Select **Provision Storage**.
 - ii. In the **View by** drop-down, select **Initiators**.
 - b. In a metro node Metro:
 - i. Select **Provision Storage**.
 - ii. In the **Provision on** drop-down, select the cluster of the port.
 - iii. In the **View by** drop-down, select **Initiators**.
2. In the **Initiators** view, select the unregistered initiator, and then click **Register**.
By default, the **Register** button is disabled if you select a registered initiator.
3. In the **Register Initiators** view, the selected initiator displays with the REGISTERED_ prefix.
4. Type a new meaningful name for the initiator, or make no changes to accept the name as it is displayed.
5. Select a host type, and click **OK**.
A message is displayed to indicate the success of the operation.
6. Click **Close**.

Supported host types

Metro node supports the following types of hosts:

- SUN-VCS
- AIX

- HPUX
- IBM-D910
- Default (use this type for all hosts not listed above)

Rediscovering an initiator

About this task

When initiators are connected directly or through a Fibre Channel fabric, metro node automatically discovers them and populates the **Initiators** view.


Steps

1. From the GUI main menu, do one of the following:
 - a. In a metro node Local:
 - i. Select **Provision Storage**.
 - ii. In the **View by** drop-down, select **Initiators**.
 - b. In a metro node Metro:
 - i. Select **Provision Storage**.
 - ii. In the **Provision on** drop-down, select the cluster of the port.
 - iii. In the **View by** drop-down, select **Initiators**.
2. In the **Initiators** view, select **MORE > Rediscover Initiators**.
3. Click **OK** to continue rediscovering initiators. This process might take some time.
4. Do one of the following:
 - If the rediscovery completes with no errors, a green check mark displays in the **Result** column. Click **Close**.
 - If an error occurs, a warning symbol displays in the **Result** column. Click **Close** and retry the rediscovery.

Manually registering an initiator

About this task

Metro node automatically discovers all connected initiators and populates the **Initiators** view. If, for some reason metro node did not discover the initiator, you must manually register the initiator. Make sure that you specify the correct host type in the **Host Type** column as this attribute cannot be changed once the registration is complete. To change the host type after registration, you must use the CLI to unregister the initiator, and then register it again using the correct host type. [Supported host types](#) lists the host types supported with metro node.

 **NOTE:** You must register each port of a host's HBA/HCA as a separate initiator.

Steps

1. From the GUI main menu, do one of the following:
 - a. In a metro node Local:
 - i. Select **Provision Storage**.
 - ii. In the **View by** drop-down, select **Ports**.
 - b. In a metro node Metro:
 - i. Select **Provision Storage**.
 - ii. In the **Provision on** drop-down, select the cluster of the port.
 - iii. In the **View by** drop-down, select **Ports**.
2. In the **Initiators** view, click **Register Initiator Manually**.
The **Register Initiators Manually** dialog box is displayed, type a meaningful name for the initiator.
3. Select the host type.
4. Type the Port WWN.
Ensure that you type a valid Port WWN. You can cut and paste into this field to avoid errors.
5. Type the Node WWN.
6. Click **OK**.
A message displays to indicate the result of the operation.

- Click **Close**.
The initiator is now visible in the **Initiators** view.

Unregistering an initiator

About this task

You can unregister an initiator that is not included in any storage views. If an initiator is included in a storage view, **Unregister** is disabled.

Steps

- From the GUI main menu, do one of the following:
 - In a metro node Local:
 - Select **Provision Storage**.
 - In the **View by** drop-down, select **Initiators**.
 - In a metro node Metro:
 - Select **Provision Storage**.
 - In the **Provision on** drop-down, select the cluster of the port.
 - In the **View by** drop-down, select **Initiators**.
- In the **Initiators** view, select the initiators to unregister.
A confirmation dialog box is displayed.
- Review the initiators to unregister, and if satisfied, click **OK**.
The results dialog box is displayed.
- Review the results of the operation, and then click **Close**.
You can sort the results in this dialog box.

Metro node port status

The following table defines the port status.

Status	Definition
OK	The port is functioning normally.
Pending	The port is enabled, but is not communicating with any components.
Suspended	The port is down or disconnected.
Hardware error	The port is experiencing hardware problems.

Initiator properties

Initiator Name	The name given to the initiator. You can change the name in this dialog box.
Host Type	Identifies the host operating system. To change the host type after the initiator has been registered, use the CLI.
Node WWN	A unique network identifier for the HBA's interface card. All ports on the card share this World Wide Name (WWN).
Port WWN	A unique network identifier for the port.
Target Port	The list of metro node ports visible to the initiator.

Add capacity to virtual volumes

Use this dialog box to expand a virtual volume using concatenation.

Add Capacity to Virtual Volume	The selected virtual volume that you want to expand.
Current Capacity	The size of the selected virtual volume.
Additional Capacity Needed	Enter the amount of space to add to this virtual volume. Metro node will look for volumes with at least that available capacity. To enter the value using a different unit, select the correct unit from the drop-down list. When you enter the amount of space needed, you can click Show Available to list the storage volumes with a minimum capacity equal to the additional capacity needed. If you change the capacity needed, click Show Available again to update Available Storage Volumes list.
Total New Capacity	The new size of the virtual volume based on the additional capacity needed.
Available storage volumes	The list of available storage volumes to choose from based on the additional capacity needed. The name of the storage volume and the largest free chunk available in the storage volume are displayed.

Storage views

About storage views

A storage view defines which hosts access which virtual volumes on which metro node ports. It consists of the following components:

- Registered initiators (hosts) to access the storage
- metro node ports (front-end) to export the storage
- One or more virtual volumes to export (with their respective LUN numbers)

Typically, one storage view is created for all hosts that require access to the same storage. The process of exporting storage requires the following steps:

1. Create a storage view
2. Add registered initiators
3. Add metro node ports
4. Add virtual volumes

Initiators

For an initiator (host) to see the virtual volumes in a storage view, it must be registered and included in the storage view's registered initiator list. The initiator must also be able to communicate with the front-end ports over Fibre Channel connections—either through a direct connection or through a fabric.

Metro node ports

Typically, a storage view has one metro node port assigned to it from each director in the cluster to provide maximum redundancy. When volumes are added to a storage view, they are exported on all metro node ports in the view, using the same LUN numbers.

Virtual volumes

The virtual volumes added to a storage view can be created on devices, distributed devices, or can be a member of a consistency group. The Virtual Volumes view contains a list of all volumes in the cluster, including the volume's Service status, which indicates whether or not the virtual volume is exported to a host.

Storage view status

Until all three components (registered initiators, metro node ports, and virtual volumes) are added to a storage view, the view is Inactive. When all three components are added to a storage view, it automatically becomes active, and hosts can access the virtual volumes in the storage view.

Storage view maps

You can view a storage hierarchy map to see the components in the storage view, and drill down further to see the health and properties of all underlying storage objects. [Storage view map](#) provides more information.

Using the Storage Views screen

The **Storage Views** screen lists all storage views in a selected cluster. Use this screen to search for a storage view, assess the health of a storage view, or show associated initiators, ports, and virtual volumes. By default, storage views are sorted by name, in ascending alphabetical order. The arrow in the default sort column indicates the direction of the sort; ascending (



) or descending (



). You can [sort](#) columns of data, [select](#) one or more objects in the list, and [show or hide](#) on the screen.

This screen shows the following information about storage views:

Column	Description
Storage View	The name of the storage view.
Operational Status	Indicates how the storage view is functioning in the cluster. Storage view status provides more information.
Active	Indicates the status of the storage view: Active or Inactive.
Initiators	The number of initiators using this storage view.
Volumes	The number of virtual volumes in this storage view.
Ports	The number of ports connecting this view.

Using the Storage View Properties panel




This screen displays the Storage View properties in a separate panel at the right side of the screen. The panel shows the details of storage view, status of all initiators, metro node ports, and virtual volumes in the selected storage view.

Use the following tabs to view the status and properties of the associated components:

- **Storage View**—Shows the details and status of storage view.
- **Initiators**—Shows the status of all initiators in the selected storage view. Click an initiator to see its properties.
- **Ports**—Shows the status of all ports that are associated with the storage view. Click a port to see its properties.
- **Virtual volumes**—Shows the status of all virtual volumes in the selected storage view. Click a virtual volume to see its properties. Also, you can select the LUN, VPD ID, Capacity (or all three check boxes) to show the associated information for each virtual volume listed in the panel.

Additional tasks

You can also perform the following tasks from the **Storage Views** screen:

- Search for a storage view—Click the Filter icon  corresponding to the element and type the full or the partial name in the field that appears. Press **Enter**.
- Refresh screen data—Click the Refresh icon .
- Rename a storage view—Select the storage view, and then click **Rename** under the **STORAGE VIEW** tab in the **Storage View Properties** panel. Click  to save the changes.
- View properties for initiators, ports, and virtual volumes in a storage view—Select the storage view name to populate the panel on the right. Click a tab to display a component, and then click the required tab to view its properties.
- Add or remove initiators—Select the storage view name to display its components in the panel on the right. In the panel, click **Initiators** and then click **ADD/REMOVE INITIATORS**.
- Add or remove ports—Select the storage view name to display its components in the panel on the right. In the panel, click **Ports**, and then click **ADD/REMOVE PORTS**.

- Add or remove virtual volumes—Select the storage view name to display its components in the panel on the right. In the panel, click **Virtual Volumes**, and then click **Add** or **Remove**.

Opening the Storage Views screen

From the GUI main menu, do one of the following:

1. In a metro node Local:
 - a. Select **Provision Storage**.
 - b. In the **View by** drop-down, select **Storage Views**.
2. In a metro node Metro:
 - a. Select **Provision Storage**.
 - b. In the **Provision on** drop-down, select the cluster of the storage view.
 - c. In the **View by** drop-down, select **Storage Views**.

The Create Storage View wizard

Use the Create Storage View wizard to create a storage view.

Best Practices

- Create a separate storage view for each host that needs access to the same storage, and then add the volumes to each view.
- For redundancy, at least two initiators from a host must be added to the storage view.

Storage view status

To create a storage view, you must give the storage view a name and add at least one metro node port. You can return to the wizard at any time to add the other components. Until all components are added, the storage view is inactive. It will display in the **Storage Views** screen, but hosts cannot access the storage in the view because it is incomplete.

Selecting storage view components

When selecting initiators in the Create Storage View wizard, if a selected initiator has ports associated with it, those ports are automatically selected for you. You also have the option of manually selecting ports. When selecting virtual volumes, you can view the consistency group to which the volume belongs and sort the virtual volumes by consistency group.

Manually setting LUN numbers

By default, the system automatically assigns a LUN number to each selected virtual volume (0 - 16383). When creating a storage view, you can manually assign a LUN number if desired, or automatically assign LUN numbers by specifying a range of numbers.

Creating a storage view

About this task

To export storage, you must first create a storage view, and then add metro node front-end ports and virtual volumes to the view. Virtual volumes are not visible to hosts until they are in a storage view with associated ports and initiators. Use the Create Storage View wizard to create a storage view, and add initiators, ports and virtual volumes to the view. When you create a storage view with all required components, it automatically becomes active. When a storage view is active, hosts can see the storage and begin I/O to the virtual volumes in the storage view. [The Create Storage View wizard](#) provides more information on best practices and other information to consider when creating a storage view.

Steps

1. From the GUI main menu:
 - a. In a metro node Local:
 - i. Select **Provision Storage**.
 - ii. In the **View by** drop-down, select **Storage Views**.
 - b. In a metro node Metro:
 - i. Select **Provision Storage**.
 - ii. In the **Provision on** drop-down, select the cluster of the storage view.
 - iii. In the **View by** drop-down, select **Storage Views**.
2. In the **Storage Views** screen, click **CREATE**.
The Create Storage Views wizard is displayed.
3. Complete the wizard. For additional help using the wizard, click **Help**.

Deleting a storage view

About this task

You can select one or more storage views to delete. If the storage view is active, a warning appears, and you must confirm before proceeding.

Steps

1. From the GUI main menu, do one of the following:
 - a. In a metro node Local:
 - i. Select **Provision Storage**.
 - ii. In the **View by** drop-down, select **Storage Views**.
 - b. In a metro node Metro:
 - i. Select **Provision Storage**.
 - ii. In the **Provision on** drop-down, select the cluster of the storage view.
 - iii. In the **View by** drop-down, select **Storage Views**.
2. In the **Storage Views** screen, select the storage views to delete, and then click **DELETE**.
A confirmation dialog box is displayed. If the storage view is active, a warning appears.
3. Review the storage views to delete, and then click **OK**.
The results dialog box is displayed.
4. Review the results, and then click **Close**.
You can sort the results in the dialog box.


Renaming a storage view

About this task

You can rename a storage view at any time.

Steps


1. From the GUI main menu, do one of the following:
 - a. In a metro node Local:
 - i. Select **Provision Storage**.
 - ii. In the **View by** drop-down, select **Storage Views**.
 - b. In a metro node Metro:
 - i. Select **Provision Storage**.
 - ii. In the **Provision on** drop-down, select the cluster of the storage view.
 - iii. In the **View by** drop-down, select **Storage Views**.
2. In the **Storage Views** screen, click the storage view name link.
The **Storage View Properties** dialog box is displayed.

3. Type the new name in the **Storage View Name** field, and then click  to save the changes.
4. Verify that the name has been changed in the **Storage Views** screen.

Adding or removing initiators from a storage view


About this task

Add an initiator to an existing storage view to allow the associated host to access the virtual volumes in the view. Remove an initiator when the host no longer needs to access the virtual volumes in the storage view. If a storage view is inactive, you can add or remove initiators with no effect on the system.

 **CAUTION:** If a storage view is active, removing an initiator can cause the data on all virtual volumes in the storage view to be unavailable to the initiator. Before removing an initiator, ensure that it is no longer accessing the virtual volumes in the storage view.

Steps

1. From the GUI main menu, do one of the following:
 - a. In a metro node Local:
 - i. Select **Provision Storage**.
 - ii. In the **View by** drop-down, select **Storage Views**.
 - b. In a metro node Metro:
 - i. Select **Provision Storage**.
 - ii. In the **Provision on** drop-down, select the cluster of the storage view.
 - iii. In the **View by** drop-down, select **Storage Views**.
2. In the **Storage Views** screen, select the storage view to display its components in the panel on the right.
3. Click **INITIATORS** in the panel to display the initiators (if any) in the storage view.
4. Click **ADD/REMOVE INITIATORS**.
 - To add an initiator, select the initiator in the Available initiators list, and then click **Add** to move it to the **Selected initiators** list. The added initiator displays in bold font.
 - To remove an initiator, select the initiator in the Selected initiators list, and then click **Remove** to move it to the Available initiators list. The removed initiator displays in bold font.

 **NOTE:** You can drag-and-drop items between lists.
5. Click **OK**.


A dialog box displays the results of the operation.
6. Review the results, and then click **Close**.

Adding virtual volumes to a storage view

About this task

Add a virtual volume to a storage view to allow the hosts in the view to access the volume. If a storage view is inactive, you can add or remove virtual volumes with no effect on the system.

By default, when adding a virtual volume to a storage view, metro node automatically assigns the LUN numbers, beginning with 0 (0 - 16383). However, you can manually assign the LUN number, or view the LUN number of volumes already added to the storage view.

 **NOTE:** You cannot change the LUN number of a virtual volume already added to a storage view. To change the LUN number of a virtual volume already added to a storage view, remove the virtual volume from the storage view, and then add it again using the desired LUN number. The volume cannot be actively receiving I/O.

