Quantum

User's Guide

Lattus 3.3.X

6-68017-01 Rev A

Lattus User’s Guide, 6-68017-01 Rev A, October 2013, Product of USA.

Contents

Chapter 1

Introduction

Overview

This User's Guide includes the following topics:

- Product Overview on page 3: High-level overview of the Lattus systems and their components.
- Basic Operations on page 7: Instructions on how to perform basic operations such as logging into the Lattus CMC, shutting down the Lattus system/nodes, and powering on the Lattus system.
- Monitoring Lattus on page 17: Instructions on how to monitor your Lattus system.
- Troubleshooting Lattus on page 45: Instructions on how to identify and resolve common problems.
- Maintaining Lattus on page 53: Instructions on how to perform maintenance tasks to prevent problems.
- Replacing CRUs on page 57: Instructions on how to replace failed power supplies and decommissioned disks.
- Growing Lattus Capacity and Performance on page 69: Instructions for determining when you should contact Quantum for help with increasing your Lattus system's capacity and performance.

- Lattus Reference on page 89: Overview of the Lattus systems and their hardware, definitions of Lattus terms and concepts, information about licensing, and a list of available documentation and training resources.
- Getting Help on page 113: Instructions on how to contact Quantum Support for help and locate your system serial number, which is required for obtaining technical support.

Audience

This document is intended for users of the Quantum Lattus Object Storage system.

Notes, Cautions, and Warnings

The following formats used throughout this guide indicate important information:

Note: A Note emphasizes important information related to the main topic. There are no hazardous or damaging consequences.

Caution: A Caution indicates potential hazards to equipment or data. Failure to take or avoid this action could result in loss of data or harm to equipment.

WARNING: A Warning indicates potential hazards to personal safety. Failure to take or avoid this action could result in physical harm to the user or hardware.

Chapter 2
Product Overview

Quantum Lattus™ Object Storage is disk-based storage that meets the extreme scalability, durability, and access requirements of large-scale Big Data archives.

Lattus object storage offers several different access methods to put data in and get data out of the Lattus storage, including StorNext access, Lattus A10 Access for NFS/CIFS, and HTTP/HTTPS REST access for customers who have ported their applications to use the Lattus REST APIs.

A Lattus-X base system includes:

- 3 Quantum Lattus C10 Controller Nodes (one node acts as the Management Node)
- 6 or 20 Quantum Lattus S10 Storage Nodes
- 2 Quantum Lattus Rack Switches
- 1 Quantum Lattus A10 Access Node (enables communication between the hardware components and Lattus A10 software)

A Lattus-M base system includes:

- 3 Quantum Lattus C10 Controller Nodes (one node acts as the Management Node)
- 6 or 20 Quantum Lattus S10 Storage Nodes

- 6 or 20 Quantum Lattus-M Feature Keys for S10 Storage Nodes (equal to the number of Lattus S10 Storage Nodes)
- 2 Quantum Lattus Rack Switches

Required Add-On Components for StorNext Integration with Lattus

You must purchase the following add-on components separately to integrate StorNext with Lattus:

- **Lattus Interconnect Switch**: Lattus-M requires an interconnect switch to connect the StorNext MDCs/DDMs to the Lattus C10 Controller Nodes.
- **StorNext M662 Metadata Appliance**: The StorNext M662 Metadata Appliance (or a customer-supplied StorNext MDC) with Lattus-M Feature Keys provides StorNext access to the Lattus storage. According to the Storage Policies configured in StorNext, the StorNext MDC will use Lattus REST to store the appropriate files in the Lattus object storage.

Optional Add-On Components for Lattus

You can grow your Lattus system configuration to meet your needs as they expand. For example, you may want to migrate to a multi-geo setup, improve performance, or want better recovery characteristics.

To accomplish this, you may need to purchase one or more of the following add-on components for your Lattus system:

- **Lattus A10 Access Node**: Add for NAS file access to Lattus. Customers can create NFS/CIFS shares for storing their files on the Lattus A10 file system. The Lattus A10 Access Node will use Lattus REST to store those files in the Lattus object storage.
- **Lattus C10 Controller Nodes**: Add for more bandwidth into the system, additional MetaStore clusters, caching or availability.
- **Lattus S10 Storage Nodes**: Add for increased capacity, or possibly to enable performance, or increased data durability.
- **Lattus-M Feature Keys**: Add for StorNext integration with Lattus. Purchase for all or some of the Lattus S10 Storage Nodes to enable them to be used with StorNext.
- **Lattus Rack Switches**: Add when the components require another rack for the solution. There are always two Lattus Rack Switches in every Lattus rack.

- **Lattus System Switches**: Add when scaling above three racks or have more than six controllers in three racks.
- **Lattus Interconnect Switch**:
  - Lattus-M requires an interconnect switch to connect the StorNext MDCs/DDMs to the Lattus C10 Controller Nodes.
  - The interconnect switch is also recommended to connect the Lattus A10 Access Node(s) to the Lattus C10 Controller Nodes, as well as provide access to the Lattus C10 Controller Nodes for other HTTP/HTTPS REST protocol applications on the same private subnets.
- **Lattus Rack**: Add when your storage exceeds your current rack space.

As your needs change, your sales person can work with you to help determine how to grow your Lattus configuration and which add-on components may be needed.

Chapter 3

Basic Operations

This chapter covers the following basic operations that can be performed on your Lattus system.

Accessing the Lattus CMC

Many of the common tasks that you will perform on your Lattus system are accomplished through the Lattus Cloud Management Center (CMC), which is accessed by pointing a Web browser to the IP address of the Lattus C10 Controller Node that is acting as the management node (your Quantum Professional Services representative should have provided you with the IP address for the management node). Write it in the space below for future reference.

Management node IP address: _____

To access the Lattus CMC, follow these steps:

1. Enter the IP address of your Lattus C10 Controller Node that is acting as the management controller in a Web browser. The Lattus CMC login screen will appear.
2. In the Lattus CMC login screen, enter the Lattus administrator user name and password

---

Note: The default user name is **admin**, and the default password is also **admin**. If you changed the user name and/or password to something other than the defaults, write them in the spaces below for reference:

User name: _____

Password: _____

3. Press **Enter** or click **Connect**.

Activating Encryption

Caution: Please contact Quantum Support before enabling encryption for any new or existing namespace. The encryption master key(s) are stored on all three Lattus C10 Controller Nodes. Quantum recommends also keeping a backup of these key(s), which is necessary to avoid permanent and irreversible data-unavailability in the event that all three controllers are lost.

When creating or editing a namespace, you can choose whether the data written to the namespace should be encrypted. Lattus uses AES-256-CTR (256-bit Advanced Encryption Standard in CTR mode) to encrypt the data.

For details on namespace encryption, refer to Lattus Terms and Concepts on page 93.

Activating Encryption on a New Namespace

Follow the instructions in Creating New Namespaces on page 81 to launch the Add Namespace wizard.

1. Configure all of the options listed and choose **Yes** for the **Use Encryption?** option.

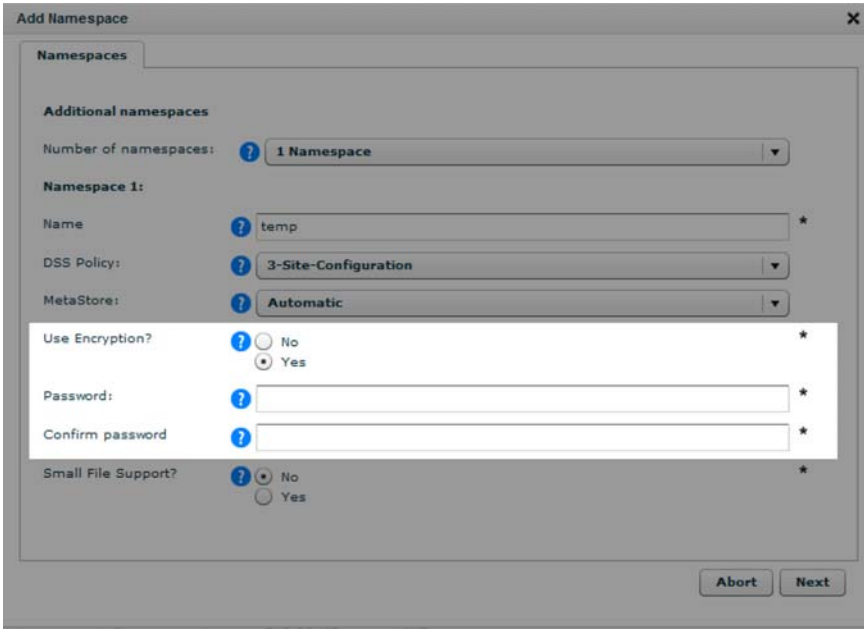

Note: If this is the first time you've activated encryption on your Lattus system, you will need to enter and confirm a password in the fields provided. You will only have to do this once. If you activate encryption for other namespaces, they will automatically use the same password.

Caution: Safeguard the entered password. This password will be used to generate a master key. This master key is used to encrypt the keys that are generated to encrypt the namespaces with. For each namespace that uses encryption, a new uniquely generated encryption key will be generated, from the same password.

2 Click **Next**.

3 In the dialog box, click **OK** to confirm your changes.

Activating Encryption on an Existing Namespace

1 Follow the instructions in Editing Existing Namespaces on page 86 to launch the **Edit Namespace** wizard.

2 Choose **Yes** for the **Use Encryption?** option.

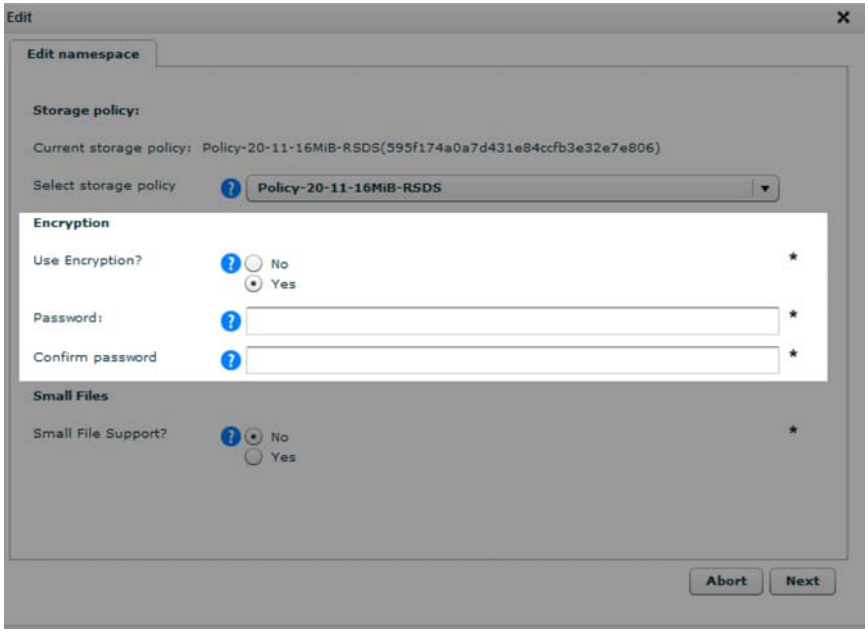

Note: If this is the first time you've activated encryption on your Lattus system, you will need to enter and confirm a password in the fields provided. You will only have to do this once. If you activate encryption for other namespaces, they will automatically use the same password.

Caution: Safeguard the entered password. This password will be used to generate a master key. This master key is used to encrypt the keys that are generated to encrypt the namespaces with. For each namespace that uses encryption, a new uniquely generated encryption key will be generated, from the same password.

3 Click **Next** to confirm your changes.

Powering On Your Lattus System

When a Lattus system is powered on, the Lattus services are started in a specific order. The power-on sequence will entirely quit if one service fails to start. This means that if a service is not running, it may not be because there is a problem with that service. Instead, it may indicate that some other service in the start order failed to start.

Note: If you observe any issues related to powering on the Lattus system, you should contact Quantum Support.

The power-on process is broken into two phases, followed by third phase that consists of monitoring the Lattus system for problems. Each phase and a short description is listed below:

- **Phase 1**

  Ensuring that all networking components (e.g., switches) are powered on, then powering on the Lattus C10 Controller Node configured as the management node, as well as all other controller nodes that participate in the object, framework, and env_metastore MetaStores.

- **Phase 2**

  Powering on any remaining Lattus C10 Controller Nodes and all Lattus S10 Storage Nodes

  As the Lattus S10 Storage Nodes power on, an event will be listed in the Lattus CMC stating that the machine was rebooted. If the machine had been gracefully shutdown, it will also raise events that its disks are offline. The management node will process these events and jobs will be created to activate these disks again.

- **Phase 3**

  The goal of this phase is simply to monitor the Live Events section of the Lattus CMC to identify any remaining issues that may need to be cleaned up. Things to look for are:

  - Lattus C10 Controller Nodes MetaStores are down.
  - Services such as the client daemons, storage daemons, and maintenance agents are down. This may indicate possible connection problems to the env_metastore.

Powering on the Lattus System

Follow these steps to power on the Lattus system:

1. Ensure that your switches and any other networking components are powered on, then power on the Lattus C10 Controller Node that was configured as the management node, as well as any other controller nodes that participate in the Lattus system's MetaStore cluster.
2. Power on any remaining Lattus C10 Controller Nodes as well as the Lattus S10 Storage Nodes.