Steps

1. From the GUI main menu, do one of the following:
 - a. In a metro node Local:
 - i. Select **Provision Storage**.
 - ii. In the **View by** drop-down, select **Storage Views**.


- b. In a metro node Metro:
 - i. Select **Provision Storage**.
 - ii. In the **Provision on** drop-down, select the cluster of the storage view.
 - iii. In the **View by** drop-down, select **Storage Views**.
2. In the **Storage Views** screen, select a storage view to display its components in the panel on the right.
3. Click the **VIRTUAL VOLUMES** tab in the panel to display the virtual volumes (if any) in the storage view.
4. Click **ADD**.
The **Add Virtual Volumes** dialog box is displayed.
5. Complete the dialog box, and then click **OK**.

Removing virtual volumes from a storage view

About this task

Before deleting virtual volumes from a storage view, note the following:

- Removing a virtual volume that is actively receiving I/O causes data to be unavailable to the volume.
- Removing a virtual volume from a storage view does not erase the data on the virtual volume.
- After deleting a volume, you can add it to another view if wanted.

 **CAUTION:** If a storage view is active, deleting a virtual volume causes the data on the virtual volume to be unavailable to applications that may be trying to access the virtual volume. Before deleting a virtual volume, ensure that no initiators are accessing the data on the virtual volume.

Steps

1. From the GUI main menu, do one of the following:
 - a. In a metro node Local:
 - i. Select **Provision Storage**.
 - ii. In the **View by** drop-down, select **Storage Views**.
 - b. In a metro node Metro:
 - i. Select **Provision Storage**.
 - ii. In the **Provision on** drop-down, select the cluster of the storage view.
 - iii. In the **View by** drop-down, select **Storage Views**.
2. In the **Storage Views** screen, select a storage view to display its components in the panel on the right.
3. Click the **VIRTUAL VOLUMES** tab in the panel to display the virtual volumes (if any) in the storage view.
4. Select the virtual volume from the list, and then click **REMOVE**.
A confirmation dialog box is displayed.
5. Review the virtual volume to delete, and then click **Yes**.

Adding or removing metro node ports from a storage view


About this task

Add metro node ports to an existing storage view to increase redundancy or processing capability, or remove ports that may have been replaced. If a storage view is inactive, you can add or remove ports with no effect on the system.

 **CAUTION:** If a storage view is active, removing a port can cause the data on all virtual volumes in the storage view to be unavailable to the initiators if no other ports in the storage view are active.

Steps

1. From the GUI main menu, do one of the following:
 - a. In a metro node Local:
 - i. Select **Provision Storage**.
 - ii. In the **View by** drop-down, select **Storage Views**.
 - b. In a metro node Metro:
 - i. Select **Provision Storage**.

- ii. In the **Provision on** drop-down, select the cluster of the storage view.
 - iii. In the **View by** drop-down, select **Storage Views**.
 2. In the **Storage Views** screen, select the storage view to display its components in the panel on the right.
 3. Click **PORTS** in the components panel to display the ports (if any) in the storage view.
 4. Click **ADD/REMOVE PORTS**. The **Add/Remove Ports** dialog box displays.
 - To add a port, select the port in the Available Ports list, and then click **Add** to move it to the Selected Ports list. The added port appears in bold font.
 - To remove a port, select the port in the Selected Ports list, and then click **Remove** to move it to the Available Ports list. The removed port appears in bold font.
-  **NOTE:** You can drag-and-drop items between lists
5. Click **OK**. The **Ports Added/Removed** dialog displays the result of the operation.

Storage view status

The following table defines a storage view's Operational states.

Operational status	Description
OK	The storage view is functioning normally.
Degraded	One or more target ports and/or virtual volumes in the storage view is unhealthy.
Error	For each target port in the view, there is an unhealthy virtual volume and/or the target port's Export status is not OK.
Stopped	There are no initiators and/or virtual volumes in the storage view.

Storage group properties

Storage Group Name	The name of the storage group.
Storage Array	The name of the array on which the storage group was created.
Array Management Provider	The name of the array management provider managing the array.
Policy	The name of the policy associated with the storage group.
Number of Volumes	The number of volumes currently in the storage group.
Parent	The name of the parent storage group (if applicable).
Masking Views	The name of the masking view associated with the storage group (if applicable).

Director properties

Director Name	The system-generated identifier for the director.
Hostname	Host name or IP address of the remote system.
Director ID	The system-generated unique identifier for this director.
Serial Number	The serial number of this director.
Port	The remote system port used for communication.
Director Type	Describes the type of director.

Director Family	The hardware family to which this director belongs.
Director Count	The number of directors this director expects at its cluster.
Auto Boot	Indicates whether the application is automatically booted on power up.
Auto Restart	Indicates whether the application is automatically rebooted after a failure.
Communication Status	Indicates the connectivity status between this instance of the CLI and the firmware.
Health State	Indicates the overall health of the director.
Health Indications	Indicates the reason for the health state that is not OK.
Operational Status	Indicates the current status of the director.

Cluster properties

Cluster Name	The name of the cluster.
Cluster ID	The ID number of the cluster.
Serial Number	The serial number of the cluster.
Island ID	The ID of the island that this cluster is currently part of.
Connected	Indicates whether or not the cluster is connected.
Directors	The list of directors in the cluster.
Transition Indications	If the operational status is not OK, this field lists the reasons why.
Transition Progress	Displays the progress of the transition, if available.
Health	Indicates the health of the cluster. See metro node cluster status .
Health Indications	Lists the reasons for the unhealthy state.
Operational Status	Indicates how the cluster is functioning. See metro node cluster status .

To view version details for the cluster, click **Version Details** at the bottom of the window.

Moving data

This section describes how to move data within or across clusters.

Topics:

- [Mobility](#)
- [Move Data Within Cluster](#)
- [Move Data Across Clusters](#)
- [Create Mobility Job wizards](#)
- [Mobility job transfer size](#)
- [Creating a mobility job](#)
- [Viewing job details](#)
- [Committing a job](#)
- [Canceling a job](#)
- [Pausing a job](#)
- [Resuming a job](#)
- [Removing the record of a job](#)
- [Changing a job transfer size](#)
- [Searching for a job](#)
- [Mobility job status](#)

Mobility

The Data Mobility feature allows you to non-disruptively move data on extents or devices to other extents or devices in the same cluster, or from devices in one cluster to devices in another cluster. **Mobility** contains wizards that allow you to create mobility jobs, and view the status and progress of those jobs. You create a mobility job to specify where to move data from (source) and where to move data to (target). Use data mobility to perform these tasks:

- Consolidate data centers.
- Move and relocate applications and data over distance.
- Migrate data faster between data centers.
- Alleviate heavy I/O to a specific storage volume that was originally configured with a large number of extents.
- Retire a storage volume.
- Relocate a "hot" device from slower to faster storage.

Data mobility workflow


1. Launch the appropriate wizard to create and start the mobility job.
2. Monitor the status of the job through completion.
3. Once completed, commit the job to complete the operation and move the data.
4. Remove the record of the mobility job (optional).

Types of data mobility operations

- **Device** — Moves data from one device to another device (within a cluster or across clusters).
- **Batch** — Groups extent or device mobility jobs into a batch job, which is then executed as a single job. The GUI does not support batch mobility jobs. To create a batch mobility job, use the CLI. See the *CLI Guide for metro node* for more information.

How data mobility works

When a mobility job begins, a temporary device is created. For extents and devices, a temporary RAID-1 is created for each source device or extent. The target extent or device is attached as a mirror leg of the temporary device, and synchronization between the source and the target begins. Once synchronization completes, both source and target are fully synchronized and you can commit (or cancel) the mobility job. Committing the job permanently transfers the data from the source to the target and completes the mobility operation. Cancelling the job detaches the target device and keeps the data on the source device. Once committed or cancelled, the temporary mobility device is deleted.



 **CAUTION:** If there is I/O to the devices being moved, and latency to the target cluster is equal to or greater than 5ms, significant performance degradation may occur.

Depending on the metro node environment, some devices cannot be moved across clusters. The Create Device Mobility Jobs wizard automatically filters out those devices that cannot be moved. If your device is not shown in the Available Devices list it could be for one of these reasons:

- The device has exported virtual volumes. Un-export the virtual volume to move this device.
- The device has virtual volumes that belong to a consistency group. Remove the virtual volume from its consistency group to move this device.

See the *Administration Guide for metro node* for more information.


Move Data Within Cluster

The **Move Data Within Cluster** option under **Mobility** provides a central location for you to create, view, and manage the status of your mobility jobs. The **Status** column is automatically sorted by job status, allowing you to see all jobs in a particular state. The arrow in the default sort column indicates the direction of the sort; ascending () or descending (). You can [sort](#) additional columns of data, [select](#) one or more jobs, and [show/hide](#) on the screen.


This screen shows the following information about your mobility jobs:

Column	Description
Mobility Job	The name assigned to the mobility job.
Type	Type of mobility job. This could be a device mobility job, or an extent mobility job.
Status	The status of the job. Mobility job status lists and describes the job states and the possible actions you can take on the job in that state.
Percentage Done	Displays the percentage of the mobility job that is complete.
Source	The name of the device from which the data is being moved.
Target	The name of the device to which the data is being moved.
Source Cluster	The cluster that the data comes from.
Target Cluster	The cluster to which the mobility job is moving data.
Transfer size	The speed of the mobility job. When creating a mobility job, you can specify the speed at which the data transfer will take place to balance your required completion time with the load restrictions on your system.
Start Time	The time at which the mobility job started. N/A may sometimes appear in this column if you are connected to one cluster while viewing a mobility job that was started on another cluster.

Filtering data


Use the Filter icon  to filter the jobs displayed on the screen.

Refreshing screen data

The time of the last data refresh displays at the upper-right of the screen. To manually refresh the data at any time, click the Refresh icon  in the upper-right corner of the screen.

Additional tasks

You can also perform the following tasks from this screen:



- Create a device mobility job — Select **CREATE > Create Device Mobility Jobs** to open the Create Device Mobility Jobs wizard.
- Perform actions on a job (pause, resume, commit, cancel, and remove) — For Pause and Resume, select the job, and then click the appropriate button. For Commit, Cancel, and Remove, select the job, and then select the appropriate option from the **MORE** drop-down.
- Search for a job — Click the Search icon  in the upper-right corner of the screen to open the **Search** page.

Opening the Move Data Within Cluster screen

From the GUI main menu, do the following:

1. Click **Mobility**.
2. In the Mobility drop-down, select the **Move Data Within Cluster**.


Move Data Across Clusters

The **Move Data Across Clusters** option under **Mobility** provides a central location for you to create, view, and manage the status of your inter-cluster mobility jobs. The Status column is automatically sorted by job status, allowing you to see all jobs in a particular state. The arrow in the default sort column indicates the direction of the sort; ascending () or descending (). You can [sort](#) columns of data, [select](#) one or more jobs, and [show/hide](#) on the screen.


This screen shows the following information about your cross-cluster mobility jobs:

Column	Description
Mobility Job Name	The name assigned to the mobility job.
Type	The type of mobility job. This could be a device mobility job, or an extent mobility job.
Status	The status of the job. Mobility job status lists and describes the job states and the possible actions you can take on the job in that state.
Percent Done	Displays the percentage of the mobility job that is complete.
Source	The name of the device from which the data is being moved.
Target	The name of the device to which the data is being moved.
Source Cluster	The cluster that the data comes from.
Target Cluster	The cluster to which the mobility job is moving data.
Transfer size	The speed of the mobility job. When creating a mobility job, you can specify the speed at which the data transfer will take place to balance your required completion time with the load restrictions on your system.
Start Time	The time at which the mobility job began. N/A may sometimes appear in the Start Time column if you are connected to one cluster while viewing a mobility job that was started on another cluster.

Filtering data


Use the Filter icon  to filter the jobs displayed on the screen.

Refreshing screen data

The time of the last data refresh displays at the upper-right of the screen. To manually refresh the data at any time, click the Refresh icon  in the upper-right corner of the screen.

Additional tasks

You can also perform the following tasks from this screen:

- Create a device mobility job — Select **CREATE > Create Device Mobility Jobs** to open the Create Device Mobility Jobs wizard.
- Perform actions on a job (pause, resume, commit, cancel, and remove) — For Pause and Resume, select the job, and then click the appropriate button. For Commit, Cancel, and Remove, select the job, and then select the appropriate option from the **MORE** drop-down.
- Search for a job — Click the Search icon  in the upper-right corner of the screen to open the **Search** page.


Opening the Move Data Across Cluster screen

From the GUI main menu, do the following:

1. Click **Mobility**.
2. In the Mobility drop-down, select the **Move Data Across Cluster**.

Create Mobility Job wizards

Mobility contains two wizards (Create Extent Mobility wizard and Create Device Mobility wizard) that allow you to create the wanted type of mobility job. Once a mobility job has started, you can view the progress of the job and perform other actions such as pause, resume, cancel, commit, and so forth.

 **CAUTION:** If there is I/O to the devices being migrated, and latency to the target cluster is equal to or greater than 5 ms, significant performance degradation may occur.

Selecting source and target devices

Keep the following in mind:

- The source device can have active I/O.
- The target device cannot be in use (no virtual volumes that are created on it).
- The target extent/device must be the same size or larger than the source extent/device.
- The target extent cannot be in use (no devices that are created on it).

Renaming mobility jobs

Extent and device mobility jobs are given a default name that you can change to a more meaningful name if wanted. To change the default name, type a new name in the **Base Name for Jobs** field in the wizard, and then click **Apply**. The new name is used as a prefix for all jobs to be created. The name is followed by an underscore and an incremented number. For example:

- *job-name_0*
- *job-name_1*
- *job-name_2*

Moving devices across clusters

Depending on the metro node environment, some devices cannot be moved across clusters. The **Create Device Mobility Jobs** wizard automatically filters out those devices that cannot be moved. If your device is not shown in the Available Devices list, it could be for one of these reasons:

- The device has exported virtual volumes. Unexport the virtual volume to move this device.
- The device has virtual volumes that belong to a consistency group. Remove the virtual volume from its consistency group to move this device.

See the *Administration Guide* for more information.

Mobility job transfer size

The transfer size determines the maximum number of bytes of data transferred at a time from the source to the target. When creating a mobility job, you can control this transfer size. The greater the size, the greater the impact on host I/O. A smaller transfer size results in the mobility job taking longer time to complete, but has a lower impact on host I/O. Monitor the mobility job's progress and its effect on host I/O. If the job is progressing too slowly, or I/O is greatly impacted, adjust the transfer size accordingly.

You can [change the transfer size](#) of a job while the job is in the queue or in progress. The change takes effect immediately. By default, the GUI transfer size is the lowest, which translates to 128 KB. This transfer size provides the best throughput while maintaining the best front-end performance in most environments.


The following table shows the different transfer-size for the mobility:

transfer-size
128 KB
2 MB
8 MB
16 MB
32 MB (ASAP)

Creating a mobility job

About this task

Creating a mobility job begins the process of moving data from one extent or device to another extent or device in the same cluster, or from one device to another device in another cluster. The procedure for moving extents and devices is the same—you need only to launch the appropriate wizard. Although you can create and start multiple mobility jobs, only the first 25 jobs run concurrently. The system allocates resources and queues the remaining jobs.

 **CAUTION:** If there is I/O to the devices being migrated, and latency to the target cluster is equal to or greater than 5 ms, significant performance degradation may occur.

Steps

1. From the GUI main menu, click **Mobility**:
 - To move data from one extent to another extent within the cluster, select **Move Data Within Cluster** from **Mobility** drop-down, and then select **Extent Mobility Jobs** from the **CREATE** drop-down.
 - To move data from one device to another device within the cluster, select **Move Data Within Cluster** from **Mobility** drop-down, and then select **Device Mobility Jobs** from the **CREATE** drop-down.
 - To move data from one device to another device across clusters, select **Move Data Across Clusters** from **Mobility** drop-down, and then select **Device Mobility Jobs** from the **CREATE** drop-down.
2. Follow the instructions on the wizard. For additional help using the wizard, click **Help**.

Viewing job details

About this task

You can view additional details for a mobility job in the properties dialog box.

Steps

1. In the GUI main menu, click **Mobility**, and then select **Move Data Within Cluster** or **Move Data Across Clusters**.
The selected view is displayed.
2. To view the **Mobility Job Properties**, click the job name link.
3. View the details of the job, and to hide the job details, click the arrow that is on the upper left of the **Device Mobility Job Properties** panel.
For help on any of the fields in the panel, click **Help**.

Committing a job

About this task

Commit a job to complete the operation and permanently move the data from the source to the target. Commit a job only after you have verified and are satisfied with the results of the operation. When you commit a job, the following occurs:

- The source extents/devices are transferred to the target extents/devices.
- The source extent/device is detached.
- The temporary extents/devices used during the migration are deleted.

Steps

1. In the GUI main menu, click **Mobility**, and then select **Move Data Within Cluster** or **Move Data Across Clusters**.
The selected view is displayed.
2. Select one or more jobs to commit.
3. Select **Commit** from **MORE** drop-down.
A confirmation dialog box is displayed.
4. Review one or more jobs to commit, and then click **OK**.

Canceling a job

About this task

Cancel a mobility job to return the source extent or device to its previous state. You can cancel a mobility job at any time until the job is committed. If the job is in progress, canceling stops the operation, and then frees any resources being used.

Steps

1. In the GUI main menu, click **Mobility**, and then select **Move Data Within Cluster** or **Move Data Across Clusters**.
The selected view is displayed.
2. Select one or more jobs to cancel.
3. Select **Cancel** from **MORE** drop-down.
A confirmation dialog box is displayed.
4. Review one or more jobs to cancel, and then click **OK**.

Pausing a job

About this task

If the performance impact of a job is too high, consider pausing the job during the day and resuming it at night or during off-peak hours to reduce the potential performance impact. You can pause and resume a job any number of times until the job completes.

Steps

1. In the GUI main menu, click **Mobility** , and then select **Move Data Within Cluster** or **Move Data Across Clusters**.
The selected view is displayed.
2. Select one or more jobs to pause.
3. Click **Pause**.
A confirmation dialog box is displayed.
4. Review one or more jobs to pause, and then click **OK**.

Resuming a job

About this task

Resume a paused mobility job to continue the mobility operation.

Steps

1. In the GUI main menu, click **Mobility** , and then select **Move Data Within Cluster** or **Move Data Across Clusters**.
The selected view is displayed.
2. Select one or more jobs to resume.
3. Click **Resume**.
A confirmation dialog box is displayed.
4. Review one or more jobs to resume, and then click **OK**.

Removing the record of a job

About this task

The system keeps a record of all mobility jobs performed. After committing or cancelling a mobility job, you can remove a record of the job. The job must be cancelled or committed before you can remove it.

Steps

1. In the GUI main menu, click **Mobility** , and then select **Move Data Within Cluster** or **Move Data Across Clusters**.
The selected view is displayed.
2. Select one or more jobs for which you want to remove the record.
3. Select **Remove** from the **MORE** drop-down.
A confirmation dialog box is displayed.
4. Review one or more jobs to remove, and then click **OK**.

Changing a job transfer size

About this task

You can change the transfer size of a job while the job status is In-progress or Queued. The new transfer size takes effect immediately. [Mobility job transfer size](#) provides more information.

Steps


1. In the GUI main menu, click **Mobility** , and then select **Move Data Within Cluster** or **Move Data Across Clusters**.
The selected view is displayed.
2. Click the job name link.
The **Mobility Job Properties** panel is displayed.
3. In the properties panel, select a new transfer size from the drop-down.
The change takes effect immediately.
4. Click **OK**.

Searching for a job

About this task

You can search for a specific job on the **Move Data Within Clusters** and **Move Data Across Clusters** screen.

Steps

1. In the GUI main menu, click **Mobility** , and then select **Move Data Within Cluster** or **Move Data Across Clusters**.
The selected view is displayed.
2. Click the Search icon  in the upper-right corner of the screen.
The **Search** page is displayed.
3. Type the job name in the text box, and then click Search icon.
If a match is found, it is displayed in the list. If no matches are found, a message is displayed.

Mobility job status

The following table lists and describes the job states, and the possible actions you can take on the job in that state:

Job state	Description	Possible actions
Ready	The job is ready to be queued.	No action is necessary.
Queued	The job is in the queue.	Cancel, Pause, or change the Transfer Size.
In-progress	The job is in progress.	Cancel, Pause, or change Transfer Size.
Paused	The job was paused while in progress.	Resume or Cancel.
Commit Pending	The job is complete. (In the CLI, this state displays as complete.)	Commit or Cancel.
Committed	The job has been committed.	Remove
Partially-committed	The commit operation has failed.	Commit the job again.
Error	The job has encountered one or more errors. The source or target of the job may be unreachable.	Cancel
Cancelled	The job was cancelled.	Remove
Partially-cancelled	The cancel operation has failed.	Cancel or Remove


Notifications

This section describes about the notifications which is a consolidated area of all appliance alerts and events for the cluster.

Topics:

- [Platform alerts view](#)
- [Hardware alerts view](#)
- [Configure alerts view](#)
- [Platform view properties](#)
- [Hardware view properties](#)

Platform alerts view

The **PLATFORM ALERTS** view lists all platform alerts generated through notification service for a corresponding NSFW event . By default, platform alerts are sorted in descending order by the last updated time. With the **Show/Hide Columns** drop-down  , you can select all columns or only required columns to show in a table.

This screen shows the following information about platform alerts:

Column	Description
Severity	<p>Indicates the urgency of alert:</p> <p>CRITICAL-A condition has occurred that can obstruct the functionality or can lead to failure of the other components.</p> <p>ERROR-An error has occurred that has a significant impact on the system and must be rectified immediately.</p> <p>WARNING-An error has occurred that you should be aware of but does not have a significant impact on the system. For example, a component is working, but its performance may not be optimum.</p> <p>INFO-An event has occurred that does not impact system functions. No action is required.</p> <p>CLEAR-When a corresponding clear event is generated for the raised alert, and the issue is resolved at the system level.</p>
Message	Indicates the cause of an event for which the alert is generated.
Scope	Represents whether it is a cluster-level or a director-level alert.
State	<p>Represents the state of the alert-OPEN, CLOSED, or ACK state.</p> <p>OPEN-The alert is active and no action has been taken on it. The alert generating condition still persists in the system. If the system administrator wants to reopen to pay attention to that event, the alert status can also be changed to open from acknowledge state.</p> <ul style="list-style-type: none"> • To reopen an alert: Select the checkbox corresponding to the alert, and then click OPEN. <p>ACK-After you view an alert and understand its contents, you can acknowledge that you have read through the alert message.</p> <ul style="list-style-type: none"> • To acknowledge an alert: Select the checkbox corresponding to the alert, and then click ACKNOWLEDGE . <p>CLOSED-Once the problem condition is resolved, the event state is updated as closed. Once closed the alerts cannot be reopened.</p> <ul style="list-style-type: none"> • To close an alert: Select the checkbox corresponding to the alert, and then click CLOSE.
Condition ID	Indicates unique ID of all defined alert definitions .
Component	System component that caused the event. Intended for service personnel.

Column	Description
Resource	It shows the actual resource for which the issue has occurred.
Count	It represents the number of times the same alert is generated over the selected period. This column is available for historical alerts only.
Event Source	Represents the context of the event. For example, Virtual Volume.
Event Source ID	Unique ID for the source of the event. Helps to narrow down to the final component.
Creation Date (UTC)	Date and time when the alert is generated.
Last Updated (UTC)	Date and time when the status of the alert is last changed.
User Notes	It shows the notes which are added by user.

Events and alerts



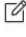
Events provide information about changes happening to the system. Alerts are events that require attention by the system administrator or user. Most alerts indicate that there is a problem with the system that must be rectified to attain the best performance from the system. The system displays live and historical alerts for Platform and Hardware in the dashboard.

Live Alerts: Alerts generated within the window of last 48 hours.

Historic Alerts: Consolidated data of the generated alerts over a period of last 6 months.

Other Tasks

From the **Platform alerts** view, you can:

- Open-Select the acknowledged platform alerts to move in the open state.
- Acknowledge-Select the platform alerts to move in the acknowledge state.
- Close-Select the platform alerts to move in the closed state.
- Configure Alerts-Select **MORE > Configure alerts**. It goes to the **Configure Alerts** page in the notifications under **Settings**.
- Disable Alerts-Select **MORE > Disable alerts**. It disables the selected platform alerts.
- Refresh-To refresh the page, click the  icon .
- Export-To export the page as a CSV file, click the  icon.
- User Notes-To add notes in the User Notes column, click the empty cell of a selected alert, it opens the Platform Alert Properties panel, go to the Notes property and click Add Notes, type the message and save it. An icon  is displayed once the note is added in the properties panel.
- iDRAC GUI-To get more details on iDRAC GUI, click the iDRAC GUI link from the top-right corner.
- Director node-By default, it shows the logged in director node (for example, director-1-1-A). To change the node, select the node from the drop-down.

NOTE: On switching a director node, you receive an error message `The token is expired or unavailable.` Please authorize certificates at "`<director>`" before you proceed. So, authorize the certificate before you switch the director node to get peer node alert details.
- View historical data-To view the data, the filters are available up to 30 days. On selection of a particular filter, it displays the data for the number of days in the filter along with the current day data (for example, on selection of 7 days filter, it displays the data for last seven days along with the current day data). To view the data beyond 30 days, select the **Customized Dates** from the drop-down.

NOTE: The customized date filters contain the data from 00:00:00 hours start date to 23:59:59 hours end date. For example, if the start date is selected as 1 Jul 2020 and end date as 15 Jul 2020, the user see the alerts from Jul 1 00:00:00 hours to Jul 15 23:59:59 hours.
- Platform Alerts Properties panel-To view the platform alert properties in detail, select the platform alert to open the panel, and then you can view the properties in detail at the right side of the screen.

Opening the Platform alerts view

From the UI main menu, do the following:

1. In a metro node:
 - a. Select **Notifications**.
 - b. In the **Notifications**, select **PLATFORM ALERTS** tab.

Hardware alerts view

The **HARDWARE ALERTS** view lists all hardware alerts generated through notification service for a corresponding iDRAC or Monitor event. By default, hardware alerts are sorted in descending order by the last updated time. With the **Show/Hide**

Columns drop-down , you can select all columns or only required columns to show in a table.

This screen shows the following information about platform alerts:

Column	Description
Severity	Indicates the urgency of alert: CRITICAL -A condition has occurred that can obstruct the functionality or can lead to failure of the other components. ERROR -An error has occurred that has a significant impact on the system and must be rectified immediately. WARNING -An error has occurred that you should be aware of but does not have a significant impact on the system. For example, a component is working, but its performance may not be optimum.
Message	Indicates the cause of an event for which the alert is generated.
State	Represents the state of the alert-OPEN, CLOSED, or ACK state. OPEN -The alert is active and no action has been taken on it. The alert generating condition still persists in the system. If the system administrator wants to reopen to pay attention to that event, the alert status can also be changed to open from acknowledge state. <ul style="list-style-type: none">● To reopen an alert: Select the checkbox corresponding to the alert, and then click OPEN. ACK -After you view an alert and understand its contents, you can acknowledge that you have read through the alert message. <ul style="list-style-type: none">● To acknowledge an alert: Select the checkbox corresponding to the alert, and then click ACKNOWLEDGE . CLOSED -Once the problem condition is resolved, the event state is updated as closed. Once closed the alerts cannot be reopened. <ul style="list-style-type: none">● To close an alert: Select the checkbox corresponding to the alert, and then click CLOSE.
Severity Code	Represents the numerical code for the corresponding severity of the alert.
Version	Represents version of the certificate.
Category	Represents category of the event (for example, Systems).
Message ID	Indicates ID of the event record.
Count	It represents the number of times the same alert is generated over the selected period. This column is available for historical alerts only.
Appname	Represents the device or application that originated the message.
Host Name	Represents IP address or network name of the remote host.
Last Updated (UTC)	Date and time when the status of the alert is last changed.
Creation Date (UTC)	Date and time when the alert is generated.

Column	Description
User Notes	It represents the notes which are added by the user.

Events and alerts




Events provide information about changes happening to the system. Alerts are events that require attention by the system administrator or user. Most alerts indicate that there is a problem with the system that must be rectified to attain the best performance from the system. The system displays live and historical alerts for Platform and Hardware in the dashboard.

Live Alerts: Alerts generated within the window of last 48 hours.

Historic Alerts: Consolidated data of the generated alerts over a period of last 6 months.

Other Tasks

From the **Hardware alerts** view, you can:

- Open-Select the acknowledged hardware alerts to move in the open state.
- Acknowledge-Select the hardware alerts to move in the acknowledge state.
- Close-Select the hardware alerts to move in the closed state.
- Alert Source Type: Select the source type of the alerts either iDRAC or Monitor.
- Refresh-To refresh the page, click the  icon .
- Export-To export the page as a CSV file, click the  icon.
- User Notes-To add notes in the User Notes column, click the empty cell of a selected alert, it opens the Platform Alert Properties panel, go to the Notes property and click Add Notes, type the message and save it. An icon  is displayed once the note is added in the properties panel.
- iDRAC GUI-To get more details on iDRAC GUI, click the iDRAC GUI link from the top-right corner.
- Director node-By default, it shows the logged in director node (for example, director-1-1-A). To change the node, select the node from the drop-down.

NOTE: On switching a director node, you receive an error message The token is expired or unavailable. Please authorize certificates at "<director>" before you proceed. So, authorize the certificate before you switch the director node to get peer node alert details.
- View historical data-To view the data, the filters are available up to 30 days. On selection of a particular filter, it displays the data for the number of days in the filter along with the current day data (for example, on selection of 7 days filter, it displays the data for last seven days along with the current day data). To view the data beyond 30 days, select the **Customized Dates** from the drop-down.

NOTE: The customized date filters contain the data from 00:00:00 hours start date to 23:59:59 hours end date. For example, if the start date is selected as 1 Jul 2020 and end date as 15 Jul 2020, the user see the alerts from Jul 1 00:00:00 hours to Jul 15 23:59:59 hours.
- Hardware Alerts Properties panel-To view the hardware alert properties in detail, select the hardware alert to open the panel, and then you can view the properties in detail at the right side of the screen.

Opening the Hardware alerts view

From the UI main menu, do the following:

1. In a metro node:
 - a. Select **Notifications**.
 - b. In the **Notifications**, select **HARDWARE ALERTS** tab, and the select the **Alert source type**.

Configure alerts view


The **Configure alerts** view lists all alert definitions in the system. By default, configure alerts are sorted in ascending order by condition ID. You can enable and disable one or multiple alerts to control the logging of alerts in the system.