   As the Lattus S10 Storage Nodes power on, an event will be listed in the Lattus CMC stating that the machine was rebooted. If the machine was gracefully shutdown, it will also raise events that its disks are offline. The management node will process these events, and jobs will be created to activate these disks again.
3. Monitor the Lattus CMC Dashboard for problems.

Caution: Contact Quantum Support if there are problems powering on the Lattus system.

Shutting Down Your Lattus System

Shutting down a Lattus system refers to shutting down all nodes in a single data center or rack, or shutting down individual nodes. A Lattus system, rack, or node shutdown should be issued through the Lattus CMC. Shutting down a machine or stopping all services is ordered internally. You do not need to take the order into account when selecting multiple nodes, racks or data centers.

Note: Typically, shutting down a node is for maintenance. During this time, your disk safety will drop, and any objects stored on this node will have a degraded durability. Objects will need to be retrieved from all nodes still in the system and it is not recommended that you shutdown more nodes than are configured for your safety, or objects will not be able to be retrieved until the required number of nodes is returned.

Caution: If a node is not functioning properly, contact Quantum Support and decommission the node for replacement.

Shutting Down a Data Center, Rack, or Individual Nodes

Follow these steps:

1. In the Lattus CMC, go to **Administration > Lattus Management > Locations > Datacenter Management** or **Rack Management**.
2. Select the desired data center or rack in the list.

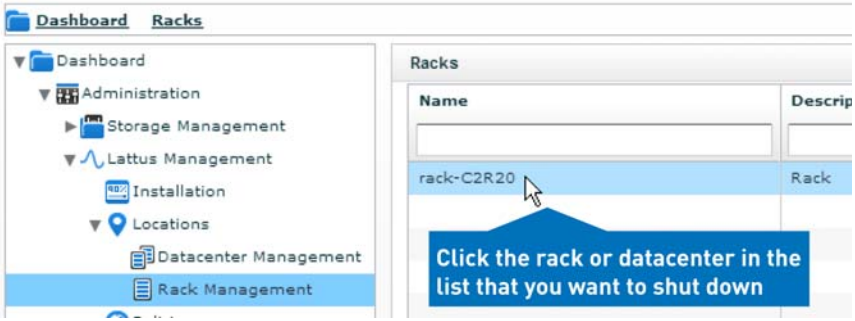

3. To shut down the selected data center or rack, click **Shutdown** in the **Commands** pane.

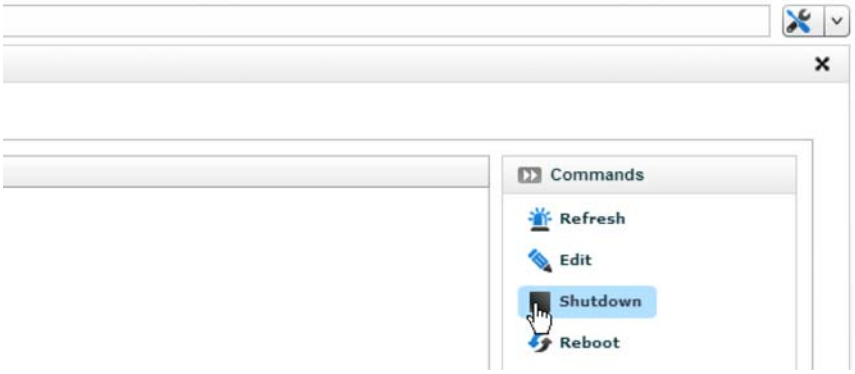

Shutting Down a Single Node

If you don't want to shut down an entire data center or rack and only want to shut down specific nodes, follow these steps:

1 Select the rack containing the nodes.

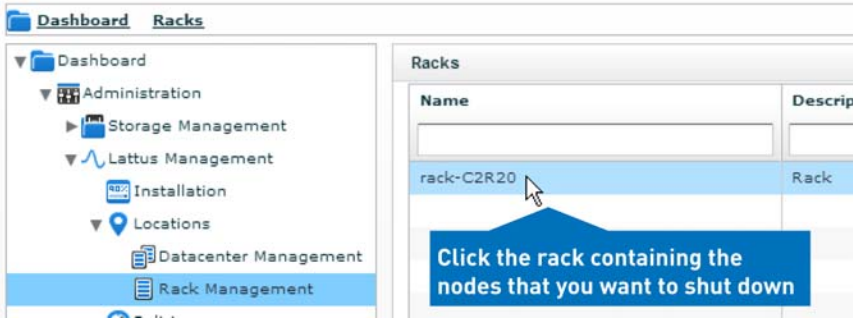

2 Select the node you want to shut down by clicking it in the **Machines** list.

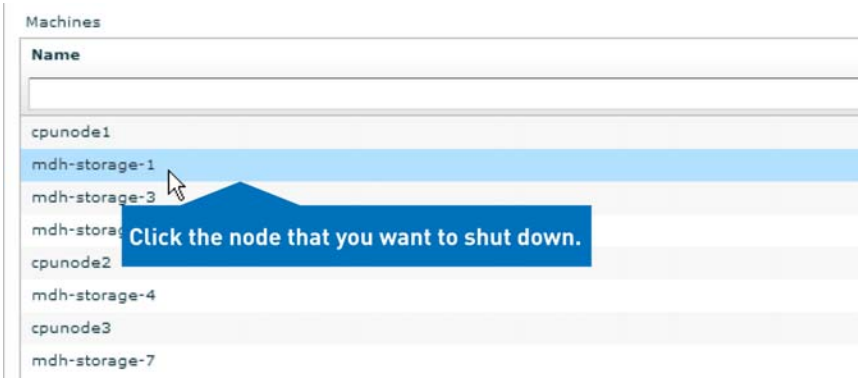

3 Click **Shutdown** in the **Commands** pane.

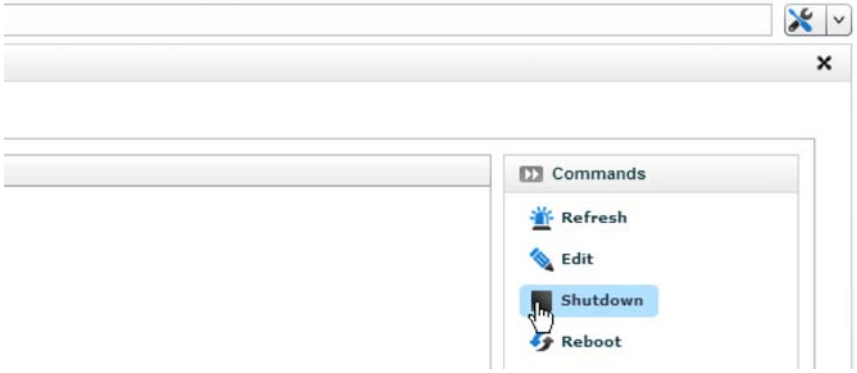

4 Repeat the previous three steps for each node you wish to shut down.

Note: If you're shutting down the whole system, but doing it one node at a time, it's recommended that you shut down the Lattus C10 Controller Node that's acting as the management node last. However, if you select an entire rack or data center to be shut down, the system takes care of this internally

Chapter 4
Monitoring Lattus

Knowing the health status of your Lattus system is necessary for ensuring that your data remains protected. To keep you updated on your Lattus system's health, Lattus provides status information through e-mail notifications and the Lattus CMC.

This chapter covers the following topics related to monitoring your Lattus system:

How the Lattus System Monitors Its Status

Each node in the environment has a monitoring agent. This monitoring agent monitors the node once per minute (a monitoring cycle). The resulting monitoring info will be updated in the local database. Only the modified objects (compared to the last cycle) will be saved in the local database. The Monitoring Policy of the Lattus C10 Controller Node configured as the management node will pull the monitoring info from all the local databases (including its own) and synchronize it with the data in the central monitoring database.

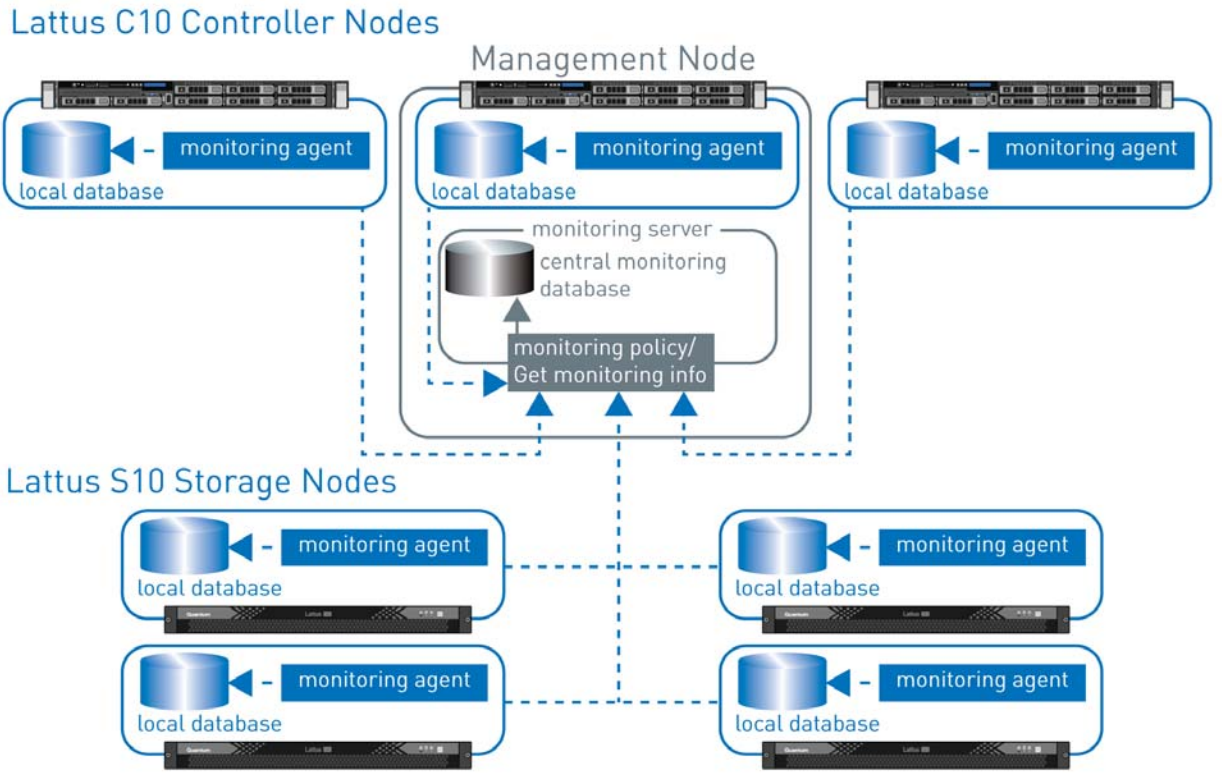

The Monitoring Policy also does the following:

- It generates information about the aggregated storage pool usage.
- It generates blacklists graphs.
- It checks the status of the remote nodes, monitoring agents and node agents.

The bulk of the time needed for this cycle is the update of the monitoring database, once the data is pulled from the node. Depending upon the load on the Lattus C10 Controller Node that has been configured as the management node, this can take seconds to a minute.

Each monitoring agent will take note of events happening on the node and send those events to the Lattus C10 Controller Node that has been configured as the management node. Depending on the configuration of the Monitoring Policies, the monitoring server will collect and escalate those when necessary. The monitoring policies will check the frequency of those events and send the necessary events to the configured email address.

Monitoring E-mail Notifications

The Lattus system uses the administrator e-mail address, which was configured during installation, to send real time event notifications and phone home notifications to communicate information about events on the Lattus system. These e-mail notifications are intended to help both you and Quantum Support monitor the Lattus system for potential problems.

Caution: If you receive notifications about events with severities of Error, Critical, or Unknown, these may indicate problems. Log in to the Lattus CMC and review the events in the Dashboard for more information, and contact Quantum Support for assistance.

Monitoring Real Time Notifications

Real time notifications are sent from your system's administrator e-mail account to the system's administrator e-mail account whenever certain events are detected within the system. Not all events trigger real time notifications; Lattus determines which events should trigger real-time event notifications.

For example, events like the Lattus storage pool reaching over 70% capacity, objects having low disk safeties, or power supplies failing will trigger real time event notifications.

However, several events, such as various services coming up, do not trigger real time notifications. Although they will still show up in the Lattus CMC Dashboard so you are aware of them.

Monitoring Phone Home Notifications

Phone home notifications are sent once every twenty-four hours and include a summary of all events of the specified severity that have occurred since the last phone home notification was sent. By default, the system only sends phone home notifications to Quantum Support. However, you can configure phone home notifications to be sent to additional recipients if necessary.

Monitoring Events

Everything that happens in the Lattus system is considered an “event”. Some events are informational, like a service coming up, while others may indicate problems, such as objects having low disk safeties.

Lattus assigns a severity to each type of event. The severity levels are (ranging from least to most severe):

Severity	Icon
Information	
Warning	
Error	
Urgent	
Critical	
Unknown	

All events that occur on the Lattus system will be displayed in the Lattus CMC Dashboard. Some events will also trigger real time email notifications, and events of Critical and Unknown severities will be included in phone home e-mail notifications that are sent once every 24 hours.

The following sections cover more about using e-mail notifications and the Lattus CMC Dashboard to monitor events happening on the Lattus system.

Monitoring Lattus from the Dashboard

The Dashboard is the first screen you'll see when you log in to the Lattus CMC, and it's the first place you should check if you receive e-mail notifications about problems. It provides statistical information related to the system, as well as alerts about potential problems.

Dashboard sections 1-8 with descriptions are shown below.

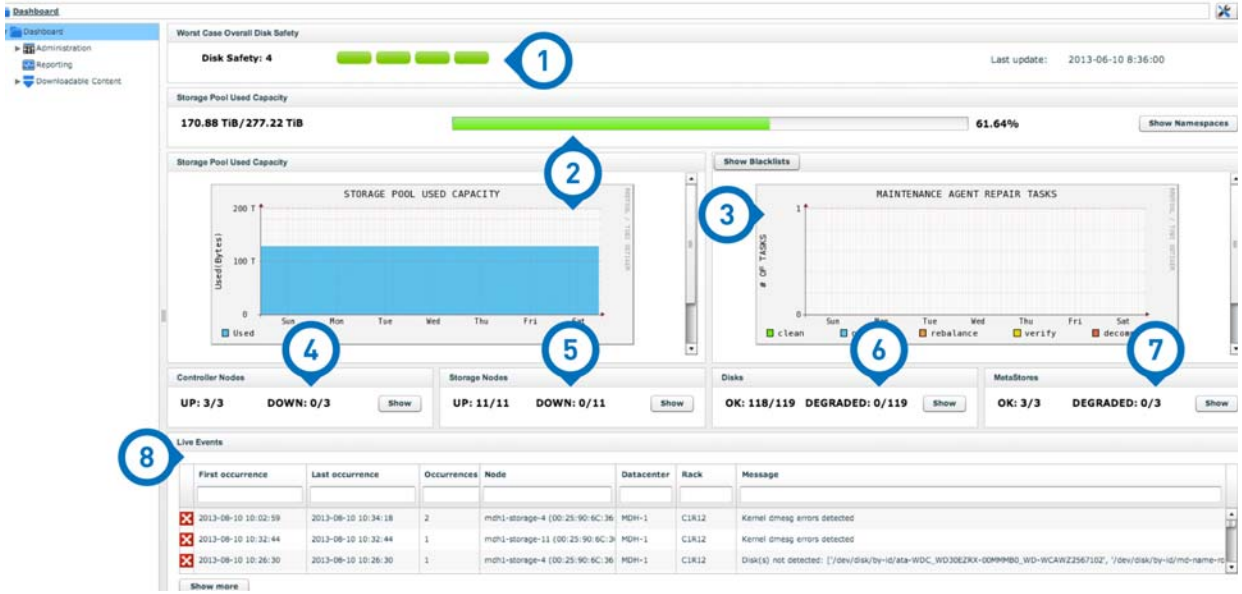

Description of Dashboard Sections

1 **Worst Case Overall Disk Safety**: Displays the current status of the object stored with the lowest disk safety. For example, if you have just a single 20/4 durability policy configured (spread width of 20 and a disk safety of 4) and all of the disks used to store the objects are working, the Worst Case Overall Disk Safety will be 4. But if one of the one of the disks used to store an object is degraded, the Worst Case Overall Disk Safety will be 3. (Quantum recommends a minimum disk safety of four.)

If the number and green dots displayed is less than four, you should check the Dashboard's **Disks** and **Live Events** sections for degraded disks, test any degraded disks (see Handling Degraded Disks in Lattus S10 Storage Nodes on page 50 for instructions), and decommission degraded disks if testing indicates they cannot be recovered (see Decommissioning Degraded Disks in the Lattus S10 Storage Node on page 53 for instructions).

Caution: If the Worst Case Overall Disk Safety drops below 2, or you have problems decommissioning your degraded disks, contact Quantum Support for assistance.

You can also view the Disk Safety Details for your individual durability policies to determine whether or not degraded disks in your Lattus system are affecting the disk safety of objects stored by specific policies. Since durability policies can use any of the available disks in the storage pool for spreading an object (with some exceptions depending on the policy's settings, such as whether the policy will disperse objects across machines/racks/data centers), it is likely that a degraded disk will affect some objects and not others.

Viewing Disk Safety Details

From the Lattus CMC, go to **Dashboard > Administration > Storage Management > Storage Policies**.

In the **Policy Management** screen, compare the **Storage Safety** (which displays the policies configured disk safety) and **Lowest disk safety** values for the policies. If the number in the **Lowest disk safety** column is less than the number in the **Storage Safety** column, this indicates that one or more objects stored by the policy are being affected by a problem, such as degraded disks.

2 **Storage Pool Used Capacity**: The status bar shows you how much of your Lattus system's capacity is currently being used out of the total capacity available. The graph shows the capacity used over time and can help you gauge how quickly the storage pool capacity is being filled.

- The system will raise WARNING events if the used capacity reaches 70%.

- The system will raise ERROR events if the used capacity reaches 80%.

  Caution: If this happens, Quantum recommends that you contact Quantum Support to purchase more Lattus S10 Storage Nodes.

- The system will raise CRITICAL events if the used capacity reaches 90%.

  Caution: It is imperative that you contact Quantum Support if your Lattus system's used capacity reaches above 90%.

3 **Maintenance Agent Repair Tasks/Blacklists**: The Maintenance Agents Repair Tasks graph (the default view of this graph) displays the number of repair tasks (e.g., decommissioning, rebalancing) that have occurred over the last week. Clicking the **Show Blacklists** button displays the number of failed reads and writes ("blacklists") to the object storage over the last week.

   Occasional blacklists are not a problem, but if the number of blacklists starts to increase, refer to the "Identifying Events that Require Action" section on page 20 for instructions on handling blacklists.

4 **Controller Nodes**: Provides status information for the Lattus C10 Controller Nodes. The ratios indicate the number of Lattus C10 Controller Nodes that are up and running (Up), and the number that have problems (**Down**) over the total number of nodes in the environment. For more details, click the Controller Nodes **Show** button.

5 **Storage Nodes**: Provides status information for the Lattus S10 Storage Nodes. The ratios indicate the number of Lattus S10 Storage Nodes that are up and running (Up), and the number that have problems (**Down**) over the total number of Lattus S10 Storage Nodes in the environment. For more details, click the Storage Nodes **Show** button.

6 **Disks**: Provides status information for both the storage disks in the Lattus S10 Storage Nodes, as well as the HDDs and SSDs in the Lattus C10 Controller Nodes. The ratios indicate the number of disks that are up and running (**OK**), and the number that have I/O errors (**Degraded**) over the total number of disks in the environment. The system regularly checks all of the disks and will mark a disk as degraded as soon as it detects any I/O errors. For more details, click the Disks **Show** button.

7 **MetaStores**: Provides status information for all of the MetaStores located on the Lattus C10 Controller Nodes. The ratios in the Dashboard indicate the number of MetaStores that have a full compliment of three participating SSDs (**OK**), and the number that have less than three participating SSDs (**Degraded**) over the total number of MetaStores in the environment.

Caution: If a MetaStore becomes degraded, contact Quantum Support for assistance.

For more details about your MetaStores, click the MetaStores **Show** button.

8 **Live Events**: Events are things that happen on the system that users should be aware of. Some events simply provide information, such as a service coming up, while others indicate problems, such as disk failures, or a network interface being down.

The **Live Events** section in the Dashboard displays only the events that are currently applicable. An event is 'applicable' as long as the monitoring agent detects it.

An event remains in the **Live Events** list:

- As long as the event is still applicable.
- For two occurrences of the monitoring interval if the event is no longer applicable.
- For 2,000 seconds if the monitoring interval is not specified and the event is no longer applicable.

To display an expanded list of events (up to 10,000), click the **Show More** button.

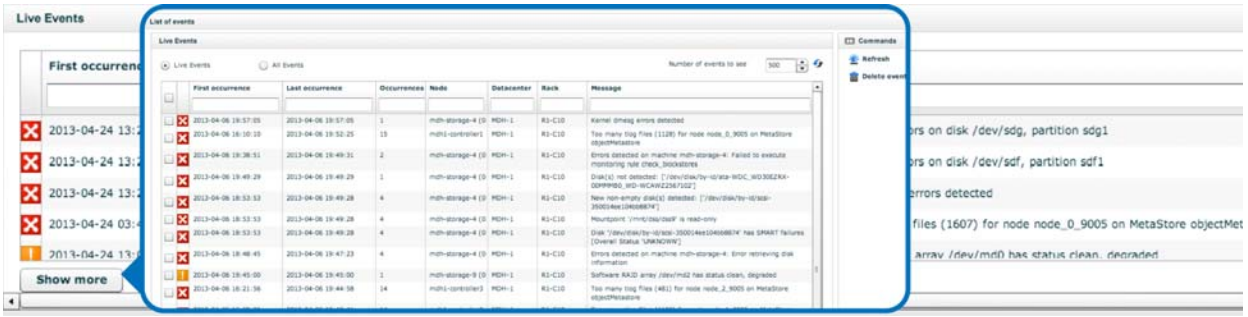

Check the Dashboard if you receive e-mail notifications about problems.

Caution: If you see ERROR, CRITICAL, and UNKNOWN severity events, contact Quantum Support for assistance. Instructions for contacting Quantum Support are located in Contacting Quantum Support on page 115.

Identifying Events that Require Action

Some events are simply informational and do not require any action. However, other events, especially those related to low disk safety or the used storage capacity being high, will require you to take action.

Caution: If you see any events that are of ERROR, CRITICAL, or UNKNOWN severity, you should contact Quantum Support for assistance.

The following list describes some of the common events you should watch for and the actions you should take if you see them:

Degraded Storage Disks and Lowered Disk Safety

If you see events related to lowered disk safety or degraded disks, check the Lattus CMC Dashboard. From there, you can see what your overall worst case disk safety is, as well display which of your disks are degraded.

Examples of Worst Case Overall Disk Safety displayed in the Lattus CMC:

- **Status**: Good. No objects have a disk safety less than 4.

- **Status**: Less than optimal. One or more objects have a disk safety of 2.

- **Status**: Critical. One or more objects have a disk safety of 0.

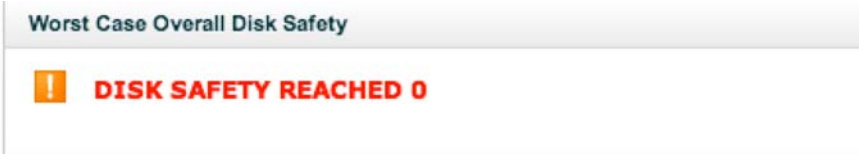

If you have degraded disks in your Lattus S10 Storage Nodes, you should follow the steps for testing them to determine whether they need to be decommissioned. (See Identifying Degraded Disks in Lattus S10 Storage Nodes on page 50 for details.)

If testing determines that the degraded disks should be decommissioned, decommission them as soon as possible to restore your disk safety and avoid potential data loss. (See Decommissioning Degraded Disks in the Lattus S10 Storage Node on page 53 for details.)

Blacklists

An increasing number of blacklists could also be an indication that one or more of your storage disks are degraded.

Caution: If you see an increasing number of blacklists in the Lattus CMC Dashboard Blacklists graph, contact Quantum Support and continue monitoring your Lattus system for degraded disks.

You can see how many blacklists have been recorded for the individual disks by looking at the Monitoring tab for a specific Lattus S10 Storage Node.

To do this, select **Dashboard > Administration > Hardware > Servers > Storage Nodes > [storage node name]**, then select the **Monitoring** tab. Blacklists are listed in the **DSS Blockstores** section.

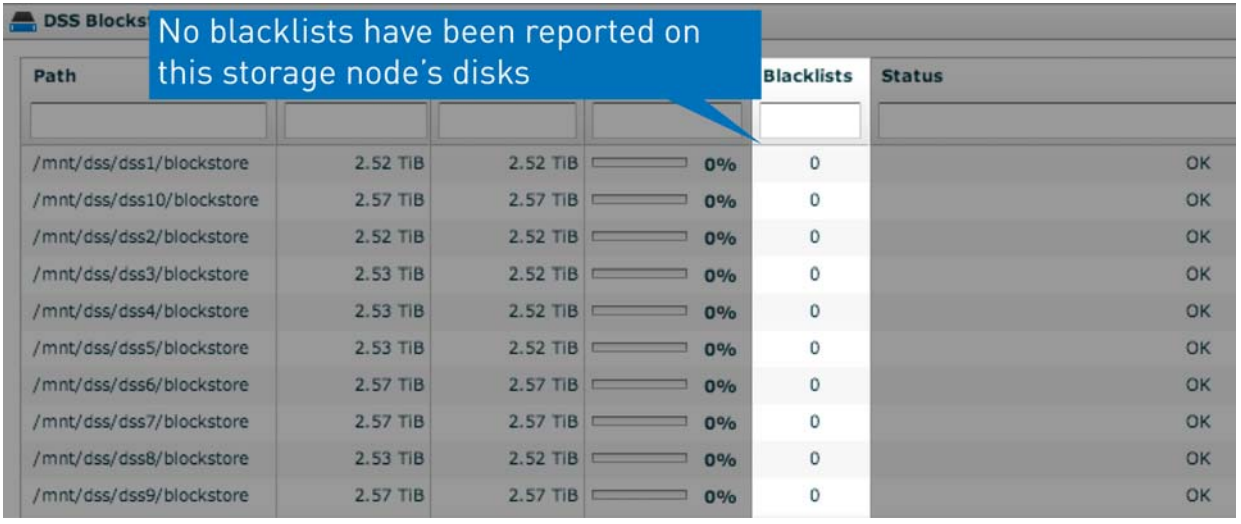

Path				Blacklists	Status
/mnt/dss/dss1/blockstore	2.52 TiB	2.52 TiB	0%	0	OK
/mnt/dss/dss10/blockstore	2.57 TiB	2.57 TiB	0%	0	OK
/mnt/dss/dss2/blockstore	2.52 TiB	2.52 TiB	0%	0	OK
/mnt/dss/dss3/blockstore	2.53 TiB	2.52 TiB	0%	0	OK
/mnt/dss/dss4/blockstore	2.53 TiB	2.52 TiB	0%	0	OK
/mnt/dss/dss5/blockstore	2.53 TiB	2.52 TiB	0%	0	OK
/mnt/dss/dss6/blockstore	2.57 TiB	2.57 TiB	0%	0	OK
/mnt/dss/dss7/blockstore	2.57 TiB	2.57 TiB	0%	0	OK
/mnt/dss/dss8/blockstore	2.53 TiB	2.52 TiB	0%	0	OK
/mnt/dss/dss9/blockstore	2.57 TiB	2.57 TiB	0%	0	OK

Storage Pool Used Capacity

As the Lattus object storage fills up, Lattus will display how much of the overall storage pool capacity has been used out of the total available capacity in the Lattus CMC's Storage Pool Used Capacity status bar and graph.