This screen shows the following information about configure alerts:

Column	Description
Condition ID	Indicates Unique ID of all defined alert definitions.
Name	Name of the alert definition.
Enabled	Indicates if the alert definition is enabled or disabled.
Event Source	Defines the source of the alert definition.
Callhome Enabled	It indicates whether the alerts are CallHome supported or not.
Description	Explains about the defined alert definition.

Other Tasks

From the **Configure alerts** view, you can:

- Enable Alerts-Select the alert to enable the alert definitions.
- Disable Alerts-Select the alert to disable the alert definitions.
- Configure Notifications- Select the alert to configure the notifications.
 - NOTE:** It controls the enable and disable of sources of alerts. If you disable the **Event Alert Feature**, then it does not show future alerts that are generated on live listing pages of Platform and Hardware alerts. Similarly, if you disable **Hardware Alerts** or **Platform Alerts**, it does not show respective future alerts on live listing pages.
- Configure Email Notifications-Email notifications can be configured to enable or disable as a whole service. You can also disable the email notifications platform, iDRAC, or monitor alerts individually.
- Test Alerts-Select the alert to see the result of test alert.
- Test Email Alert-Select the alert to see the test email. If Test Email Alert is not enabled, then enable the Email Notification to enable it.
- Refresh-To refresh the page, click the  icon .

Opening the Configure alerts view

From the UI main menu, do the following:

1. In a metro node:
 - a. Select **Notifications**.
 - b. In the **Notifications**, select **PLATFORM ALERTS** tab.
 - c. Select the **Configure Alerts** from the **MORE** drop-down. It goes to the **Configure Alerts** page in the notifications under **Settings**
2. Select the **Notifications** from the **Settings** drop-down. It goes to the **Configure Alerts** page in the notifications under **Settings**

Configure notifications

Using Configure Notification feature, alert notifications can be enabled or disabled. It controls the enable and disable of sources of alerts. If you disable the Event Alert Feature, then it does not show future alerts that are generated on live listing pages of Platform and Hardware alerts. Similarly, if you disable Hardware Alerts or Platform Alerts, it does not show respective future alerts on live listing pages. Disabling the alerts from UI means that the user does not want to see the alerts in the UI. It does not mean that the alerts are not generated. The alerts keep on generating and are stored in the database.

You can enable or disable email notification at multiple levels:

- Event and Alert Feature level
- Component Group level
- Component level

Event and Alert feature


Disabling the Event and Alert feature disables all the notifications and you do not receive any notifications that are generated for alerts. To receive the notifications for generated alerts in future, enable this feature. To enable or disable this feature, go to **Settings > Notifications > Configure Alerts > Configure Notification**, and enable or disable the **Event and Alert** feature depending upon the requirement.

Component group level

The alerts can be enabled or disabled at the component group level. The component group can be either Platform alerts or Hardware alerts. Enabling or disabling this feature enables or disables the alerts that are associated with all components of that component group. The user can enable or disable both the component group or either one of them. To enable or disable this feature, go to **Settings > Notifications > Configure Alerts > Configure Notification**, and enable or disable the **Hardware Alerts** or **Platform Alerts** feature or both depending upon the requirement.

Component level

The email notification can be enabled or disabled at the component level. Enabling the email notification at component level allows notifications for generated alerts that are associated with the respective components. Disabling the email notification at component level do not allow notifications for generated alerts that are associated with the respective components. To enable or disable the email notification at component level, go to **Settings > Notifications > Configure Alerts > Configure Notification**, then enable or disable the email notification of the individual component depending upon the requirement.

 **NOTE:** Enabling or disabling of Platform alerts at component level affects respectively on related Condition IDs. The following table shows the Component—>Condition ID mapping.

Component	Condition ID	Event Source
Certificate	0x8A4a61F6	CERTIFICATES
	0x8A4a61F7	CERTIFICATES
	0x8A4a91F8	CERTIFICATES
	0x8A4a91F9	CERTIFICATES
Cluster	0x30013	CLUSTER
	0x60001	CLUSTER
	0x60005	CLUSTER
	0x70001	CLUSTER
	0x70002	CLUSTER
	0x70003	CLUSTER
	0x70004	CLUSTER
	0x70005	CLUSTER
	0x70006	CLUSTER
	0x70007	CLUSTER
Consistency Group	0xd0002	CONSISTENCYGROUP
Device	0x3000b	DEVICE
	0x30015	DEVICE
Director	0x100001	DIRECTOR
	0x100002	DIRECTOR
	0x60002	DIRECTOR
	0x60003	DIRECTOR

Component	Condition ID	Event Source
	0xc0001	DIRECTOR
	0xd0001	DIRECTOR
	0x40001	DIRECTOR
	0x40002	DIRECTOR
	0x40003	DIRECTOR
Distributed Device	0x30014	DISTRIBUTEDDEVICE
	0x3001b	DISTRIBUTEDDEVICE
Fabric Port	0x90001	TARGETPORT
	0x9000e	TARGETPORT
IP Port	0x110001	IPPORT
	0x120003	COMMUNICATIONSPATH
	0x150008	VIRTUALPORT
	0x15000d	VIRTUALPORT
	0x150010	VIRTUALPORT
	0x150011	INTERFACE
	0x150018	VIRTUALPORT
	0x150019	INTERFACE
	0x15001d	INTERFACE
	0x15001e	INTERFACE
IP Port	0x1300c8	COMMUNICATIONSPATH
	0x1300c9	COMMUNICATIONSPATH
	0x1300ca	COMMUNICATIONSPATH
	0x1300cb	COMMUNICATIONSPATH
	0x1701f5	COMMUNICATIONSPATH
	0xa0000	INTERFACE
	0xa0001	INTERFACE
	0xa0002	INTERFACE
	0xa0003	INTERFACE
	0xa0004	INTERFACE
	0xa0005	INTERFACE
	0x13006f	PORT
	0x150001	VIRTUALPORT
	0x15000e	VIRTUALPORT
	0x15000f	VIRTUALPORT
	0x170192	VIRTUALPORT
	0x170193	VIRTUALPORT
	0x170194	VIRTUALPORT
Logical Unit	0x20004	LOGICALUNIT

Component	Condition ID	Event Source
	0x20005	LOGICALUNIT
	0x20006	LOGICALUNIT
	0x20009	LOGICALUNIT
	0x2000d	LOGICALUNIT
Meta Volume	0x30005	METAVOLUME
	0x30010	METAVOLUME
	0x30011	METAVOLUME
	0x30012	METAVOLUME
	0x30017	METAVOLUME
	0x30019	METAVOLUME
	0x3001c	METAVOLUME
	0x3001d	METAVOLUME
	0x3001e	METAVOLUME
	0x3001f	METAVOLUME
	0x30020	METAVOLUME
	0xf0001	METAVOLUME
	0xf0002	METAVOLUME
	0xf0003	METAVOLUME
	0xf0004	METAVOLUME
	0x8A00010E	METAVOLUME
	0x8A00010F	METAVOLUME
	0x8A000110	METAVOLUME
	0x8A000111	METAVOLUME
Storage Array	0x20001	ARRAY
	0x20003	ARRAY
	0x20007	ARRAY
	0x20008	ARRAY
	0x2000a	ARRAY
	0x2000c	ARRAY
	0x2000b	ITLNEXUS
	0x20002	ITNEXUS
	0x2000e	TARGET
Storage View	0x9000c	STORAGEVIEW
	0x9000d	STORAGEVIEW
Storage Volume	0x30004	STORAGEVOLUME
	0x3000e	STORAGEVOLUME
	0x3000f	STORAGEVOLUME
	0x30018	STORAGEVOLUME

Component	Condition ID	Event Source
	0x3001a	STORAGEVOLUME
	0x180001	DISKID
	0x180002	DISKID
	0x190001	DISKID
System	0xb0001	SYSTEM
	0xb0001	SYSTEM
Virtual Volume	0x30003	VIRTUALVOLUME
	0x30006	VIRTUALVOLUME
	0x3000a	VIRTUALVOLUME
	0x3000c	VIRTUALVOLUME
	0x3000d	VIRTUALVOLUME
	0x30016	VIRTUALVOLUME
	0x90002	VIRTUALVOLUME
	0x90003	VIRTUALVOLUME
	0x90004	VIRTUALVOLUME
	0x90005	VIRTUALVOLUME
	0x90006	VIRTUALVOLUME
	0x90007	VIRTUALVOLUME
	0x90008	VIRTUALVOLUME
	0x90009	VIRTUALVOLUME
	0x9000a	VIRTUALVOLUME
	0x9000b	VIRTUALVOLUME

Configure email notifications

Using configure Email Notification feature, email notifications can be enabled or disabled. It controls email notification of alerts. On disabling the Email Notification feature, you do not receive the email messages regarding the alerts.

You can enable or disable email notification at multiple levels:

- Email Notification feature level
- Component level

Email notification feature level

If Email Notification feature is enabled, then you receive the email messages for generated alerts. If you do not want to receive any alerts email messages in future, then disable this feature. To enable or disable this feature, go to **Settings > Notifications > Configure Email Notifications**, then enable or disable the **Email Notification** feature depending upon the requirement.

Component level

The email notification can be enabled or disabled at the component level. Enabling the email notification at component level allows notifications for generated alerts that are associated with the respective components. Disabling the email notification at component level do not allow notifications for generated alerts that are associated with the respective components. To enable or disable the email notification at component level, go to **Settings > Notifications > Configure Email Notification**, then enable or disable the email notification of the individual component depending upon the requirement.

Test Alerts

The Test Alerts provides the information about the alerts are generating or not.


Prerequisites

The **Event Alert Feature** should be enabled.

Steps

1. From the UI, go to the **Settings > Notifications** or go to the **Notifications > Platform Alerts > MORE > Configure Alerts**. The **Configure alerts** page is displayed.
2. Select **More > Test Alerts** . A window is displayed with the list of following alerts to test:
 - a. Platform alert
The two test alerts are generated from firmware:
 - Cluster scope critical alarm
 - Director scope error operational events
 - b. Monitor alert
 - c. SMS alert
3. Select the required alerts from the list and click **YES**. A **Test Alert Result** window is displayed which shows the result of selected test alerts.

The state of these alerts is **OPEN**.

 **NOTE:** Before triggering test alerts, close all existing test alerts.

Attribute	Explanation
Alert type	It is the name of alert type.
Details	It displays the status of generated test alert.
Message	It provides the description of the status for generated alerts.

4. Click the **CLOSE** button. It closes the window and changes the state of the alerts from **OPEN** to **CLOSE**.

Test Email


It is used to check that the email alerts are generating or not.

Prerequisites

- The **Email Notification Feature** should be enabled.
- SMTP details must be configured.

Steps

1. From the UI, go to the **Settings > Notifications** or go to the **Notifications > Platform Alerts > MORE > Configure Alerts**. The **Configure alerts** page is displayed.
2. Select **MORE > Test Email Alert**. An email alert is sent to the respective configured email ID.

 **NOTE:** If **Email Notification Feature** is not enabled, then selection of **Test Email Alert** is disabled for a user.

Platform view properties

Platform alert tab

Severity	<p>Indicates the urgency of alert:</p> <p>CRITICAL-A condition has occurred that can obstruct the functionality or can lead to failure of the other components.</p> <p>ERROR-An error has occurred that has a significant impact on the system and must be rectified immediately.</p> <p>WARNING-An error has occurred that you should be aware of but does not have a significant impact on the system. For example, a component is working, but its performance may not be optimum.</p> <p>INFO-An event has occurred that does not impact system functions. No action is required.</p> <p>CLEAR-When a corresponding clear event is generated for the raised alert, and the issue is resolved at the system level.</p>
State	<p>Represents the state of the alert-OPEN, CLOSED, or ACK state.</p> <p>OPEN-The alert is active and no action has been taken on it. The alert generating condition still persists in the system. If the system administrator wants to reopen to pay attention to that event, the alert status can also be changed to open from acknowledge state.</p> <ul style="list-style-type: none">• To reopen an alert: Select the checkbox corresponding to the alert, and then click OPEN. <p>ACK-After you view an alert and understand its contents, you can acknowledge that you have read through the alert message.</p> <ul style="list-style-type: none">• To acknowledge an alert: Select the checkbox corresponding to the alert, and then click ACKNOWLEDGE . <p>CLOSED-Once the problem condition is resolved, the event state is updated as closed. Once closed the alerts cannot be reopened.</p> <ul style="list-style-type: none">• To close an alert: Select the checkbox corresponding to the alert, and then click CLOSE.
Message	Indicates the cause of an event for which the alert is generated.
Category	Displays the event category.
Enabled	Indicates whether the alert is enabled or not.
Alert ID	Represents the unique ID assigned to each alert .
Alert Name	Represents the name of the alert.
Description	Describes the platform alert.
Corrective Action	Action to eliminate the cause of event.
Scope	Represents whether it is a cluster-level or a director-level alert.
Condition ID	Indicates Unique ID of all defined alert definitions.
Component	System component that caused the event. Intended for service personnel.
Event Source	Represents the context of the event. For example, Virtual Volume.
Event Source ID	Unique ID for the source of the event. Helps to narrow down to the final component .
Resource	It shows the resource for which the issue has occurred.
Count	It represents the number of times the same alert is generated over the selected period. This column is available for historical alerts only.
Scope Incarnation	Represents unique ID assigned to each director and cluster.
Creation Date	Date and time when the alert is generated.
Last Updated	Date and time when the status of the alert is last changed.

External RCA	Represents external root cause analysis of the issue.
Notes	It shows the notes which are added by user.
Debug Events	Associated legacy debug events generated through firmware.

Additional details tab

To view additional information about the alert in name-value pair, click the Additional details tab.

Logs tab

It is for historical alerts only. It provides the list of creation and last updated time for the same alert has occurred.

Hardware view properties

Hardware alert tab

Severity	<p>Indicates the urgency of alert:</p> <p>CRITICAL-A condition has occurred that can obstruct the functionality or can lead to failure of the other components.</p> <p>ERROR-An error has occurred that has a significant impact on the system and must be rectified immediately.</p> <p>WARNING-An error has occurred that you should be aware of but does not have a significant impact on the system. For example, a component is working, but its performance may not be optimum.</p>
State	<p>Represents the state of the alert-OPEN, CLOSED, or ACK state.</p> <p>OPEN-The alert is active and no action has been taken on it. The alert generating condition still persists in the system. If the system administrator wants to reopen to pay attention to that event, the alert status can also be changed to open from acknowledge state.</p> <ul style="list-style-type: none"> • To reopen an alert: Select the checkbox corresponding to the alert, and then click OPEN. <p>ACK-After you view an alert and understand its contents, you can acknowledge that you have read through the alert message.</p> <ul style="list-style-type: none"> • To acknowledge an alert: Select the checkbox corresponding to the alert, and then click ACKNOWLEDGE . <p>CLOSED-Once the problem condition is resolved, the event state is updated as closed. Once closed the alerts cannot be reopened.</p> <ul style="list-style-type: none"> • To close an alert: Select the checkbox corresponding to the alert, and then click CLOSE.
Message	Indicates the cause of an event for which the alert is generated.
Enabled	Indicates whether the alert is enabled or not.
Alert ID	Represents the unique ID assigned to each alert .
Severity Code	Represents the numerical code for the corresponding severity of the alert.
Version	Represents version of the certificate.
Message ID	Indicates ID of the event record.
Host	Display the host details.
Host Name	Represents IP address or network name of the remote host.
AppName	Represents the device or application that originated the message.
Creation Date	Date and time when the alert is generated.

Last Updated	Date and time when the status of the alert is last changed.
Notes	It shows the notes which are added by user.

Logs tab

It is for historical alerts only. It provides the list of creation and last updated time for the same alert has occurred.


Topics:

- [Create Mobility Job wizards](#)
- [Supported host types](#)
- [Storage group properties](#)
- [Director properties](#)
- [Cluster properties](#)
- [Metro node Witness status](#)
- [Configure Settings](#)
- [Metro node cluster status](#)
- [Product version](#)
- [Step 1: Select or create a consistency group for the virtual volume](#)
- [Step 1: Create a consistency group](#)
- [Step 2: Select volume options](#)
- [Step 3: Select a storage pool](#)
- [Step 3: Select a pool for each mirror on the second cluster](#)
- [Step 3: Select a pool for each mirror in the cluster](#)
- [Step3: Create thin virtual volumes](#)
- [Select a storage view for the virtual volume\(s\) \(optional\)](#)
- [Step 5: Review your selections](#)
- [Step 6: View results](#)
- [Step 2: Select volume options](#)
- [Step 2: Select volume options](#)
- [Step 3: Select a storage volume to create the virtual volume](#)
- [Step 3: Select a source and target storage volume](#)
- [Step 3: Create thin volumes](#)
- [Step 3: Select a target storage volume on the remote cluster](#)
- [Step 3: Select target storage on the remote cluster](#)
- [Step 6: View results](#)
- [Show Logical Units](#)
- [The Claim Storage wizard](#)
- [Metro node overview](#)
- [Pool properties](#)
- [Virtual volume properties](#)
- [Show ITLs dialog box](#)
- [Consistency group properties](#)
- [Device component properties](#)
- [Device properties](#)
- [Distributed device properties](#)
- [Register Array Management Provider dialog box](#)
- [Add capacity to virtual volumes](#)
- [Extent properties](#)
- [Initiator properties](#)
- [Logical unit properties](#)
- [ALUA Support field values](#)
- [Visibility field values](#)
- [Extent or Device mobility job properties](#)
- [Metro node port properties](#)
- [Storage array properties](#)

- [Storage view properties](#)
- [Storage volume properties](#)
- [Provision Job properties](#)
- [The Distributed Consistency Groups view](#)
- [Create Virtual Volumes dialog box](#)
- [Platform view properties](#)
- [Hardware view properties](#)
- [Configure alerts properties](#)

Create Mobility Job wizards

Mobility contains two wizards (Create Extent Mobility wizard and Create Device Mobility wizard) that allow you to create the wanted type of mobility job. Once a mobility job has started, you can view the progress of the job and perform other actions such as pause, resume, cancel, commit, and so forth.

 **CAUTION:** If there is I/O to the devices being migrated, and latency to the target cluster is equal to or greater than 5 ms, significant performance degradation may occur.

Selecting source and target devices

Keep the following in mind:

- The source device can have active I/O.
- The target device cannot be in use (no virtual volumes that are created on it).
- The target extent/device must be the same size or larger than the source extent/device.
- The target extent cannot be in use (no devices that are created on it).

Renaming mobility jobs

Extent and device mobility jobs are given a default name that you can change to a more meaningful name if wanted. To change the default name, type a new name in the **Base Name for Jobs** field in the wizard, and then click **Apply**. The new name is used as a prefix for all jobs to be created. The name is followed by an underscore and an incremented number. For example:

- `job-name_0`
- `job-name_1`
- `job-name_2`

Moving devices across clusters

Depending on the metro node environment, some devices cannot be moved across clusters. The **Create Device Mobility Jobs** wizard automatically filters out those devices that cannot be moved. If your device is not shown in the Available Devices list, it could be for one of these reasons:

- The device has exported virtual volumes. Unexport the virtual volume to move this device.
- The device has virtual volumes that belong to a consistency group. Remove the virtual volume from its consistency group to move this device.

See the *Administration Guide* for more information.

Supported host types

Metro node supports the following types of hosts:

- SUN-VCS
- AIX
- HPUX
- IBM-D910

- Default (use this type for all hosts not listed above)

Storage group properties

Storage Group Name	The name of the storage group.
Storage Array	The name of the array on which the storage group was created.
Array Management Provider	The name of the array management provider managing the array.
Policy	The name of the policy associated with the storage group.
Number of Volumes	The number of volumes currently in the storage group.
Parent	The name of the parent storage group (if applicable).
Masking Views	The name of the masking view associated with the storage group (if applicable).

Director properties

Director Name	The system-generated identifier for the director.
Hostname	Host name or IP address of the remote system.
Director ID	The system-generated unique identifier for this director.
Serial Number	The serial number of this director.
Port	The remote system port used for communication.
Director Type	Describes the type of director.
Director Family	The hardware family to which this director belongs.
Director Count	The number of directors this director expects at its cluster.
Auto Boot	Indicates whether the application is automatically booted on power up.
Auto Restart	Indicates whether the application is automatically rebooted after a failure.
Communication Status	Indicates the connectivity status between this instance of the CLI and the firmware.
Health State	Indicates the overall health of the director.
Health Indications	Indicates the reason for the health state that is not OK.
Operational Status	Indicates the current status of the director.

Cluster properties

Cluster Name	The name of the cluster.
Cluster ID	The ID number of the cluster.
Serial Number	The serial number of the cluster.
Island ID	The ID of the island that this cluster is currently part of.
Connected	Indicates whether or not the cluster is connected.
Directors	The list of directors in the cluster.

Transition Indications	If the operational status is not OK, this field lists the reasons why.
Transition Progress	Displays the progress of the transition, if available.
Health	Indicates the health of the cluster. See metro node cluster status .
Health Indications	Lists the reasons for the unhealthy state.
Operational Status	Indicates how the cluster is functioning. See metro node cluster status .

To view version details for the cluster, click **Version Details** at the bottom of the window.

Metro node Witness status

The metro node Witness status view shows the following information:

Name	The name of the component. For a metro node cluster, this is the name of the cluster in the metro node Metro configuration. For the metro node Witness server this is server.
Private IP Address	Private IP address of the Cluster Witness Server VM used for cluster witness-specific traffic.
Public IP Address	Public IP address of the Cluster Witness Server VM used as an endpoint of the IPsec tunnel.
Admin State	Identifies whether metro node Witness is enabled or disabled. The Admin states are defined in the following table.
Operational state	Operational state of the corresponding component. The Operational states are defined in one of the following tables.
Management Connectivity	Reachability of the specified Witness component over the IP management network from the management server where the CLI command is run. The Management Connectivity states are defined in one of the following tables.
Diagnostic information	String generated based on the analysis of the data and status information reported by the corresponding component. The diagnostic information is defined in one of the following tables.

Admin state	Definition
Enabled	Metro node Witness functionality is enabled on this component.
Disabled	Metro node Witness functionality is disabled on this component.
Inconsistent	All Cluster Witness components are reachable over the management network but some components report their administrative state as disabled while others report it as enabled. This is a rare state which may result failure during enabling or disabling.
Unknown	This component is not reachable and its administrative state cannot be determined.

Operational state	Definition
For Cluster Witness server:	
clusters in contact	Both clusters are in contact with each other over the inter-cluster network.
cluster partition	The clusters are partitioned from each other over the inter-cluster network, and the Cluster Witness server can still talk to each cluster.
cluster unreachable	One cluster has either failed or become isolated (that is, partitioned from its peer cluster and disconnected from the Cluster Witness server).
unknown	Cluster Witness server does not know the states of one or both of the clusters and needs to learn them before it can start making decisions. The Cluster Witness server assumes this state upon startup.
For metro node clusters:	

Operational state	Definition
in contact	This cluster is in contact with its peer over the inter-cluster network. Rebuilds may be in progress. Subject to other system-wide restrictions, I/O to all distributed virtual volumes in all synchronous consistency groups is allowed from the perspective of metro node Witness.
cluster partition	This cluster is not in contact with its peer and the Cluster Witness server has declared that two clusters partitioned. Subject to other system-wide restrictions, I/O to all distributed virtual volumes in all synchronous consistency groups (with specific preference rule set) is allowed from the perspective of metro node Witness.
remote cluster isolated or dead	This cluster is not in contact with its peer and the Cluster Witness Server has declared that the remote cluster (the peer) was isolated or dead. Subject to other system-wide restrictions, I/O to all distributed virtual volumes in all synchronous consistency groups (with specific preference rule set) is allowed from the perspective of metro node Witness.
local cluster isolated	This cluster is not in contact with its peer and the Cluster Witness Server has declared that the remote cluster (the peer) is the only proceeding cluster. This cluster must suspend I/O to all distributed virtual volumes in all synchronous consistency groups regardless of cluster preference.
cluster state unknown	This cluster is not in contact with its peer over the inter-cluster network and is awaiting guidance from the Cluster Witness Server. I/O to all distributed volumes in all consistency groups is suspended regardless of bias.


Management Connectivity state	Definition
ok	The component is reachable.
not reachable	The cluster is not reachable.
server not reachable	The server is not reachable.

Diagnostic information	Definition
Local cluster-x hasn't yet established connectivity with the server	The cluster has never connected to Cluster Witness Server.
Remote cluster-x hasn't yet established connectivity with the server	The cluster has never connected to Cluster Witness Server.
Cluster-x has been out of touch from the server for X days, Y secs	Cluster Witness Server has not received messages from a given cluster for longer than 60 seconds.
Cluster witness server has been out of touch for X days, Y secs	Cluster witness server has been out of touch for X days, Y secs.
Cluster Witness is not enabled on component-X, so no diagnostic information is available	Cluster Witness Server or either of the clusters is disabled.

Configure Settings

Use this dialog box to configure the inactivity timeout for your connection to the GUI, or to define the data reload policy when navigating through provisioning storage object views.

GUI idle timeout

 **WARNING:** A potential security risk can be exposed by setting idle time to greater than 30 minutes. Therefore, if you set the idle timeout to greater than 30 minutes, a warning will display on this dialog box.

The idle timeout is based on the URL used to access the GUI. Therefore when you access the GUI on cluster 1 and set the idle timeout, it will apply only to cluster 1. If you log into cluster 2, the idle timeout dialog will not have the same value as that of cluster 1, because the URL is different.

The value of the timeout field can be an integer value between 1 and 12 if the field unit value is set to hours. The value of the Timeout field can be an integer value between 1 and 720 if the field unit value is set to minutes.


NOTE: The value for Timeout is stored in a file named using the URL with which you accessed the GUI. If you access the GUI for a cluster using different URLs (For example, using the DNS name of the cluster or using the IP address), the timeout value for each URL used will be different.

Provision storage data reload policy

Select one of the following options to define how often data is reloaded when navigating through provisioning views.

- **Always** — Always reload data upon navigation to a storage object view. This is the default setting.
- **Never** — Never reload data upon navigation to a storage object view.
- **After** — The number of hours or minutes after which to reload data when you navigate to a storage object view. Enter a number in the field and then select the time interval. Once you navigate to an object view, the data will automatically reload at the specified time interval.

The data reload policy setting applies only to the URL and client used to access the GUI when you set the policy. Therefore, if you log in to the GUI again using a different client (laptop versus desktop) or a different URL (cluster 1 versus cluster 2), the reload policy setting will not be applied to that GUI session.

To manually refresh the data, click the Refresh icon  in the upper-right corner of the screen title bar or the **Last Refreshed** link at the bottom of the screen.

Metro node cluster status

The following table defines the values for the metro node cluster states.

Status Type	Values	Definition
Operational Status	OK	The cluster is functioning normally.
	Degraded	The cluster is not operating as configured. This may indicate one or more of the following problems: degraded redundancy level (for example, a director is dead), suspended virtual volumes or exported volumes, storage volumes not visible from all directors, or the metadata volume is not yet processed.
	Unknown	Metro node encountered a problem determining the operational status of the cluster. This may indicate a degraded state, since it usually means that at least one of the directors is not responding or is communicating abnormally.
	Transitioning	The cluster is reacting to external events and may not be operating as configured. I/O may be suspended during the transition period.
	Isolated	The cluster is not communicating with any other clusters.
	Not-running	All directors in any one of the clusters are in Stopped state and when the VPN is stopped between clusters.
Transitioning Indication	Suspended exports	Cluster lost local quorum during the unvaulting state. This is possibly due to a disk not visible from all directors, no active meta-volume, delayed cluster transition, metadata, not all directors are present, or no quorum.
	Suspended volumes	The distributed device is not accepting new I/O; pending I/O requests are frozen.
Health state	OK	The cluster is functioning normally.

Status Type	Values	Definition
	Degraded	The cluster is not functioning at an optimal level. This may indicate non-functioning remote virtual volumes, unhealthy devices or storage volumes, suspended devices, conflicting director count configuration values, out-of-date devices, and so forth.
	Unknown	Metro node cannot determine the cluster's health state, or the state is invalid.
	Critical failure	The cluster is not functioning and may have failed completely. This may indicate a complete loss of back-end connectivity through all directors in the cluster.
	Major failure	<p>The cluster is failing and some functionality may be degraded or unavailable. This may indicate a loss of back-end connectivity through one or more (but not all) directors. When any director has issues with field replacement or if the required number of uninterrupted power supplies is not detected. If any one of the directors is waiting to gain quorum or unvaulting. Example:</p> <pre>Major-failure: engine-2-1 : director-2-1-A : Director has not gained quorum Major-failure: director is unvaulting</pre>
	Minor failure	The cluster is functioning, but some functionality may be degraded. This may indicate one or more unreachable storage volumes.
LOCAL_COM	OK	All LOCAL-COM links have the expected connectivity: this port-group is operating correctly.
	warning	Some links have unexpected connectivity. This port-group is operational but not properly configured. Performance may not be optimal.
	error	Some connectivity is missing from this port-group. It is not operating correctly.
	fail	All connectivity is missing from this port-group. LOCAL-COM is not operational.
WAN_COM	OK	The WAN COM switch is operating normally
	full	All port-groups have a status of either ok or warning. wan-com connectivity is complete through minor configuration errors may still exist. See individual port-group statuses.
	partial	Some port-groups have a status of error or fail, but at least one port-group has a status of OK or warning. wan-com is operating (possibly minimally) through at least one channel. Performance is degraded.
	none	All port-groups have a status of either error or fail. wan-com is not operational.
	not-applicable	The system is a single-cluster (i.e. metro node Local) system. Validating wan-com connectivity is not applicable.

Product version

The **Product Version** dialog box shows the following product version details for the cluster:

Product Version	The product version of GeoSynchrony running on the cluster.
SMSv2	The version of SMSv2 running on the management server.
Mgmt server base	The version of base software running on the management server.

Mgmt server software	The version of management server software running.
Cluster Witness Server Software	The version of Cluster Witness server software running (if Cluster Witness is installed).
Directors	Lists the name and version number of each director in the cluster.

Step 1: Select or create a consistency group for the virtual volume

Select or create a consistency group for the virtual volume.

To view the topology of a consistency group, select the group in the list. In the map that displays to the right, the Storage icon



and blue shading around a cluster indicate that the selected consistency group uses storage from that cluster.

- To add the virtual volume to an existing consistency group, select **Add storage to existing consistency group** and then select the consistency group in the list.
- To filter the list of consistency groups by name, type the full or partial name of the consistency group in the **Filter** text box. To filter by storage location, click the drop-down box to the right of the **Filter** text box, and then select the storage location.
- To create a new consistency group for the virtual volume, select **Create new consistency group**.


Step 1: Create a consistency group



Complete the following fields to create a consistency group for the virtual volume.

Consistency Group Name — Type a name for the consistency group.

Consistency Group Type — Select the type of consistency group to create: Local or Distributed.