The event severities according to the used capacity are listed below:

- WARNING when 70% of the storage capacity is used. At this point, you may want to consider purchasing more Lattus S10 Storage Nodes.
- ERROR when 80% of the storage capacity is used. Contact your Quantum sales representative to discuss purchasing additional Lattus S10 Storage Nodes.
- CRITICAL when 90% of the storage capacity is used. Contact your Quantum sales representative.

The following example shows the Storage Pool Used Capacity for a Lattus system whose used capacity is over 90%.

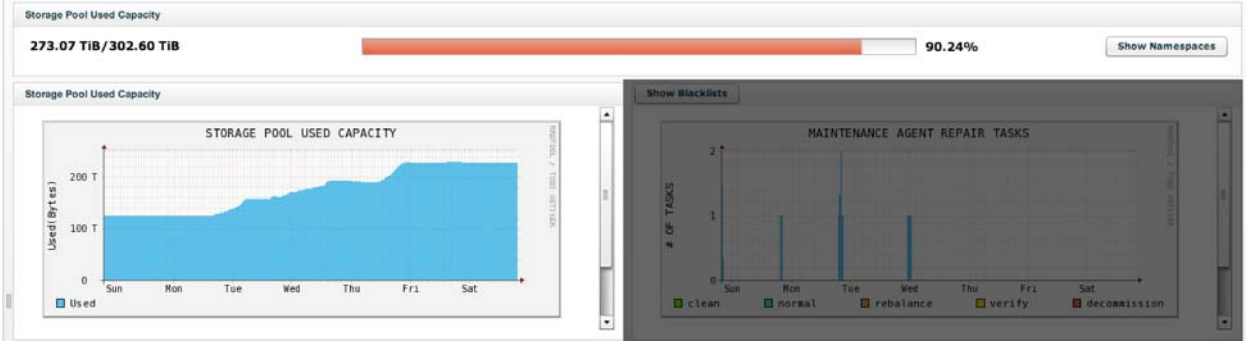

Blockstore/Disk Capacity

If individual disks/blockstores begin to fill up, this will also raise events.

For example, you will see WARNING events if a disk is 90% full, ERROR events if it reaches 96% full, and CRITICAL events if it reaches 98%.

Additionally, you will see events related to the number of Check-Blocks on a single S10 Drive:

- WARNING events at 10,000,000 Check-Blocks
- ERROR events 12,000,000 Check-Blocks
- CRITICAL events at 15,000,000 Check-Blocks

If you start seeing events related to blockstores/disks filling up, contact your Quantum sales representative.

MetaStore-Related Events

Degraded MetaStores

Each of your Lattus system's MetaStores consists of three SSDs on three distinct Lattus C10 Controller Nodes for high availability. The store is considered a “cluster”, each participating controller is considered a “node”, and each service is considered an “instance”.

To write and retrieve data from a MetaStore, you need a majority of the participating instances available. This means that, since a MetaStore spans three Lattus C10 Controller Nodes, two instances need to be running.

Caution: If you see events about degraded MetaStores, or you see in the Lattus CMC Dashboard that you have degraded MetaStores (see the example below), contact Quantum Support.

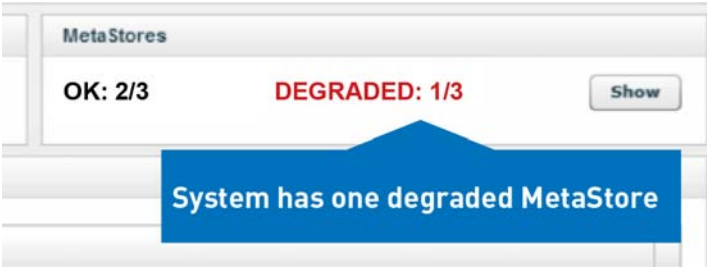

MetaStore Capacity

Every superblock written to the Lattus storage requires a metadata entry in a MetaStore. The maximum size of a MetaStore is 200 GB. As the amount of metadata stored in the MetaStore reaches certain thresholds, Lattus will begin to alert you through events like the one below:

Lattus-MON-ARAKOON-0010 Database partition for MetaStore node <metastore_name>:<node_name> is more than <x>% full WARNING / ERROR / CRITICAL

The following levels determine whether the event will be a warning, error, or critical:

- WARNING events at 95%
- ERROR events 96%

- CRITICAL events at 98%

If you want to see the total and used capacity for a MetaStore, follow these steps in the Lattus CMC:

1 Select **Dashboard > Administration > Storage Management > MetaStores**.

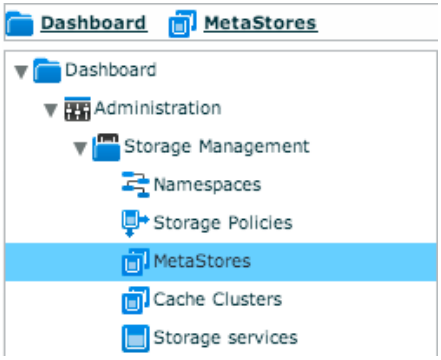

2 Select the MetaStore whose total and used capacity you would like to review.

3 In the **MetaStore: [MetaStore name]** screen, note the SSD on which the MetaStore is located in the **SSD Disk** column, then select one of the Lattus C10 Controller Nodes in the **Machine** column.

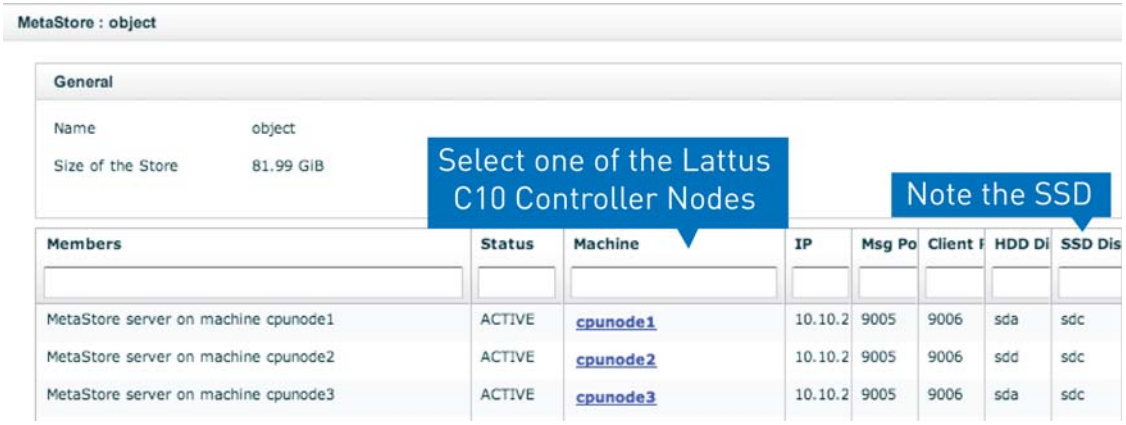

4 In the **Controller Node: [Controller node name]** screen, select the **Disks** tab, and then select the MetaStore's SSD.

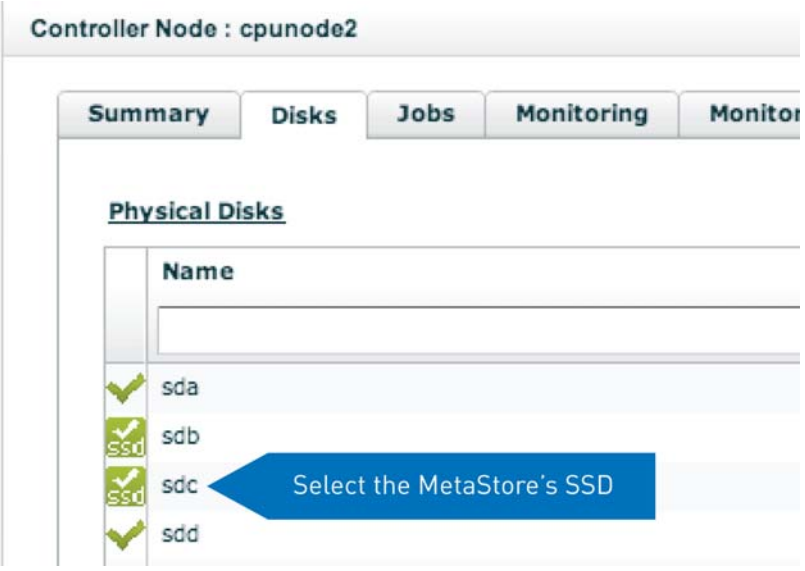

5 In the **Disk: [Disk name]** screen, select the **Monitoring** tab and scroll through the **Partition Usage** section until you see your MetaStore. The **Capacity (GB)** field displays the configured capacity

for the MetaStore in GB, and the **Used (GB)** field displays how much of the MetaStore's capacity has been used in GB.

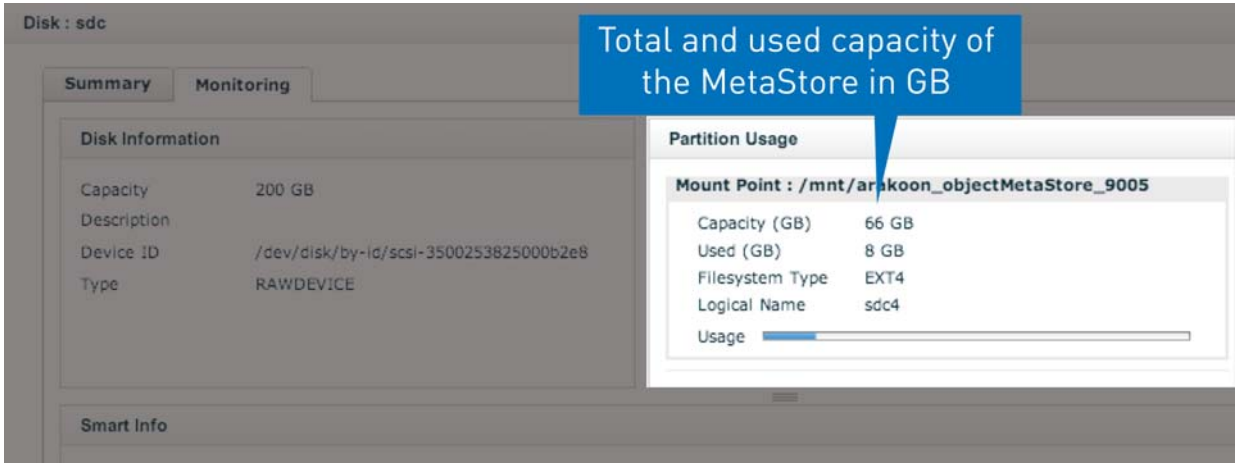

If you see events related to a MetaStore's capacity filling up, contact Quantum Support for assistance.

Monitoring Jobs, Events, and Policies

There are three Logging areas that you should check daily: **Jobs**, **Events**, and **Policies**.

- **Jobs**: Provides information about both automated and user-created workflows that have to happen within the environment.
- **Events**: Provides an extended list of event history for the system. This is the same list that's displayed if you click **Show All** under the Dashboard's **Events** section.
- **Policies**: Displays what the system's policies are doing and whether any of them are in error Policies are specific workflows that the system executes to make sure it is in an optimal state.

To access the **Logging** screens, follow these steps:

1 In the Lattus CMC, go to **Dashboard > Administration > Lattus Management > Logging**.

2 Depending on what logging information you want to see, select **Jobs**, **Events**, or **Policies**.

> **Note:** Clicking **Events** displays the same list of events clicking the Show More button under Live Events on the Dashboard.