- For local consistency groups, select the cluster on which to create the consistency group.
- For Distributed consistency groups, select a detach rule. The detach rule determines if IO continues and at which cluster if an inter-cluster link failure occurs.

 **NOTE:** Detach rules are applied per consistency group.

Detach rule	Cluster-1	Cluster-2
No automatic winner	Suspends I/O  NOTE: This detach rule is generally used if you have created external scripts or use other programs to determine which cluster detaches and continues I/O during a link failure.	Suspends I/O  NOTE: This detach rule is generally used if you have created external scripts or use other programs to determine which cluster detaches and continues I/O during a link failure.
Winner: <cluster-1> (5 seconds)	Suspends I/O	Suspends I/O
Winner: <cluster-2> (5 seconds)	Suspends I/O	Suspends I/O

Step 2: Select volume options


Select mirroring on the cluster where the virtual volumes will be created (optional), the quantity, size and base name for the virtual volumes, and then enter a name for the provisioning job (optional).

High Availability — Select this check box to enable mirroring at the cluster indicated. The option to mirror at a cluster is disabled if you have less than two arrays in the cluster.

- For a local consistency group, select the check box to enable mirroring at the local cluster.

- For a distributed consistency group, select the appropriate check box to enable mirroring at one cluster or both clusters, or leave both check boxes clear to disable mirroring.

Number of Volumes — The number of volumes to create. You can create a maximum of 10 volumes at a time. Each volume will have the same RAID geometry.

 **NOTE:** For Unity arrays, you can provision only one volume at a time.

Capacity Per Volume — The size of each volume. Select the size in MB, GB or TB. The minimum size is 150MB. Each volume will have the same capacity.

Total Capacity — The total capacity of all the virtual volumes to be created.

Volume Base Name — Type a base name for the virtual volumes. A number will be added to the base name for each volume created. For example, using the base name Test, volumes will be named Test_1, Test_2, Test_3, and so forth. You can reuse the base name to provision additional volumes. Metro node will continue to increment the number at the end of the name, beginning with the highest existing numeric suffix.

Job Name — Select the **Specify a job name** check box, and then type a name for the provisioning job to be created. This field is optional. If you do not specify a job name, the volume base name will be used. If the job name you enter is already in use, a numeric suffix will be added to the job name to make it unique, for example, *jobname_n*. Where n is the first available number that makes the name unique. Or, you can choose to use a different name. You can see currently used job names in the **Provision Jobs Status** page under **Monitoring**. The job name must meet the following criteria:

- Maximum of 64 characters
- Must begin with a letter
- Must include only numbers, letters, underscores, or dashes

To discard all selections made on the previous screen and restart the wizard from Step 1, click **Start Over** in the lower-right corner of the screen.

To go back to a specific step, click **Back** until you reach that step.

Step 3: Select a storage pool

Select a storage array, and then select a storage pool and storage group on the array for the volumes to be created.

V3GR arrays only: Selecting a policy/workload

For V3GR arrays, policies define a level of performance (based on an average response time) required by the storage workload. Refer to the array documentation for more information.


Note the following:

- If you select a policy, storage pools and storage groups with the same policy are available for selection.
- If you select **None**, pools with any type of policy, and workloads with no policy are available for selection.
- If you select **Optimized**, only system-optimized pools and storage groups with optimized policies are available for selection.
- Selecting a workload is optional, however, if you select a policy and a workload, the storage pools and storage groups with the same Policy+Workload are listed for selection.

Selecting a storage pool


The Storage Pool section shows the following information about the storage pools on the selected array:

- **Name** — The name of the pool on the array.
- **Policy** — The type of policy associated with the storage pool.
- **Type** — Indicates whether the pool is capable of provisioning thick volumes, thin volumes, or both. If you select a thin pool, a thin rebuild will be automatically set for the storage volume that is created during the provisioning process.
- **Free** — The amount of storage available for provisioning.
- **Total** — The total amount of storage available from the pool.

 **NOTE:** To create a thin volume, the selected pools must be thin capable and the supporting arrays must belong to the same array family.

Selecting a storage group

By default, metro node will automatically select a storage group for the newly created volumes, based on compatible policies, and the lowest amount of volumes in the group. However, if you want to select the storage group, click the **Storage Group** drop-down box, and then select `Select from list` to show the available storage groups.

 **NOTE:** For V3GR arrays only, the **Storage Group** drop-down box changes to a selection list because a storage group is required. For all other arrays, the storage group is optional.

The Storage Group section shows the following information about available storage groups:

- **Name** — The name of a storage group on the selected array.
- **Policy** — The type of policy associated with the storage pool (if any). metro node lists only storage groups with policies that match that of the selected pool, and storage groups with no associated policies.
- **Parent** — The name of the parent storage group (if applicable).
- **Masking views** — The name of the masking view associated with the storage group (if applicable).
- **# Volumes** — The number of volumes currently in the storage group.

To discard all selections made on previous screens and restart the wizard from Step 1, click **Start Over** in the lower-right corner of the screen.

To go back to a specific step, click **Back** until you reach that step.

Step 3: Select a pool for each mirror on the second cluster

Select a storage pool on the second cluster to provide storage for the virtual volumes to be created.

V3GR arrays only: Selecting a policy or workload

For V3GR arrays, policies define a level of performance (based on an average response time) required by the storage workload. Refer to the array documentation for more information.


Note the following:

- When you select a policy, storage pools and storage groups with the same policy are available for selection.
- If you select `None`, pools with any type of policy, and workloads with no policy are listed for selected for selection.
- If you select `Optimized`, only system-optimized pools and storage groups with optimized policies are available for selection.
- Selecting a workload is optional, however, if you select a policy and a workload, the storage pools and storage groups with the same Policy+Workload are listed for selection.

Selecting a storage pool


The Storage Pool section shows the following information about the storage pools on the selected array:

- **Name** — The name of the pool on the array.
- **Policy** — The type of policy associated with the storage pool.
- **Type** — Indicates whether the pool is capable of provisioning thick volumes, thin volumes, or both. If you select a thin pool, a thin rebuild will be automatically set for the storage volume that is created during the provisioning process.
- **Free** — The amount of storage available for provisioning from the pool.
- **Total** — The total amount of storage available from the pool.

 **NOTE:** To create a thin volume, the selected pools must be thin capable and the supporting arrays must belong to the same family of arrays.

Selecting a storage group

By default, metro node will automatically select a storage group for the newly created volumes, based on compatible policies, and the lowest amount of volumes in the group. However, if you want to select a storage group, click the **Storage Group** drop-down box, and then select `Select from list` to show the available storage groups to choose from.

 **NOTE:** For V3GR arrays only, the **Storage Group** drop-down box changes to a selection list because a storage group is required. For all other arrays, the storage group is optional.

The Storage Group section shows the following information about available storage groups:

- **Name** — The name of a storage group on the selected array.
- **Policy** — The type of policy associated with the storage pool (if any). Metro node lists only storage groups with policies that match that of the selected pool, and storage groups with no associated policies.
- **Parent** — The name of the parent storage group (if applicable).
- **Masking views** — The name of the masking view associated with the storage group (if applicable).
- **# Volumes** — The number of volumes currently in the storage group.

To discard all selections made on previous screens and restart the wizard from Step 1, click **Start Over** in the lower-right corner of the screen.

To go back to a specific step, click **Back** until you reach that step.

Step 3: Select a pool for each mirror in the cluster

Select an array to provide storage for each mirror in the cluster, and then select a pool from the list of available pools on the array. To ensure high availability, the array selected for one mirror is automatically filtered from the list of arrays available for other mirror.

V3GR arrays only: Selecting a policy or workload

For V3GR arrays, policies define a level of performance (based on an average response time) required by the storage workload. Refer to your array documentation for more information.


Note the following:

- When you select a policy, storage pools and storage groups with the same policy are available for selection.
- If you select `None`, pools with any type of policy, and workloads with no policy are listed for selected for selection.
- If you select `Optimized`, only system-optimized pools and storage groups with optimized policies are available for selection.
- Selecting a workload is optional, however, if you select a policy and a workload, the storage pools and storage groups with the same `Policy+Workload` are listed for selection.

Selecting a storage pool

The Storage Pool section shows the following information about the storage pools on the selected array:

- **Name** — The name of the pool on the array.
- **Policy** — The type of policy associated with the storage pool.
- **Type** — Indicates whether the pool is capable of provisioning thick volumes, thin volumes, or both. If you select a thin pool, a thin rebuild will be automatically set for the storage volume that is created during the provisioning process.
- **Free** — The amount of storage available for provisioning from the pool.
- **Total** — The total amount of storage available from the pool.

 **NOTE:** To create thin volumes, the selected pools must be thin capable and the supporting arrays must belong to the same array family.

Selecting a storage group

By default, metro node will automatically select a storage group for the newly created volumes, based on compatible policies, and the lowest amount of volumes in the group. However, if you want to select a storage group, click the **Storage Group** drop-down box, and then select `Select from list` to show the available storage groups to choose from.

NOTE: For V3GR arrays only, the **Storage Group** drop-down box changes to a selection list because a storage group is required. For all other arrays, the storage group is optional.

The Storage Group section shows the following information about available storage groups:

- **Name** — The name of a storage group on the selected array.
- **Policy** — The type of policy associated with the storage pool (if any). Metro node lists only storage groups with policies that match that of the selected pool, and storage groups with no associated policies.
- **Parent** — The name of the parent storage group (if applicable).
- **Masking views** — The name of the masking view associated with the storage group (if applicable).
- **# Volumes** — The number of volumes currently in the storage group.

To discard all selections made on previous screens and restart the wizard from Step 1, click **Start Over** in the lower-right corner of the screen.

To go back to a specific step, click **Back** until you reach that step.

Step3: Create thin virtual volumes

The pools you have selected support creating thin virtual volumes. To create thin volumes, Select the **Create thin virtual volumes** check box, otherwise, click **Next** to continue without creating thin volumes.

NOTE: To create thin virtual volumes, the selected storage arrays must belong to the same array family, and pools must support creating thin volumes. For Unity and XtremIO arrays, the storage volumes are always thin.

Select a storage view for the virtual volume(s) (optional)

Select a storage view for the virtual volume(s), or select `None` to continue and select a storage view later.

- If you select an active storage view, the virtual volume is added to that storage view and automatically becomes visible to hosts with access to the storage view. In some cases a rescan of the host may be required. If the storage view is inactive, go to the **Storage Views** screen and add the components necessary to activate the storage view. [Creating a storage view](#) provides more information.
- If you select `None`, when you are ready to expose the virtual volume to hosts, you must manually add the virtual volume to a storage view. [Adding or removing virtual volumes from a storage view](#) provides more information.

To discard all selections made on previous screens and restart the wizard from the Step 1, click **Start Over** in the lower-right corner of the screen.

To go back to a specific step, click **Back** until you reach that step.

Step 5: Review your selections

This screen shows a summary of your selections in the wizard and a topology map of the volume to be created.

Review your selections, and if satisfied, click **Provision Storage** in the lower-right corner of the screen to submit the provisioning request.

To make changes to a specific step, click **Back** until you reach that step.

To discard all selections made on previous screens and restart the wizard from the Step 1, click **Start Over** in the lower-right corner of the screen.

Step 6: View results

This screen shows the results of the provisioning request.

- To manage the provisioning job, click the **View job Status** link to open the **Provisioning Jobs Status** screen.
- To provision more storage, click **Provision More** in the lower-right corner of the screen to go back to Step 1 in the wizard, otherwise, click **Done** to close the wizard.

Step 2: Select volume options

Select high availability options, synchronization options, and then enter a name for the volume. If a cluster has only one array, the ability to mirror is disabled for that cluster.

- **Source Cluster** — Select the cluster where you want to create the virtual volume.
- **High Availability** — Select a check box to enable mirroring at the cluster indicated. You can select a single cluster, both clusters, or leave both check boxes unselected to disable mirroring at both clusters. When you select the check box to enable mirroring at a cluster, *Local Mirroring* appears above the cluster in the map.
- **Synchronization Options** — Select the appropriate option to indicate whether or not you want to synchronize data from the source storage volume to the target storage volume.
- **Volume Name** — Type a name for the virtual volume.

To discard all selections made on the previous screen and restart the wizard from Step 1, click **Start Over** in the lower-right corner of the screen.

To go back to a specific step, click **Back** until you reach that step.

Step 2: Select volume options

Select high availability options and a name for the virtual volume. The mirroring option is enabled only for the cluster where the virtual volume will be created. If the cluster has only one array, the option to mirror is disabled.

Source Cluster — The cluster where the virtual volume will be created.

High Availability — Select the check box to mirror the virtual volume at the cluster indicated.

Volume Name — Type a name for the virtual volume.


To discard all selections made on the previous screen and restart the wizard from the Step 1, click **Start Over** in the lower-right corner of the screen.

To go back to a specific step, click **Back** until you reach that step.

Step 3: Select a storage volume to create the virtual volume

Select one storage volume to use for the virtual volume. The virtual volume will use the entire capacity of the storage volume. You can select claimed or unclaimed storage volumes.

Select an array in the drop-down box, and then select a storage volume from the list of available storage volumes on the array.

 **NOTE:** To create a thin volume, make sure that the storage volume is thin capable.

To discard all selections made on previous screens and restart the wizard from Step 1, click **Start Over** in the lower-right corner of the screen.


To go back to a specific step, click **Back** until you reach that step.

Step 3: Select a source and target storage volume

Select a source and a target storage volume on the cluster where the virtual volume will be created. The source storage is used to create the virtual volume and the target storage is used to mirror the virtual volume. The virtual volume to be created will use

the entire capacity of the source storage volume. To ensure high availability, the array selected to provide the source storage is automatically filtered from the list of arrays available for the target storage.

Select only one source and one target storage volume. The storage volume can be claimed or unclaimed.

 **NOTE:** To create a thin volume, the selected source and target storage volumes must come from the same array family, and must also be thin capable.

Source Storage — Select an array in the drop-down box, and then select from the list of available storage volumes on the array.

Target Storage — Select an array in the drop-down box, and then select a storage volume from the list of available storage volumes on the array.

To discard all selections made on previous screens and restart the wizard from Step 1, click **Start Over** in the lower-right corner of the screen.

To go back to a specific step, click **Back** until you reach that step.

Step 3: Create thin volumes

The storage volumes you selected support creating thin virtual volumes. To create thin virtual volumes, click the **Create thin virtual volumes** check box, or click **Next** to continue without creating thin volumes.

Step 3: Select a target storage volume on the remote cluster

Select a target storage volume on the remote cluster. Data written to the source storage volume will be copied onto this storage volume.

Note the following:

- You can select claimed or unclaimed storage volumes.
- The size the target volume must be greater than or equal to the size of the source volume.


Select an array in the drop-down box, and then select from the list of available storage volumes on the array.

To discard all selections made on previous screens and restart the wizard from the Step 1, click **Start Over** in the lower-right corner of the screen.

To go back to a specific step, click **Back** until you reach that step.

Step 3: Select target storage on the remote cluster

When mirroring on both clusters, select target storage volumes on the remote cluster. The data on the source storage volume will be mirrored to the target storage volumes on the remote cluster. To ensure high availability, the array selected for one target is automatically filtered from the list of arrays available for the other target.

 **NOTE:** The size of the target volume must be greater than or equal to the size of the source volume. You can select claimed or unclaimed storage.

To discard all selections made on previous screens and restart the wizard from the Step 1, click **Start Over** in the lower-right corner of the screen.

To go back to a specific step, click **Back** until you reach that step.

Step 6: View results

This screen shows the results of the provisioning request.

To provision more storage, click **Provision More** in the lower-right corner of the screen to go back to Step 1 in the wizard.

To close the wizard, click **Done**.

Show Logical Units

Name (VPD ID)	The VPD identifier for the logical unit.
Storage Volume	The storage volume to which this logical unit belongs.
LUNs	The LUN number on the array.
Connectivity	The connectivity status of the logical unit.
Visibility	Indicates how the logical unit is visible to the controllers. See possible values .
ALUA Support	Indicates if Asymmetrical Logical Unit Access (ALUA) is supported on the array. Depending on the type of controller exposed to the LUN, one of these values will display.

The Claim Storage wizard

Use the Claim Storage wizard to claim storage an array. To encapsulate storage volumes already in use, see the Encapsulation procedures in the SolVe Desktop, available on Dell EMC Online Support.

Naming storage volumes

By default, storage volumes are identified by their SCSI Vital Product Data (VPD), for example: VPD83T3:600601604a950d00fe9bc44d11b3da11. If the array requires a Name Mapping file, use this file to map the default storage volume names to more user-friendly names that allow you to easily identify the storage volume.

For arrays that do not require a Name Mapping file, the storage volumes are named using one of the following formats:

- `<array_name><last_four_digits_of_array_serial_number>_<device_number>`
- `<array_name><last_four_digits_of_array_serial_number>_<LUN_number>`

You can rename the storage volumes after claiming.

Selecting a storage type

Select a storage type to assign different categories of data to different types of storage media as a means of reducing total storage cost. Categories of data can be based on levels of protection needed, performance requirements, frequency of use, and other considerations. As an example, you could designate all storage designated as "High" to be used for highly classified or critical information, which would be stored on more expensive, high quality storage arrays. Data designated as "Medium" might be seldom-used data, which could be stored on a less expensive storage array, and so forth. By default, the storage type is included in the storage volume name.

Selecting a thin-rebuild

For thinly-provisioned storage volumes, you can select if the storage volume will use thin rebuilds. [Setting a thin rebuild for storage volumes](#) provides more information.

Metro node overview

The metro node family is a solution for federating Dell EMC and non Dell EMC storage arrays. Metro node resides between the servers and heterogeneous storage assets supporting a variety of arrays from various vendors. Metro node simplifies storage management by allowing LUNs provisioned from various arrays to be managed through a centralized management interface. Metro node removes physical barriers within, across and between data centers.

- Metro node Local provides simplified management and non-disruptive data mobility across heterogeneous arrays.
- Metro node Metro provides data access and mobility between two metro node clusters within synchronous distances.

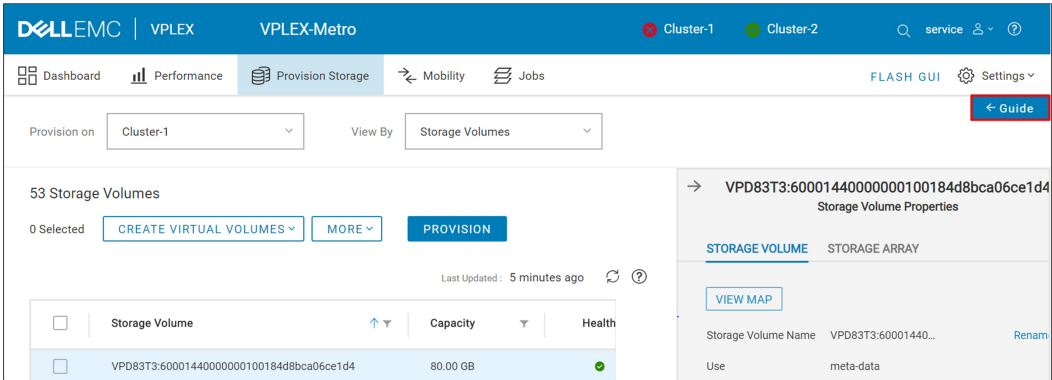
With a unique scale-up and scale-out architecture, metro node's advanced data caching and distributed cache coherency provides workload resiliency, automatic sharing, balancing and failover of storage domains, and enables both local and remote data access with predictable service levels.

System management

In addition to the web-based Graphical User Interface (UI), metro node also supports a command-line interface (CLI) for managing your system. The *CLI Guide for the metro node appliance* provides more information.

Basic workflow

To begin using metro node, provision and export storage so that hosts and applications can use the storage. For an overview of the provisioning process and step-by-step instructions for provisioning storage, see **Guide**. To view the Guide, click **Provision Storage**, and then expand the **GUIDE** at the upper-right corner of the screen.



Pool properties

Pool Name	The name of the pool available on the array.
Storage Array	The name of the managed array on which the Pool was created.
Supported Pool Types	The type of storage supported by the pool. Possible values are: Thin and/or Thick.
Policy	The name of the policy associated with the storage pool.
Free Capacity	The amount of free capacity available from the pool.
Used Capacity	The amount of capacity already used from the pool.
Total Capacity	The total capacity of the pool.

Virtual volume properties

Virtual Volume Name	The name of the virtual volume. You can change the name in this field.
Supporting Device	The device on which the virtual volume was created.
Consistency Group	Identifies the metro node consistency group to which this virtual volume belongs.
Locality	Designates where the virtual volume is located. Values can be cluster 1, cluster 2 or distributed.
Visibility	Indicates where a virtual volume is visible: Global (both clusters) or Local (only at this cluster).

Expansion Method	<p>The expansion method available for this volume. Possible values:</p> <p><code>concatenation</code> — The volume can be expanded using Concatenation or RAID-C expansion.</p> <p><code>storage-volume</code> — The volume can be expanded to the Expandable capacity using storage volume expansion.</p> <p><code>not-supported</code> — The volume does not support expansion.</p>
Expandable By	The amount of space that is available for volume expansion. A volume that has an Expandable by of 0 can still be expanded using concatenation. A volume that has an Expandable by of > 0 can be expanded using storage volume expansion.
Expansion Status	<p>The expansion status of the volume. The volume can be:</p> <ul style="list-style-type: none"> <code>-</code> (dash) — This volume can be expanded. <code>failed</code> — The last volume expansion on this volume failed. <code>unknown</code> — The volume expansion status is unknown. <code>in-progress</code> — The volume cannot be expanded because it has a volume expansion in progress.
Block Count	The number of blocks in the underlying extent.
Block Size	The size of a single block, in kilobytes.
Capacity	The size of the virtual volume.
Thin Enabled	<p>Indicates if the virtual volume is created as a thin volume. Possible values:</p> <ul style="list-style-type: none"> <code>Enabled</code> — The volume was created as a thin volume. <code>Disabled</code> — The volume was not created as a thin volume although the underlying storage is thin-capable. <code>Unavailable</code> — The underlying storage is not thin-capable. <p>You can change this property for XtremIO, VMAX and VNX virtual volumes only in the CLI. Refer to the CLI Guide for information.</p>
Storage Array Family	<p>The family of arrays that the supporting array belongs to. Possible values:</p> <ul style="list-style-type: none"> <code>XTREMIO</code> — The array belongs to the XtremIO family of arrays. <code>CLARiion</code> — The array belongs to the VNX or Unity family of arrays. <code>SYMMETRIX</code> — The array belongs to the VMAX or Symmetrix family of arrays. <code>Mixed</code> — The supporting arrays belong to different array families. For example, you might have a RAID-1 device with one leg on a VNX array and the other leg on a VMAX array. <code>Other</code> — Metro node does not recognize the array or does not support creating thin volumes on the array. <code>-</code> (dash) — The Storage Array Family property is not populated. There may be a connectivity problem.
Health	The overall status of the virtual volume. See Virtual volume states.
Health Indications	The reasons for the virtual volumes unhealthy state. See Virtual volume status.
Operational Status	Indicates how the virtual volume is functioning in the cluster. See Virtual volume status.
Service Status	Indicates how the virtual volume is functioning in the cluster. See Virtual volume status.

Show ITLs dialog box

This dialog box shows a read-only list of Initiator-Target-LUN combinations that shows the storage volume's back-end path (metro node back-end port, storage port, and LUN). you can also search for a specific initiator, target or LUN by entering the number in the **Find** text box and pressing **Enter**.

The following table describes the fields shown in the dialog box:

Field/Column	Description
Storage Volume Name	The name of the selected storage volume.

Field/Column	Description
System ID	The array's ID number for the storage volume.
Initiator	The Port WWN of the metro node backend port.
Target	The Port WWN of the storage array target port.
LUN	The Logical Unit Number for the storage volume.

Consistency group properties

Consistency Group Name	The name of the consistency group. You can change the name in the properties dialog box.
Detach Rule (metro node Metro only)	<p>The detach rule that applies to the consistency group. It designates which cluster detaches if the clusters lose connectivity. Possible values:</p> <p>no-automatic-winner — The consistency group does not select a winning cluster.</p> <p>winner <cluster-name> <delay> — The cluster specified by <cluster-name> is declared the winner after the inter-cluster link outage lasts more than the number of seconds specified by delay. If a consistency group has a detach-rule configured, the rule applies to all volumes in the consistency group, and overrides any rule-sets applied to individual volumes.</p>
Visibility	Lists the clusters at which the consistency group is visible.
Storage At Clusters	Indicates if the volumes in the consistency group use storage from one cluster, or both clusters.

Device component properties

Component Name	The name of the device's component. You can change the name from this dialog box.
Geometry	The underlying RAID structure of the device.
Health	The overall health of the component.
Health Indications	A list of reasons for the component's unhealthy state.
Operational Status	Indicates how the component is functioning in the cluster. See Operational states.
Service Status	Indicates whether I/O can proceed on the device. See Service states.

Device properties

Device Name	The name of the top-level device. You can change the name in the properties dialog box. The name must be unique across clusters.
Geometry	The underlying RAID structure of the device.
Virtual Volume	The name of the virtual volume created on the device.
Locality	Indicates where the device is visible: Local or Remote.
Cluster	The name of the cluster to which the device belongs.
Block Count	The number of blocks in the underlying extent.
Block Size	The size of the data blocks used internally when managing I/O. The block size determines the granularity of storage volume I/O and the rebuild rate of distributed devices.
Capacity	The size of the device.

Thin Capable	Indicates if the underlying storage volume is thinly-provisioned.
Storage Array Family	<p>The family of arrays that the supporting array belongs to. Possible values:</p> <ul style="list-style-type: none"> • XTREMIO — The array belongs to the XtremIO family of arrays. • CLARiION — The array belongs to the VNX or Unity family of arrays. • SYMMETRIX — The array belongs to the VMAX or Symmetrix family of arrays. • Mixed — The supporting arrays belong to different array families. For example, a RAID-1 device with one leg on a VNX array and the other leg on a VMAX array. • Other — metro node does not recognize the array or does not support creating thin volumes on the array. • - (dash) The Storage Array Family property is not populated. There may be a connectivity problem.
Stripe Depth	For RAID-0, the depth of the stripe in bytes.
Health	The overall health of the device. See Health states .
Health Indications	The reasons for the device's unhealthy state. If the Health Indications include rebuilding, you can view the rebuild status for the device by clicking the Go to Rebuild Status view link.
Operational Status	Indicates how the device is functioning in the cluster. See Operational states .
Service Status	Indicates the status of I/O to the device. See Service states .

Distributed device properties

Distributed Device Name	The name of the device. You can change the name in this dialog box.
Transfer Capacity	The maximum capacity at which a distributed device can mirror data on the underlying devices. The transfer capacity is always the same as the capacity of the smallest device, since the system cannot mirror more data than the smallest device can handle.
Virtual Volume	The name of the virtual volume created on the device. By default, the virtual volume inherits the name of the device with _vol appended at the end of the name.
Thin Capable	Indicates if the underlying storage volume is thinly-provisioned.
Storage Array Family	<p>The family of arrays that the supporting storage array belongs to. Possible values:</p> <ul style="list-style-type: none"> • XTREMIO — The array belongs to the XtremIO family of arrays. • CLARiION — The array belongs to the VNX or Unity family of arrays. • SYMMETRIX — The array belongs to the VMAX or Symmetrix family of arrays. • Mixed — The supporting arrays belong to different array families. For example, you might have a RAID-1 device with one leg on a VNX array and the other leg on a VMAX array. • Other — metro node does not recognize the array or does not support creating thin volumes on the array. • - (dash) — The Storage Array Family property is not populated. There may be a connectivity problem.
Rule Set	The name of the rule set applied to the device. You can change the rule set in this dialog box.
Health	The overall health of the device. See Health states .
Health Indications	Lists the reasons for the device's unhealthy state. If the Health Indications of the distributed device indicate that it is rebuilding, click the Go to Rebuild Status view link to monitor the rebuild status of the distributed device.
Operational Status	Indicates how the device is functioning. See Operational states .
Service Status	Indicates whether I/O can proceed on the device. See Service states .

Register Array Management Provider dialog box

Provider Type	The type of AMP used to manage the array. When provisioning from XtremIO arrays, select <code>REST</code> . For all other arrays, select <code>SMI-S</code> . For XtremIO v3.0, metro node requires one REST AMP per array. Therefore, if you are using multiple v3.0 arrays, register a REST AMP for each array. For XtremIO v4.0, one REST AMP can manage multiple v4.0 arrays. When you enter the array credentials, metro node automatically detects the array versions managed by the XMS.
Provider name	Type a user-friendly name for the AMP.
IP Address	Type the IP address of the AMP.
Port	The port number to use with the IP address to create the URL that connects to the AMP. Although a default port number is provided, you can use any valid port number.
User Name	The username used to log into the AMP.
Password	The password used to log into the AMP.
Use SSL	Determines how the secure socket layer (ssl) protocol is used. <ul style="list-style-type: none">- For SMI-S AMPs, select this check box to use https when connecting to the AMP's URL. If unchecked, http is used.- For REST AMPs, https is required, therefore, the check box is selected by default.

Add capacity to virtual volumes

Use this dialog box to expand a virtual volume using concatenation.

Add Capacity to Virtual Volume	The selected virtual volume that you want to expand.
Current Capacity	The size of the selected virtual volume.
Additional Capacity Needed	Enter the amount of space to add to this virtual volume. Metro node will look for volumes with at least that available capacity. To enter the value using a different unit, select the correct unit from the drop-down list. When you enter the amount of space needed, you can click Show Available to list the storage volumes with a minimum capacity equal to the additional capacity needed. If you change the capacity needed, click Show Available again to update Available Storage Volumes list.
Total New Capacity	The new size of the virtual volume based on the additional capacity needed.
Available storage volumes	The list of available storage volumes to choose from based on the additional capacity needed. The name of the storage volume and the largest free chunk available in the storage volume are displayed.

Extent properties

Extent Name	The name of the extent. You can change the name of the extent from this dialog box. By default, extents are named after the underlying storage volume with the prefix "extent_".
Use	Indicates whether the extent is currently in use. Possible values: Claimed or Used.
Used By	The name of the device created from the extent.
Storage Volume	The name of the underlying storage volume.
Block Count	The number of blocks in the underlying extent.
Block Offset	The starting block number for the extent. If not specified, extents are created using sequential block numbers on a storage volume. To specify the starting block number, use the CLI.
Block Size	The size of the data blocks used internally when managing I/O. The block size determines the granularity of disk I/O and the rebuild rate of distributed devices.
Capacity	The size of the extent.