3 In the list that appears, click one of the jobs, events, or policy jobs listed to see more details about it.

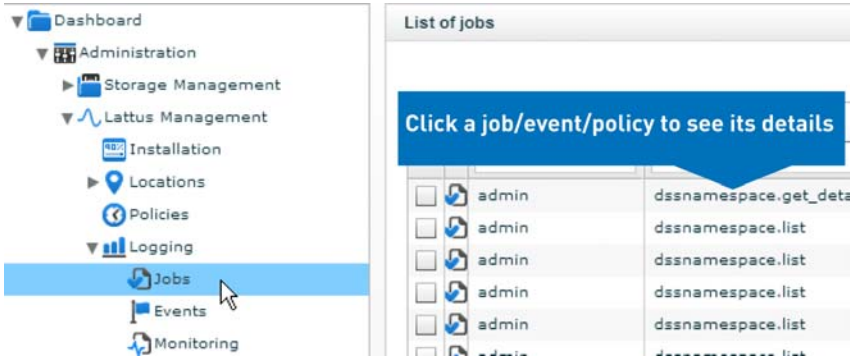

If you see failed jobs and policies listed, contact Quantum Support for assistance.

Monitoring Hardware Status

You don't need to check the hardware status regularly, but you should check it you see other indications of failures, such as a node that's down, or events listed in the Dashboard.

To review information about your system's status, follow these steps:

1 In the Lattus CMC, go to **Dashboard > Administration > Hardware**.

2 To view information about your Lattus C10 Controller Nodes or your Lattus S10 Storage Nodes, click **Servers**.

   a To view information about your Lattus C10 Controller Nodes, click Controller Nodes. Then click the individual nodes that are listed to see more information.

b To view information about your Lattus S10 Storage Nodes, click **Storage Nodes**. Then click the individual nodes that are listed to see more information.

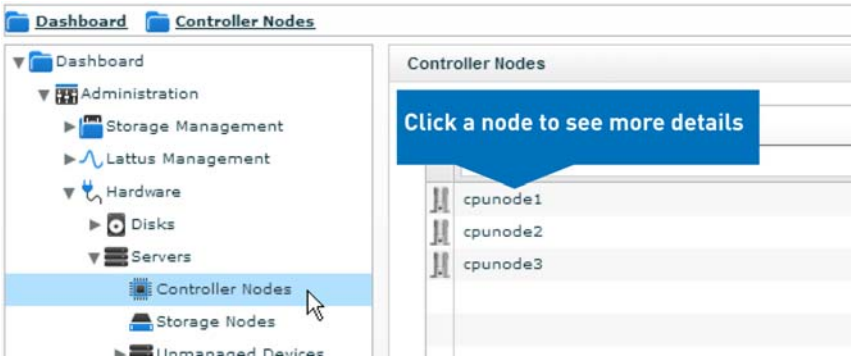

3 To view information about specific degraded or decommissioned disks, under **Hardware**, click **Disks**.

a To view information about degraded disks, click **Degraded**. Then click the individual disks that are listed to see more information. If no disks are currently degraded, none will be listed.

b To view information about decommissioned disks, click **Decommissioned**. Then click the individual disks that are listed to see more information. If no disks are currently decommissioned, none will be listed.

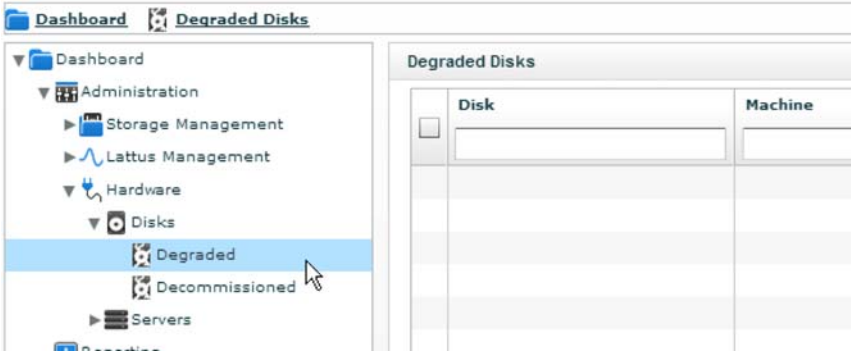

Displaying Storage Device Statistics

The system's reporting looks at the statistical records of every storage service in the environment, and then displays that information graphically over time.

To review reports on the processes running on the system's nodes, follow these steps in the Lattus CMC:

1 Click **Reporting**. The Reporting screen appears.

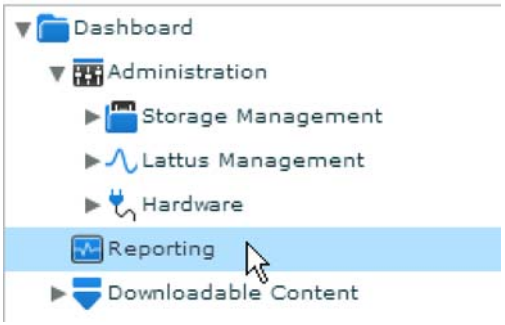

2 In the left pane, select the process whose information you want to review: Client **Daemons**, **Storage Daemons**, **Maintenance Agents**, or **MetaStores**. (See Lattus Terms and Concepts on page 93 for explanations of these processes.)

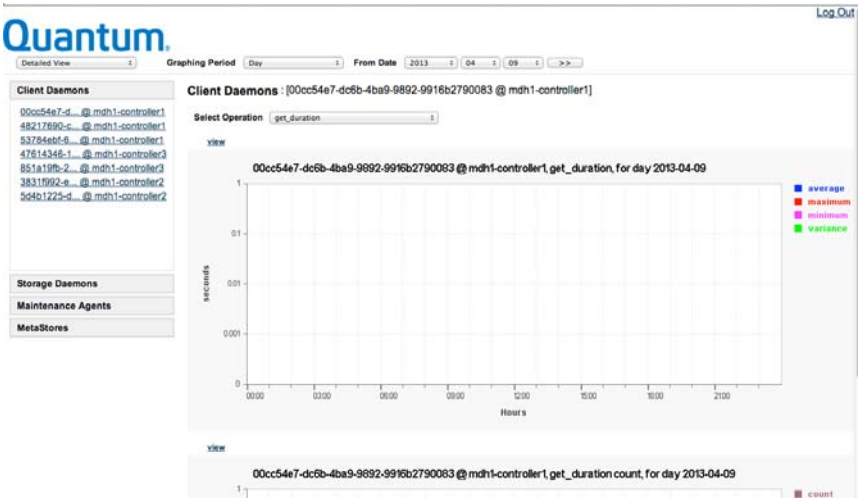

3 Use the **Graphing Period** and **From Date** drop-down lists to select the time frame whose data you want to see.

4 Use the **Select Operation** drop-down to select the operation whose statistics you want to see.

5 The default view is **Overview**. To see a more detailed report of the selected processes, select **Detailed View** in the upper left corner of the screen.

Note: Reporting is delayed, so the statistics displayed are not up-to-the-minute.

Changing E-mail Notification Settings

If you need to change the administrator e-mail address used by the system for sending notifications (and receiving real time notifications), or you need to add recipients for phone home notifications, you can easily make these changes using the Lattus CMC.

To see a video tutorial on changing Lattus e-mail notifications, go to http://www.quantum.com/lattushowtos.

Changing the Administrator E-mail Address

The administrator e-mail address sends all of the system's e-mail notifications, and receives all real time notifications.

To change the administrator e-mail address, follow these steps:

1. From the Lattus CMC, go to **Dashboard > Administration > Lattus Management > Groups and Users > Users**.
2. In the **Users Overview** screen, click **admin**. The **Clouduser: admin** screen appears.

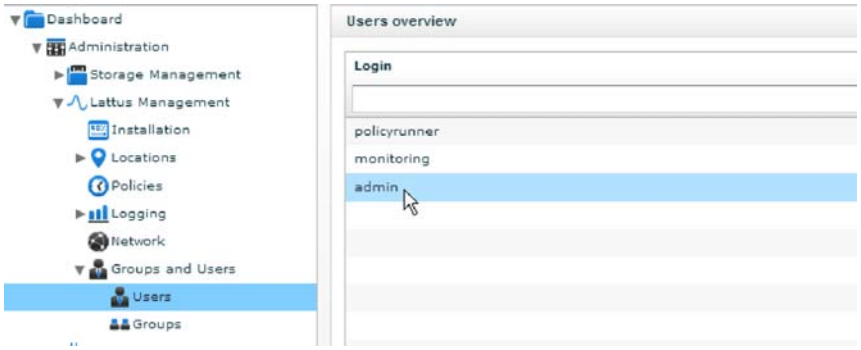

3. In the **Commands** pane, click **Edit**. The **Edit clouduser** window appears.

4 Select the **Information** tab.

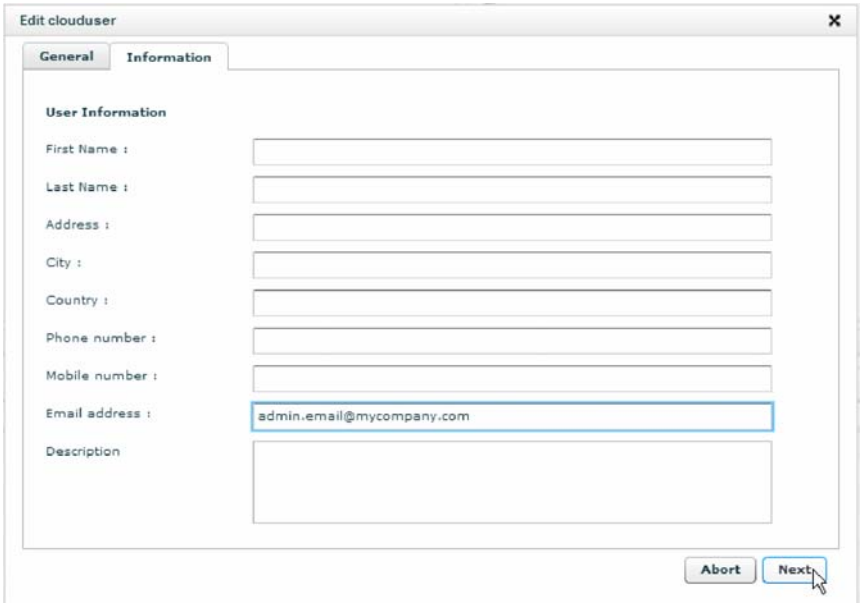

5 Enter the new administrator e-mail address in the **E-mail Address** text box.

6 Click **Next**.

7 In the dialog box, click **Yes**.

Your Lattus system will now use the new administrator e-mail address to send all future notifications.

Caution: DO NOT leave the administrator e-mail address blank. This will prevent you from receiving e-mail notifications from your Lattus system.

Changing Phone Home Notification Settings

This section describes the following tasks:

- Adding Recipients on page 40
- Changing When Phone Home Notifications are Sent on page 43

Caution: It is REQUIRED that phone home notifications be sent to Quantum Support (techsup@quantum.com). This is how phone home notifications were configured during initial installation and configuration of your Lattus system. You may add other recipients to receive phone home notifications (such as system administrators) by using a distribution list or e-mail alias, however, your distribution list/e-mail alias MUST include Quantum Support.

Removing Quantum Support from phone home notifications will prevent Quantum from being able to proactively monitor your Lattus system for problems. Additionally, your Lattus system's phone home notifications were configured to only include events of the severity level CRITICAL. This ensures that the phone home notifications only list events related to potential problems, and don't include events that are simply informational.

Adding Recipients

By default, your Lattus system was configured to send its phone home notifications to Quantum Support, which is a requirement.

If you need the system to send phone home notifications to others, such as system administrators, doing the following:

1 In the Lattus CMC, go to **Dashboard > Administration > Lattus Management > Policies**. The Policy Overview screen appears.

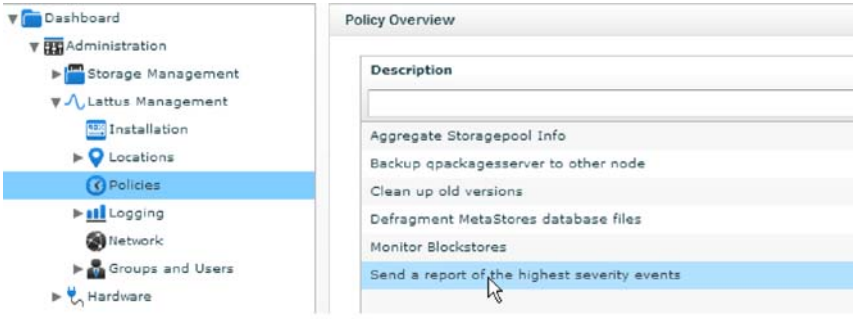

2 Click **Send a report of the highest severity events**. The **Policy: phone_home** screen appears.

3 In the **Commands** pane, click **Edit Policy**. The **Edit Policy** window appears.

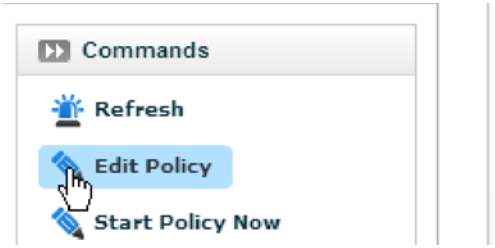

4 In the **E-mail** text box, enter the distribution list or e-mail alias address that should receive phone home notifications.

Caution: You can only enter one e-mail address, so to add recipients you must create a distribution list or e-mail alias address that includes Quantum Support (techsup@quantum.com). This is required to ensure that Quantum Support can proactively monitor your Lattus system.

5 Ensure that the **maxErrorconditionlevel** drop-down is set to **Critical**. This is the required severity level for phone home notifications sent to Quantum Support.

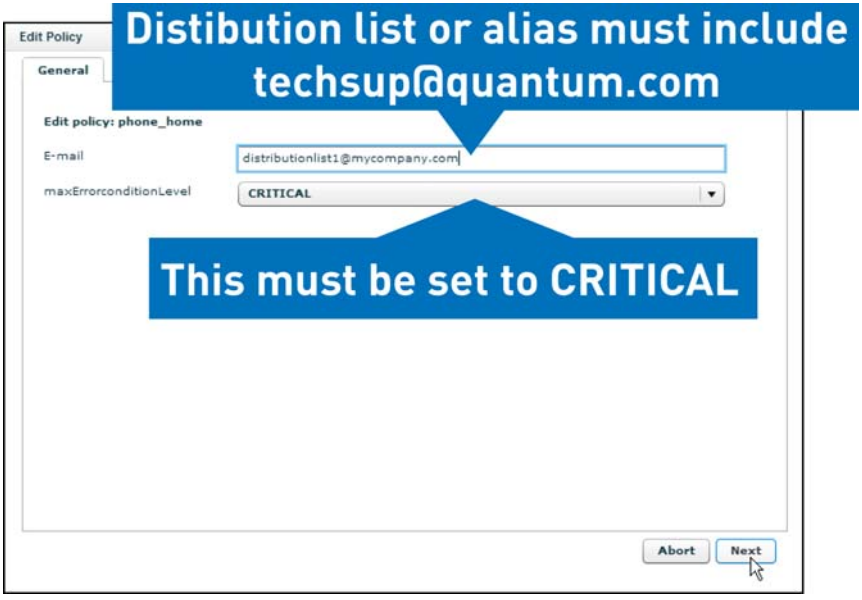

6 After making your changes, click **Next**.

7 In the dialog box, click **OK**.

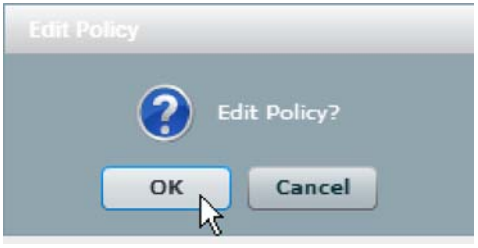

8 To test your phone home notification settings, click **Start Policy Now** in **Commands** pane.

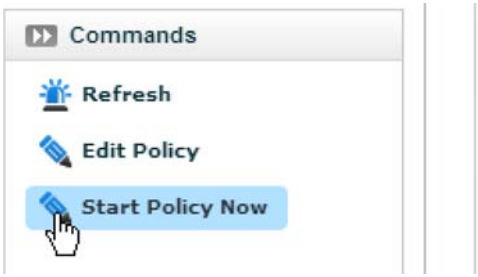

The system will immediately send a phone home notification to the recipient e-mail address.

One phone home notification will be sent every twenty-four hours starting from the time you click **Start Policy Now**. Each phone home notification will include all critical events that have occurred since the last phone home notification was sent.

Changing When Phone Home Notifications are Sent

If you wish to send/receive phone home notifications at a different time, follow these steps at the time when you want phone home notifications sent:

1 In the Lattus CMC, go to **Dashboard > Administration > Lattus Management > Policies**. The Policy Overview screen appears.

2 Click Send a report of the **highest severity events**. The **Policy: phone_home** screen appears.

3 Click **Start Policy Now** in Commands pane.

The system will immediately send a phone home notification. The system will continue to send one phone home notification every twenty-four hours starting from the time you click **Start Policy Now**.

Chapter 5

Troubleshooting Lattus

This chapter covers the following common Lattus troubleshooting tasks:

Checking Events in the Lattus CMC

From the Lattus CMC Dashboard, you can see the most recent events under the **Live Events** section. Click the **Show More** button to view an expanded history of events (up to 10,000).

Events are divided into the following severity categories with corresponding icons (see below).

Severity	Icon
Information	
Warning	
Error	
Urgent	
Critical	
Unknown	

Click an event in the Lattus CMC to view its details. (See Monitoring Lattus on page 17 for more details on monitoring Lattus for problems.)

Checking the LEDs on the Lattus S10 Storage Nodes

There are three LEDs on the front panel of each Lattus S10 Storage Node. The LEDs from left to right are:

- **Location LED**: Emits can be turned on through the Lattus CMC, making it easier to identify and replace failed nodes.
- **Disk Activity LED**: Emits when disk activity is occurring.
- **Power Connection LED**: Emits when the power is connected.

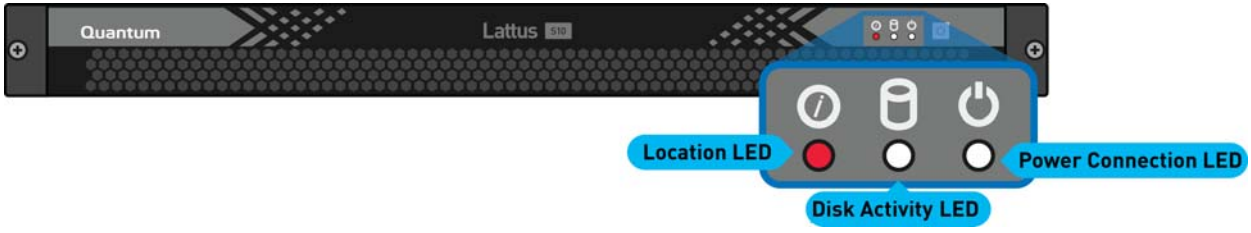

Expanding an Existing MetaStore

If the Dashboard indicates that a MetaStore has become degraded (see graphic below), it needs to be expanded. Contact your Quantum sales representative for more information about expanding.

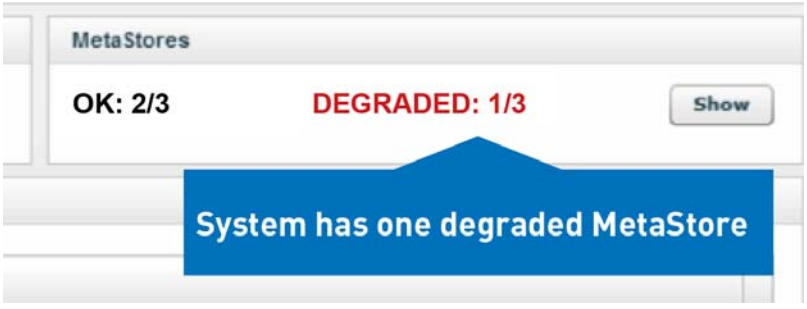

Expanding a MetaStore will require replacement of the failed SSD in the Lattus C10 Controller Node. Once the SSD has been replaced, the MetaStore can be expanded using the following steps:

1. In the Lattus CMC, go to **Dashboard > Administration > Storage Management > MetaStores.**
2. In the **MetaStores** window, click on the desired MetaStore.

3 In the **MetaStore** window, click **Add Disks** in the **Commands** pane.

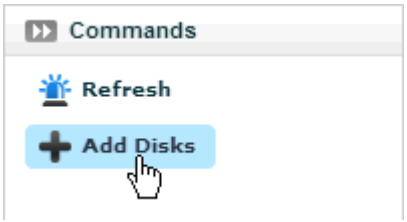

4 In the **Add disks** window, select the new SSD(s) to add to the MetaStore.

5 When finished, click **Refresh** in the **MetaStores** screen's **Commands** pane.

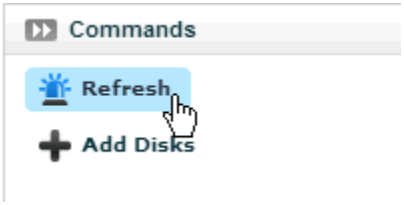

Handling Failed CRUs

Each Lattus S10 Storage Node contains the following customer replaceable units (CRUs):

- 12 hard disk drives
- 2 power supplies (referred to as “PSUs” in events and notifications)

The following instructions cover how to identify and handle failed CRUs.

Identifying Failed Power Supplies in Lattus S10 Storage Nodes

A failed power supply in a Lattus S10 Storage Node will cause an event to be displayed in the Dashboard of the Lattus CMC, as well as a real time notification e-mail to alert you about the problem.

- The event ID for a failed power supply is: ***Lattus-MON-PMACHINE-0115***
- The message will read: ***PSU is not functioning properly***

The failed power supply's status LED will also flash amber, making it easy to locate it when in your data center. The illustration below shows examples of the status LEDs for a working power supply and a failed power supply.

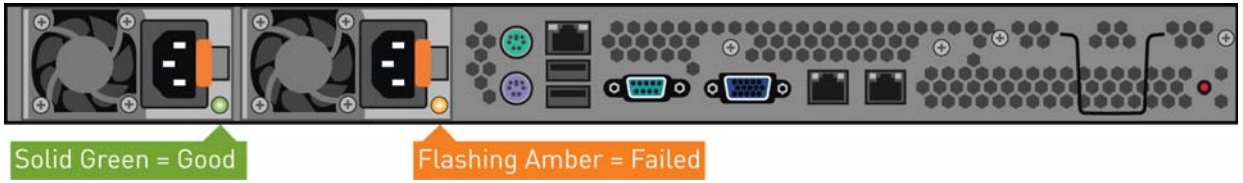

Handling Failed Power Supplies in Lattus S10 Storage Nodes

Since each Lattus S10 Storage Node has two power supplies for redundancy, the node will still have power even if one of its power supplies fails. However, if the second power supply fails, the node will lose power completely and may impact the disk safety for some of the objects in your Lattus storage.

If you receive e-mail notifications or see events in the Lattus CMC related to a failed power supply, you should replace it as soon as possible.

See Replacing the Lattus S10 Storage Node Power Supplies on page 60 for instructions on how to two replace the power supplies.

Identifying Degraded Disks in Lattus S10 Storage Nodes

The **Disks** section of the Lattus CMC Dashboard will show you if the system has degraded disks. (See the following image).

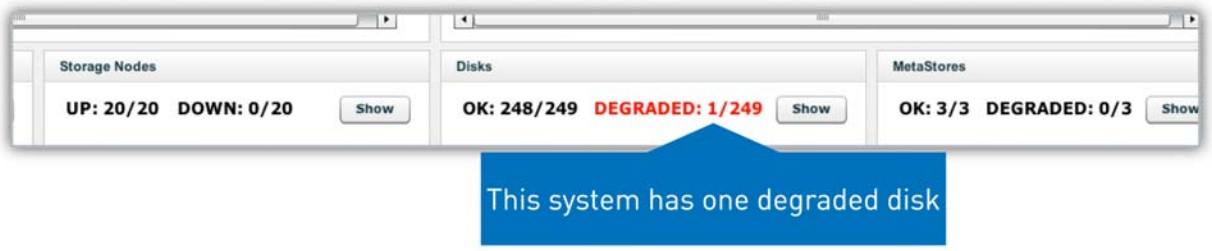

Handling Degraded Disks in Lattus S10 Storage Nodes

If you get a warning about a degraded disk, follow these steps in the Lattus CMC:

1 Navigate to **Dashboard > Administration >Hardware > Disks > Degraded**.

2 Check the number of degraded disks:

a If multiple disks are degraded on the same node, follow the instructions for Handling Multiple Degraded Disks in One Lattus S10 Storage Node on page 51.

b If only one disk per node is degraded, follow the instructions for Handling One Degraded Disk in One Lattus S10 Storage Node on page 51.

Note: Notifications about failed disks will also be made via SNMP and the Phone Home function.

Handling Multiple Degraded Disks in One Lattus S10 Storage Node

If you have multiple degraded disks in the same Lattus S10 Storage Node, proceed as follows:

1. Cold Reboot the Lattus S10 Storage Node. (Power down and wait at least 10 seconds before powering up the node)
2. Once the node has successfully rebooted, in the Lattus CMC, navigate to: **Dashboard > Administration > Hardware > Disks > Degraded**.
3. In the degraded disks list, select all degraded disks.
4. Click **Reset** in the **Commands** pane. After the reset is complete, the disks should disappear from the degraded disks list.

5. If the disks did not disappear, contact Quantum Support for assistance.

Handling One Degraded Disk in One Lattus S10 Storage Node

If you have one degraded disk in one Lattus S10 Storage Node, follow these steps to check the disk:

1. In the **Degraded Disks** screen, click **Diagnose** in the **Action** column for the degraded disk.

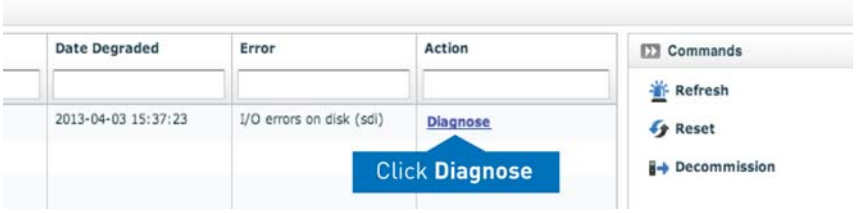

2 Select both the **SMART** and **Read/Write** tests.

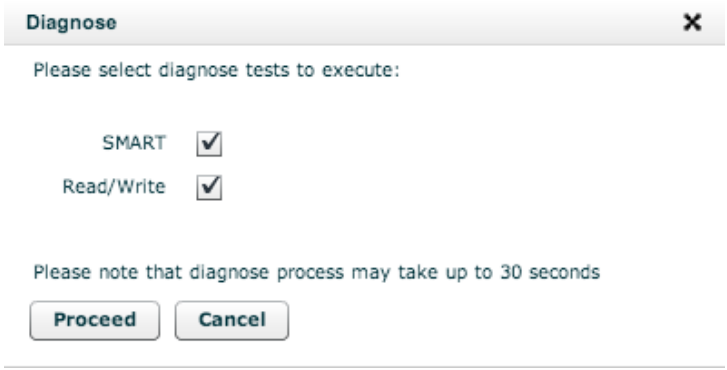

3 Click **Proceed** to run the tests. When the tests are complete, their results will be displayed.

4 Check the test results.

   a If both results are a success, proceed to step 5.

   b If the SMART or Read/Write test failed, decommission the disk. Refer to Decommissioning Degraded Disks in the Lattus S10 Storage Node on page 53 for instructions.