Thin Capable	Indicates if the underlying storage volume is thinly-provisioned.
Storage Array Family	<p>The family of arrays that the supporting array belongs to. Possible values:</p> <ul style="list-style-type: none"> • <code>XTREMIO</code> — The array belongs to the XtremIO family of arrays. • <code>CLARiION</code> — The array belongs to the VNX or Unity family of arrays. • <code>SYMMETRIX</code> — The array belongs to the VMAX or Symmetrix family of arrays. • <code>Other</code> — metro node does not recognize the array or does not support creating thin volumes on the array. • <code>-</code> (dash) — The Storage Array Family property is not populated. There may be a connectivity problem.
Health	Indicates the overall health of the extent. See Health states .
Health Indications	Displays the reasons for the extent's unhealthy state.
I/O Status	The status of I/O to the underlying storage volume. See I/O states .
Operational Status	Indicates how the extent is functioning in the cluster. See Operational states .

Initiator properties

Initiator Name	The name given to the initiator. You can change the name in this dialog box.
Host Type	Identifies the host operating system. To change the host type after the initiator has been registered, use the CLI.
Node WWN	A unique network identifier for the HBA's interface card. All ports on the card share this World Wide Name (WWN).
Port WWN	A unique network identifier for the port.
Target Port	The list of metro node ports visible to the initiator.

Logical unit properties

Name	The VPD identifier for the logical unit.
Active AAO Controller	The active (or ALUA optimized) controller for the logical unit.
Active AAO Visibility	Lists all directors that have a path to the active (or ALUA optimized) controller for this logical unit.
Passive AAN Controller	The passive (or ALUA non-optimized) controller for the logical unit.
Passive AAN Visibility	Lists all directors that have a path to the passive (or ALUA non-optimized) controller for this logical unit.
LUNs	The LUN number on the array.
Storage Volume	The storage volume to which this logical unit belongs.
ALUA Support	Indicates if Asymmetrical Logical Unit Access (ALUA) is supported on the array. Depending on the type of controller exposed to the LUN, one of these values will display. Possible values:

`Implicit` — The logical unit can change its ALUA access state internally within the array without any external input requesting a change.

`Explicit` — The logical unit can change its ALUA access state in response to a request from an initiator (metro node backend software).

`Implicit explicit` — The logical unit has both implicit and explicit capabilities.

None — ALUA is not supported on the array.

Connectivity Status The connectivity status of the logical unit.

Visibility Indicates how the logical unit is visible to the controllers. Possible values:

Active optimized — All commands and task management functions can be sent on this path; this path has the highest performance (ALUA arrays only).

Active non-optimized — All commands and task management functions can be sent on this path, but the path has a lower performance than an active optimized path (ALUA arrays only).

Active — The controller actively processing I/O requests.

Both — Both controllers process I/O and provide standby capability for each other. One controller is actively processing I/O while the other is idle in standby mode waiting to process I/O should the active primary controller fail or be taken offline.

None — ALUA is not supported on the array.

ALUA Support field values

Value	Definition
Implicit	The logical unit can change its ALUA access state internally within the array without any external input requesting a change.
Explicit	The logical unit can change its ALUA access state in response to a request from an initiator (metro node back-end software).
Implicit explicit	The logical unit has both implicit and explicit capabilities.
None	ALUA is not supported on the array.

Visibility field values

Value	Definition
Active optimized	All commands and task management functions can be sent on this path; this path has the highest performance (ALUA arrays only).
Active non-optimized	All commands and task management functions can be sent on this path, but the path has a lower performance than an active optimized path (ALUA arrays only).
Active	The controller is actively processing I/O requests.
Passive	The controller is in standby mode waiting to process I/O should the active primary controller fail or be taken offline.
Both	Both controllers process I/O and provide standby capability for each other. One controller is actively processing I/O while the other is idle in standby mode waiting to process I/O should the active primary controller fail or be taken offline.
None	ALUA is not supported on the array.

Extent or Device mobility job properties

Mobility Job Name The name of the mobility job.

Source The name of the source extent or source device.

Target The name of the target extent or target device.

Start Time	The date and time the job started. Note that N/A may display in this field if you are connected to one cluster and view the properties of a job that was started on another cluster.
Transfer Speed	The maximum number of bytes of data transferred at a time from the source to the target. You can change the transfer speed in the properties dialog box while the job is in the queue or in progress. Mobility job transfer speed provides more information on selecting the transfer speed.
Status	The status of the extent or device mobility job. Mobility job status provides more information.

Metro node port properties

Enabled Port	Indicates whether or not the port is enabled. Use the checkbox to enable or disable the port.
Port Name	The ID number for the metro node port.
Port Name	Name of the port.
Director ID	The ID of the director on which the port is located.
Port WWN	The port's World Wide Name (WWN). The WWN is a unique network identifier for the port.
Port Node WWN	The port's node World Wide Name (WWN). The WWN is a unique network identifier for the port's node.
Export Status	Indicates whether or not this port is used to export storage. See metro node port status .

Storage array properties

Storage Array Name	The name of the storage array manufacturer.
Storage Array Family	<p>The family of arrays that the storage array belongs to. Possible values:</p> <ul style="list-style-type: none"> • XTREMIO — The array belongs to the XtremIO family of arrays. • CLARiion — The array belongs to the VNX or Unity family of arrays. • SYMMETRIX — The array belongs to the VMAX or Symmetrix family of arrays. • Other — Metro node does not recognize the array or does not support creating thin volumes on the array. • – (dash) — The Storage Array Family property is not populated. There may be a connectivity problem.
Controllers	A list of controllers in the array. This list is empty if the array is not supported.
Ports	A list of all port WWNs discovered for the array.
Logical Unit Count	The number of LUNs exported from the array.
Connectivity Status	Indicates if all LUNs exported from the array are visible to both directors. See Storage array connectivity status .

Storage view properties

Storage view Name	The name of the storage view. You can change the name in this field.
Operational Status	Indicates how the storage view is functioning in the cluster. See Storage view status.
Active	Indicates whether or not the storage view is active.

Storage volume properties

Storage Volume Name	The name of the storage volume. You can change the name in this field.
Use	Indicates whether the storage volume is Claimed, Unclaimed, or Used.
Used By	The name of the extent created from the storage volume.
Storage Volume Type	Describes the type of data stored on the storage volume: Data protected, Metadata or Normal
Integrated Services	The name of the supporting storage array.
Thin Rebuild	Indicates if a thin rebuild was selected for the storage volume.
System ID	The internal system ID for the storage volume.
Block Count	The number of blocks in the storage volume.
Block Size	The size of a single block in kilobytes.
Capacity	The size of the storage volume.
Thin Capable	Indicates if the storage volume is thinly-provisioned.
Storage Array Family	<p>The family of arrays that the supporting array belongs to. Possible values:</p> <ul style="list-style-type: none">• <code>XTREMIO</code> — The array belongs to the XtremIO family of arrays.• <code>CLARiion</code> — The array belongs to the VNX or Unity family of arrays.• <code>SYMMETRIX</code> — The array belongs to the VMAX or Symmetrix family of arrays.• <code>Other</code> — Metro node does not recognize the array, or does not support creating thin volumes on the array.• <code>-</code> (dash) — The Storage Array Family property is not populated. There may be a connectivity problem.
Largest Free Chunk	The total number of bytes in the largest range of blocks available for use on the storage volume.
Health	The overall health of the storage volume. See Health states .
Health Indications	The overall health of the storage volume. See Health states .
I/O Status	The status of I/O to the storage volume. See I/O states .
Operational Status	Indicates how the storage volume is functioning in the cluster. See Operational states .

Provision Job properties

Provision Job Name	The optional name given to the provisioning job. If you did not provide a name for the job, the base volume name is used.
Target Cluster	The cluster on which the job was initiated.
Status	The status of the job.
Status Details	An explanation of the job status.
Created	The date and time the job was created.
Last Updated	The last date and time the job status was automatically updated.
Description	A description of the volumes being created.
Volume Options	The options selected for the volumes.

Consistency Group	The name of the consistency group for the volumes (if any).
Total Capacity	The total capacity of the volumes being created.
Thin	Indicates if the volumes are thin.
Resources affected at <cluster>	<p>Storage View — The storage view selected for the volumes (if any).</p> <p>Storage Array — The storage array from which the volumes were provisioned.</p> <p>Storage Pool — The storage pool from which the volumes were created.</p> <p>Storage Group — The storage group automatically assigned or selected for the virtual volumes (if applicable).</p>

The Distributed Consistency Groups view

The Distributed Consistency Groups screen shows all distributed consistency groups. Note that distributed consistency groups are used in metro node Metro configurations only. Use this view to quickly view information about a distributed consistency group.


By default, groups are sorted by name in ascending alphabetical order. The arrow in the default sort column indicates the direction of the sort; ascending (▲) or descending (▼). You can [sort](#) columns of data, [select](#) one or more objects in the list, and [show or hide properties links](#) on the screen.

This view shows the following information about distributed consistency groups. To view additional properties, click the consistency group name to open the properties dialog box.

Column	Description
Name	The name of the consistency group.
Operational Status	The overall status of the consistency group. If the consistency group is distributed and status at both clusters is not the same, the status that is not OK displays. Click the status link for detailed information about the status. If the consistency group is distributed, the GUI displays the status at each cluster. Consistency group status provides more information.
Detach Rule	The detach rule that applies to the consistency group.
# Volumes	Indicates the number of volumes in the consistency group.

Additional tasks

From the **Distributed Consistency Groups** view, you can also:

- View properties for a consistency group — Click the consistency group name to open the properties dialog box.
- View properties of a virtual volume in the group — Click the virtual volume name in the **Components of Selected Consistency Group** panel on the right.
- Create a distributed consistency group — Click **Create** at the bottom of the page.
- Delete a group — Click **Delete** at the bottom of the screen.
- Rename a group — Click the consistency group name link to open its properties dialog box. Change the name the dialog box.
- Search for a group — Type the complete or partial group name in the **Find** text box, and then press **Enter**.
- Add or remove virtual volumes from the group — Click **Add/Remove Virtual Volumes** to open the **Add/Remove Virtual Volumes** dialog box.
- Refresh screen data — Click the Refresh icon .

Opening the Distributed Consistency Groups view

From the GUI main menu, hover over **Provision Storage** and select **Distributed > Distributed Consistency Groups**.

Create Virtual Volumes dialog box

Use this dialog box to create a virtual volume from an existing device. The Available Devices list shows all available top-level devices. Select a device from the list and then click **Add** to add the device to the Selected Devices list.

To create thin virtual volumes, select the **Create thin virtual volumes** checkbox and make sure that all selected devices are thin-capable.

NOTE: Selecting a mix of devices (thin capable and not thin capable) will result in creating thin virtual volumes on thin devices and thick virtual volumes on thick devices.

Platform view properties

Platform alert tab

Severity	<p>Indicates the urgency of alert:</p> <p>CRITICAL-A condition has occurred that can obstruct the functionality or can lead to failure of the other components.</p> <p>ERROR-An error has occurred that has a significant impact on the system and must be rectified immediately.</p> <p>WARNING-An error has occurred that you should be aware of but does not have a significant impact on the system. For example, a component is working, but its performance may not be optimum.</p> <p>INFO-An event has occurred that does not impact system functions. No action is required.</p> <p>CLEAR-When a corresponding clear event is generated for the raised alert, and the issue is resolved at the system level.</p>
State	<p>Represents the state of the alert-OPEN, CLOSED, or ACK state.</p> <p>OPEN-The alert is active and no action has been taken on it. The alert generating condition still persists in the system. If the system administrator wants to reopen to pay attention to that event, the alert status can also be changed to open from acknowledge state.</p> <ul style="list-style-type: none">To reopen an alert: Select the checkbox corresponding to the alert, and then click OPEN. <p>ACK-After you view an alert and understand its contents, you can acknowledge that you have read through the alert message.</p> <ul style="list-style-type: none">To acknowledge an alert: Select the checkbox corresponding to the alert, and then click ACKNOWLEDGE . <p>CLOSED-Once the problem condition is resolved, the event state is updated as closed. Once closed the alerts cannot be reopened.</p> <ul style="list-style-type: none">To close an alert: Select the checkbox corresponding to the alert, and then click CLOSE.
Message	Indicates the cause of an event for which the alert is generated.
Category	Displays the event category.
Enabled	Indicates whether the alert is enabled or not.
Alert ID	Represents the unique ID assigned to each alert .
Alert Name	Represents the name of the alert.
Description	Describes the platform alert.
Corrective Action	Action to eliminate the cause of event.
Scope	Represents whether it is a cluster-level or a director-level alert.
Condition ID	Indicates Unique ID of all defined alert definitions.
Component	System component that caused the event. Intended for service personnel.
Event Source	Represents the context of the event. For example, Virtual Volume.

Event Source ID	Unique ID for the source of the event. Helps to narrow down to the final component .
Resource	It shows the resource for which the issue has occurred.
Count	It represents the number of times the same alert is generated over the selected period. This column is available for historical alerts only.
Scope Incarnation	Represents unique ID assigned to each director and cluster.
Creation Date	Date and time when the alert is generated.
Last Updated	Date and time when the status of the alert is last changed.
External RCA	Represents external root cause analysis of the issue.
Notes	It shows the notes which are added by user.
Debug Events	Associated legacy debug events generated through firmware.

Additional details tab

To view additional information about the alert in name-value pair, click the Additional details tab.

Logs tab

It is for historical alerts only. It provides the list of creation and last updated time for the same alert has occurred.

Hardware view properties

Hardware alert tab

Severity	<p>Indicates the urgency of alert:</p> <p>CRITICAL-A condition has occurred that can obstruct the functionality or can lead to failure of the other components.</p> <p>ERROR-An error has occurred that has a significant impact on the system and must be rectified immediately.</p> <p>WARNING-An error has occurred that you should be aware of but does not have a significant impact on the system. For example, a component is working, but its performance may not be optimum.</p>
State	<p>Represents the state of the alert-OPEN, CLOSED, or ACK state.</p> <p>OPEN-The alert is active and no action has been taken on it. The alert generating condition still persists in the system. If the system administrator wants to reopen to pay attention to that event, the alert status can also be changed to open from acknowledge state.</p> <ul style="list-style-type: none"> • To reopen an alert: Select the checkbox corresponding to the alert, and then click OPEN. <p>ACK-After you view an alert and understand its contents, you can acknowledge that you have read through the alert message.</p> <ul style="list-style-type: none"> • To acknowledge an alert: Select the checkbox corresponding to the alert, and then click ACKNOWLEDGE . <p>CLOSED-Once the problem condition is resolved, the event state is updated as closed. Once closed the alerts cannot be reopened.</p> <ul style="list-style-type: none"> • To close an alert: Select the checkbox corresponding to the alert, and then click CLOSE.
Message	Indicates the cause of an event for which the alert is generated.
Enabled	Indicates whether the alert is enabled or not.
Alert ID	Represents the unique ID assigned to each alert .

Severity Code	Represents the numerical code for the corresponding severity of the alert.
Version	Represents version of the certificate.
Message ID	Indicates ID of the event record.
Host	Display the host details.
Host Name	Represents IP address or network name of the remote host.
AppName	Represents the device or application that originated the message.
Creation Date	Date and time when the alert is generated.
Last Updated	Date and time when the status of the alert is last changed.
Notes	It shows the notes which are added by user.

Logs tab

It is for historical alerts only. It provides the list of creation and last updated time for the same alert has occurred.

Configure alerts properties

Condition ID	Indicates Unique ID of all defined alert definitions.
Name	Name of the alert definition.
Enabled	Indicates if the alert definition is enabled or disabled.
Event Source	Defines the source of the alert definition.
Description	Explains about the defined alert definition.