5 If both tests were successful, proceed as follows:

   a Note the number of the disk somewhere for future reference in case it becomes degraded again. If it becomes degraded again, you will need to decommission it.

   b f this is the first time the disk was degraded, click **Reset** in the right-hand window. The disk will disappear from the degraded disks list.

   c If this is not the first time the disk was degraded, decommission the disk. Refer to Decommissioning Degraded Disks in the Lattus S10 Storage Node on page 53 for instructions.

Chapter 6

Maintaining Lattus

The key maintenance task for Lattus is regularly decommissioning degraded disks on the Lattus S10 Storage Nodes to ensure the highest level of durability for your data.

Decommissioning Degraded Disks in the Lattus S10 Storage Node

This section describes how to decommission a disk after testing a degraded disk indicates that it needs to be decommissioned. (Refer to Handling Degraded Disks in Lattus S10 Storage Nodes on page 50 for details.)

Before You Begin

If a degraded disk contains partitions that are members of MD devices (RAID), you may receive the following event messages after decommissioning the disk:

```
Load average over the last 15 minutes is high
Software RAID array /dev/md2 has a status degraded
```

To avoid receiving these message about the degraded disk, be prepared to replace the disk at the time of decommissioning.

Use the following command to view the partition table:

```
# /sbin/parted /dev/sdx print
```

Decommissioning a Degraded Disk

Follow these steps to decommission a degraded disk:

1. In the Lattus CMC, go to **Dashboard > Administration > Hardware > Disks > Degraded**.
2. Select the degraded disk and click **Decommission** in the **Commands** pane.

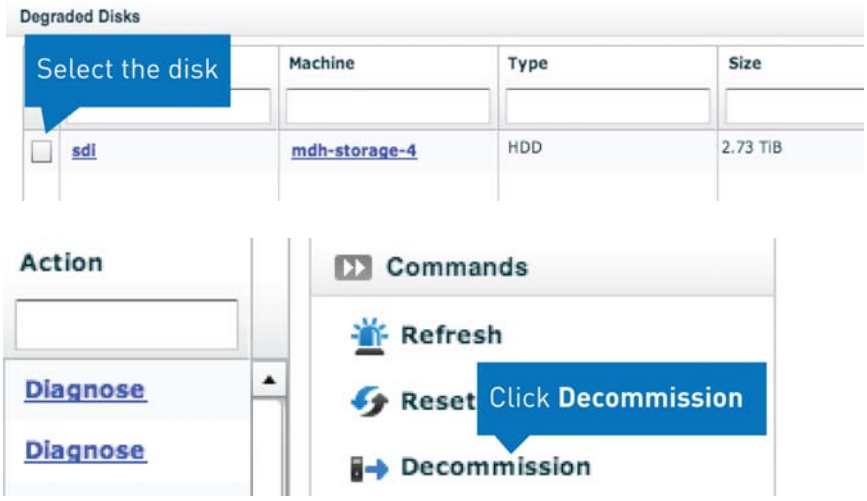

3. Choose whether or not you want to erase the disk (secure erase following NIST SP 800-88).

   Note: Secure erase can take several hours, depending on the amount of data on the disk.

4. Choose whether or not you want to halt the machine (Note: Do this only if this is the only disk and if you'll replace it immediately).

   Once the wizard is completed, the disk will be automatically moved to the Decommissioned disk section of the Lattus CMC. When the

disk gains the status “decommissioned”, a repair crawl will start for all namespaces.

Note: The repair crawl is an iteration of all objects in a namespace. The crawl will cause repair tasks to be created and put into a queue for the maintenance agents to reserve and execute.

The crawl and repair queue are managed by a storage daemon on a storage node. This will ensure Lattus starts repair activities for objects that were affected by the decommissioning immediately, rather than waiting until the next 24-hour crawl begins.

5 It is not necessary to replace decommissioned disks immediately, as decommissioning will repair the objects that were affected by the degraded disk. Replacements for your decommissioned disks will be shipped to you when required.

If you have problems decommissioning your degraded disks, contact Quantum Support for assistance.

Chapter 7
Replacing CRUs

Each Lattus S10 Storage Node contains the following customer replaceable units (CRUs):

- 12 hard disk drives
- 2 power supplies

The instructions in this chapter cover:

WARNING: Opening or removing the Lattus S10 Storage Node cover while it is powered on may expose you to a risk of electric shock.

Caution: When replacing items from the inside of the chassis, ensure that you take precautions to prevent Electrostatic Discharge (ESD).

Locating Lattus S10 Storage Nodes

When you're ready to replace CRUs in your Lattus S10 Storage Nodes, you can physically locate the storage nodes in your data center by doing the following in the Lattus CMC to illuminate the nodes' Location LEDs:

1. Determine which Lattus S10 Storage Node contains the failed CRU.
   a. If a power supply has failed, the **Live Events** section of the Dashboard will list the name of the node containing the failed power supply the **Node** column.
   b. If you have decommissioned drives that need to be replaced, navigate to **Dashboard > Administration > Hardware > Disks > Decommissioned**. The names of the nodes with decommissioned disks will be listed in the **Machine** column.
2. Navigate to the **Storage Nodes:[storage node name]** screen.
   a. If you're trying to locate a Lattus S10 Storage Node that contains decommissioned disks, just click the name of the node in the **Decommissioned Disks** screen's **Machine** column.
   b. If you need to turn on the **Location LED** for a Lattus S10 Storage Node that contains a failed power supply, select **Dashboard > Administration > Hardware > Servers > Storage Nodes** to access the **Storage Nodes** screen, then click the Lattus S10 Storage Node that you would like to locate.

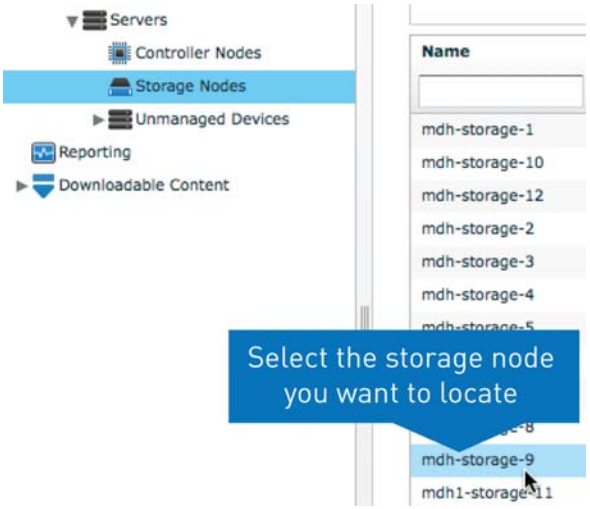

3 In the **Commands** pane of the **Storage Node: [storage node name]** screen, select the **Location LED On** radio button. This will illuminate the Lattus S10 Storage Node's Location LED. After replacing CRUs, turn off the Location LED by selecting the **Location LED OFF** radio button.

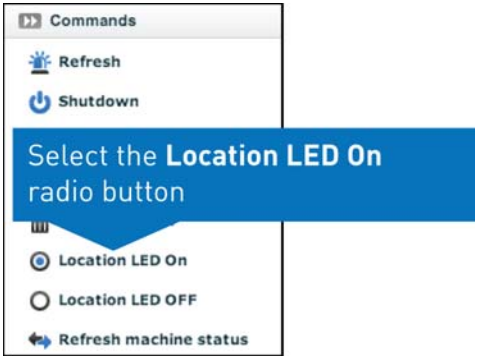

4 You can now locate the selected Lattus S10 Storage Node at the data center by its illuminated Location LED.

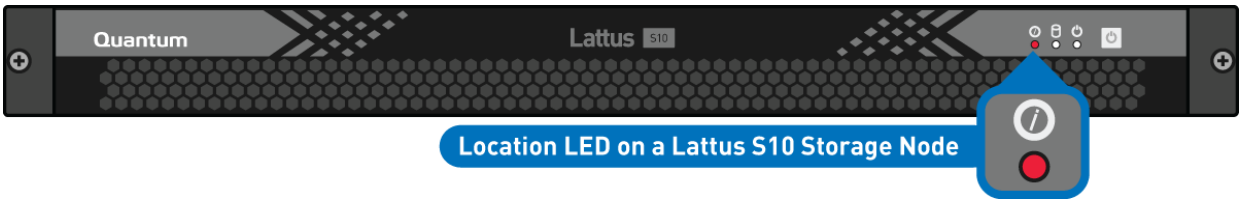

Replacing the Lattus S10 Storage Node Power Supplies

Prerequisites:

- A spare power supply

To replace a Lattus S10 Storage Node's power supply, follow these steps:

1. Turn on Lattus S10 Storage Node's Location LED by following the steps in Locating Lattus S10 Storage Nodes on page 58.
2. Go to the data center and locate the Lattus S10 Storage Node that contains failed power supply (look for the illuminated Location LED).
3. Identify the failed power supply. Its status LED will be flashing amber.

4. On the failed power supply, detach the safety on the power cable and unplug it.
5. Remove the power supply using the handle.
6. Insert the replacement power supply in the node and push it in until you hear a click.
7. Plug the power cable in and attach the safety. When installed correctly, the power supply's status LED will be green.

8 After successfully replacing the failed power supply, turn the Lattus S10 Storage Node's Location LED off from the Lattus CMC.

Replacing Decommissioned Disks in the Lattus S10 Storage Nodes

Replacements for your decommissioned disks will be shipped to you when required. To replace the disks, follow the steps below, which have been divided into seven sections:

Caution: Replace the decommissioned disks in one S10 Storage Node at a time. If you shut down too many S10 Storage Nodes at once, data unavailability may occur.

Be especially careful when replacing storage nodes disks in systems that have a small number of storage nodes, such as a 6-node system.

If the spread width is larger than the number of storage nodes, turning off one or more storage nodes can have a serious impact on Lattus availability. The moment the actual disk safety becomes negative, your Lattus installation is no longer accessible.

Follow these guidelines if you must replace disks in an environment with a small number of storage nodes:

1. Turn off only one node at a time.
2. Replace the disk(s).
3. Switch on the node.
4. When the node is fully operational, wait at least five minutes before turning off the next node.

Prerequisites:

- A screwdriver.
- A number of new storage disks equal to the number of disks that need to be replaced.
- Printout(s) of the details of the disk(s) that need to be replaced. This implies that the faulty drives have been decommissioned already (**Dashboard > Administration > Hardware > Disks, right-hand pane Decommission**).

Exporting the Decommissioned Disk Details

To properly identify decommissioned disks within a Lattus S10 Storage Node, you will need to print the details for the decommissioned disks from the Lattus CMC:

1 Navigate to **Dashboard > Administration > Hardware > Disks > Decommissioned**.

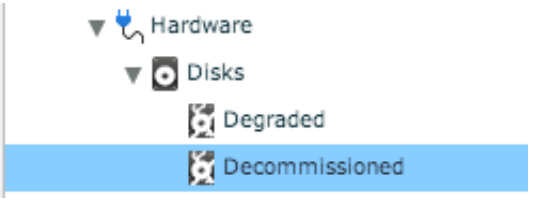

2 In the **Decommissioned Disks** screen, select the check boxes next to the decommissioned disks.

3 In the **Commands** pane, click one of the following:

a **Export details as PDF**: A PDF with all the necessary details will be generated.

b **Mail details**: specify an email address. A generated PDF with the necessary details will be sent to that email.

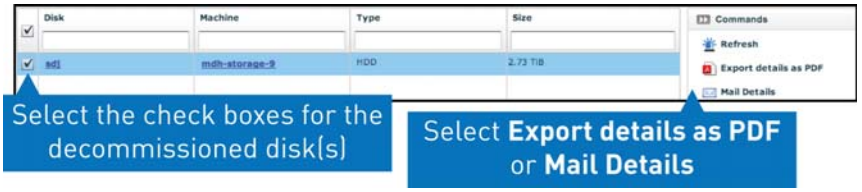

Note: Once you've determined which Lattus S10 Storage Node contains the decommissioned disks that need to be replaced, turn on its Location LED by following the steps in the "Locating Lattus S10 Storage Nodes" section on page 40.

Shutting Down the Lattus S10 Storage Node

Before replacing disks inside a Lattus S10 Storage Node, you must shut down the node that requires disk replacement. If you need to replace disks in multiple storage nodes, follow the steps below for shutting down and replacing the disks one node at a time.

Caution: Do not shut down more than one Lattus S10 Storage Node at a time. If you shut down too many S10 Storage Nodes at once, data unavailability may occur.

1. In the Lattus CMC, navigate to **Dashboard > Administration > Servers > Storage Nodes**.
2. In the **Commands** pane, click **Shutdown**.

For details, see Shutting Down Your Lattus System on page 13, and refer to the instructions for Shutting Down a Single Node on page 15.

Removing the Lattus S10 Storage Node from the Rack

To replace the disks in a Lattus S10 Storage Node, you must remove it from the rack completely. To remove it from the rack, follow these steps:

1. Once the Lattus S10 Storage Node has been successfully shut down, go to the data center and locate the Lattus S10 Storage Node that contains disks that need to be replaced (look for the illuminated Location LED).

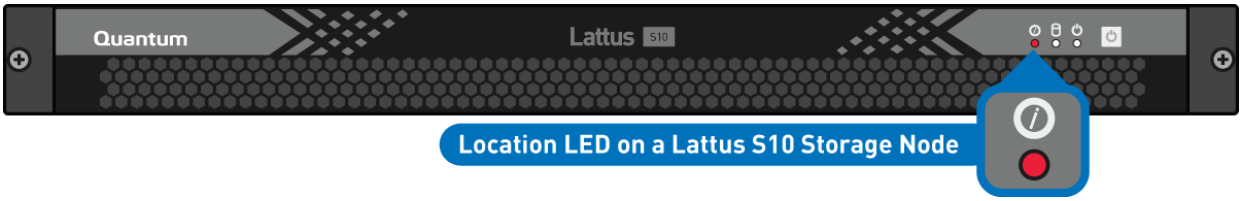

2 Detach the safety on the power cables and unplug them from the Lattus S10 Storage Node's power supplies.

Note: If they are not labeled already, mark the power cables to differentiate between the two to ensure that you reconnect the power cables to the correct power supplies when you reinstall the Lattus S10 Storage Node.

3 Disconnect the network cables from the Lattus S10 Storage Node's network ports.

Note: If labels or different colored network cables were not used to identify the network cables, mark them before unplugging them so you can differentiate between them when you reconnect them later. The example below shows how the network cables might be labeled to identify which one connects to Eth0, and which one connects to Eth1. This ensures that the ports on the Lattus S10 Storage Node are reconnected to the correct ports on the Lattus Rack Switches.

4 Unscrew the two front screws.

5 Slowly slide the Lattus S10 Storage Node out until you reach the pull-safety (you will hear a soft clicking sound).

6 Push the pull-safety on both sides of the Lattus S10 Storage Node, and slide the node out until you see the split line between the two top covers.

7 Using two people, remove the Lattus S10 Storage Node from the rack completely.

Caution: Once you pull the node past the pull-safety, don't leave the node in the rack. Otherwise the rack rails might break.

Replacing the Disks

Once you have removed the Lattus S10 Storage Node from the rack, follow these steps to replace the decommissioned disk(s):

1 Unscrew the four screws from the Lattus S10 Storage Node's top plate and carefully slide it off.

2 Using the decommissioned disk details report you exported from the Lattus CMC, identify the disks that need to be replaced.

Caution: Always check and confirm the serial numbers of the disks to verify that you're replacing the right ones. The location shown in the report's image is not guaranteed to be correct.

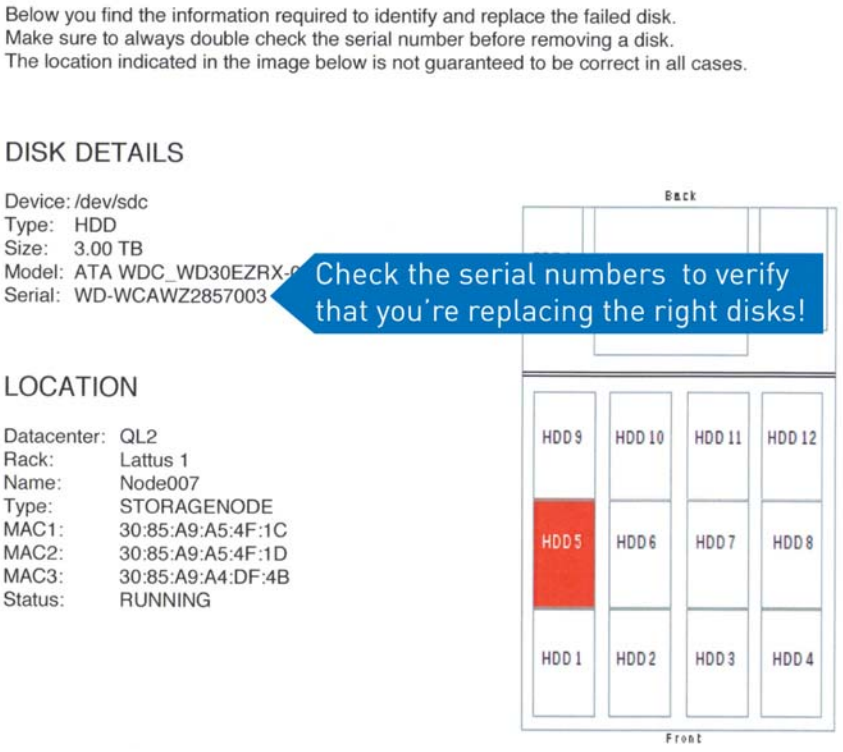

3 Unscrew the disk from the Lattus S10 Storage Node and slide it out using the handle.

4 Unscrew the disk from its protecting cover.

5 Screw the new disk in the cover and place it in the Lattus S10 Storage Node. Screw the disk into the Lattus S10 Storage Node.

6 Repeat steps 1 through 5 for each disk that must be replaced.

7 When all of the decommissioned disks have been replaced, slide the top plate back on the Lattus S10 Storage Node and screw it back on to the node tightly.

Reinstalling a Lattus S10 Storage Node in a Rack

1 Using two people, safely reinstall the Lattus S10 Storage Node back into its position in the rack and slide the node all the way into the rack.

Caution: Do not leave the Lattus S10 Storage Node partially installed in the rack. Slide it all the way into the rack. Otherwise the rack rails might break.

2 Tighten the two front screws.

3 Reconnect the network cables to the ports from which you disconnected them.

4 Reconnect the power cables to the power supplies from which you removed them, and attach the safety.

Booting the Lattus S10 Storage Node

1 Push the power button on the front of the Lattus S10 Storage Node to boot it.

2 Check if the node has successfully booted and if the disks have been initialized by doing the following:

a Look in the Events section of the **Dashboard**. When the machine has booted, you will see the two events listed below:

- ***Machine was rebooted***
- ***New empty disk(s) detected***. This event may take up to 15 minutes to appear in the Events list. If you insert a disk that does have data on it, then you will see the "New non-empty disk(s) detected" event.

b After the second event appears, go to **Dashboard > Administration > Lattus Management > Logging > Jobs**.

c In the **List of Jobs** screen, look for the job ***Initializing new disks on machine <machine name>***.

Note: If the job doesn't appear, but the ***New empty disk(s) detected*** event appeared, it could mean that you replaced the wrong disk(s). Verify that the serial number of disk you replaced matches the serial number of the decommissioned disk in the disk details PDF you exported from the Lattus CMC earlier.

d Select the job to monitor its progress. Once the job completes, its status will be displayed as Done (blue check mark; initialization was successful) or Error (red check mark; initialization failed) along with the job's start and end timestamps.

3 If the disks fail to initialize, contact Quantum support. Click the failed job to display its details in the Job Details screen; technical support will need this information to help determine the cause and resolution for the failed disk initialization.

What to Do When You're Finished Replacing Disks in the Lattus S10 Storage Node

Make sure you turn off the Lattus S10 Storage Node's Location LED.

If you have additional Lattus S10 Storage Nodes with decommissioned disks that need to be replaced, repeat the process for each of those nodes.

Caution: Replace the decommissioned disks in one Lattus S10 Storage Node at a time. If you shut down too many nodes at once, data unavailability may occur.

Chapter 8

Growing Lattus Capacity and Performance

As your capacity and performance demands grow, you can grow your Lattus can grow to meet them. This chapter discusses determining when capacity and performance growth may be necessary, and some common tasks related to growing your Lattus system.

This chapter covers the following topics:

Determining When You Need More Storage Capacity

Lattus will begin warning you, through events in the Lattus CMC Dashboard and e-mail notifications, when the used storage capacity reaches 70%.

Additionally, as the storage pool fills up, you may see events and receive notifications related to the number of Check-Blocks on the individual drives within the Lattus S10 Storage Nodes.

If you see events or receive notifications that your Lattus system's capacity 80% full or more (Storage pool is more than 80% full), Quantum recommends that you contact your Quantum sales representative to purchase more Lattus S10 Storage Nodes.

Note: If you add more Lattus S10 Storage Nodes, Quantum Support will need to rebalance your data to include the new storage nodes as part of the spread. Until rebalancing is performed, your durability policies will not use the new storage nodes to protect your data.

The image below shows the **Storage Pool Used Capacity** status bar and graph in the Lattus CMC for a Lattus system that is more than 90% full. At this point, the events raised will be of the severity level CRITICAL. This status bar and graph are helpful tools for monitoring how much of your Lattus capacity is being used.

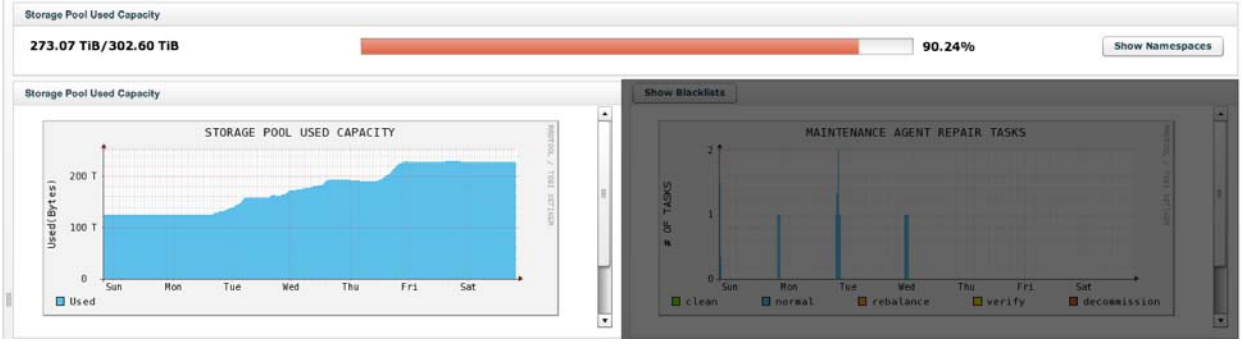

Determining the Size and Quantity of Your MetaStores

The number of objects that can be stored in a namespace is limited by the size of the MetaStore. The size of the superblocks used for storing your objects has an impact on the number of objects that can be stored in the MetaStore. A single MetaStore can only perform a certain number

of input/output operations per second (IOPS), so the maximum number of IOPS that your Lattus system can perform is dependent on the number of MetaStores available.

Details about these factors are listed below:

- A superblock entry on a MetaStore is 200-400 bytes, and a Metastore SSD cluster is 200 GB. This means you can store approximately 400,000,000 superblocks on a MetaStore.
- A MetaStore can do approximately 1,000 get operations per second. This means that for three Lattus C10 Controller Nodes and a single MetaStore, you need approximately 3 MB objects to achieve 3 GB/s.
- A Metastore can do 100 put operations per second. This means that for three Lattus C10 Controller Nodes and a single MetaStore, you need approximately 30 MB objects to achieve 3 GB/s
- Tuning your superblock is critical but has 2 effects.
    - The smaller your superblock, the more metadata entries this takes per object and the less objects you can store per MetaStore.
    - 64 MB superblock can only achieve a 4 TB maximum object, so a 256 MB superblock is needed for 16 TB objects.

If you require additional MetaStores for increased capacity and performance, you should contact your Quantum sales representative to discuss adding more MetaStores.

Creating New Durability Policies (a.k.a. Storage Policies)

A durability policy (also known as a “storage policy” in the Lattus CMC) is a user-configured policy that determines the level of protection for the objects being stored. A single durability policy may be used for several namespaces, but a single namespace can only be tied to one durability policy at a time.

As part of growing your Lattus system, you may need to create additional durability policies with settings to accommodate different

kinds of data. Work with Quantum Support and follow the steps below if you need to create new durability policies.

1 In the Lattus CMC, go to **Dashboard > Administration > Storage Management > Storage Policies**.

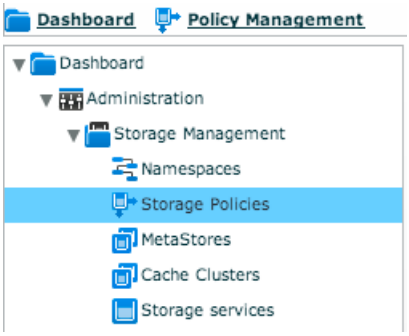

2 In the **Commands** pane, click **Add storage policy**. The **Create New storage policy** wizard appears

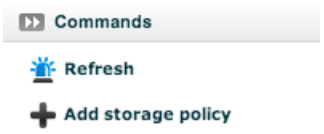

On the General Settings Tab:

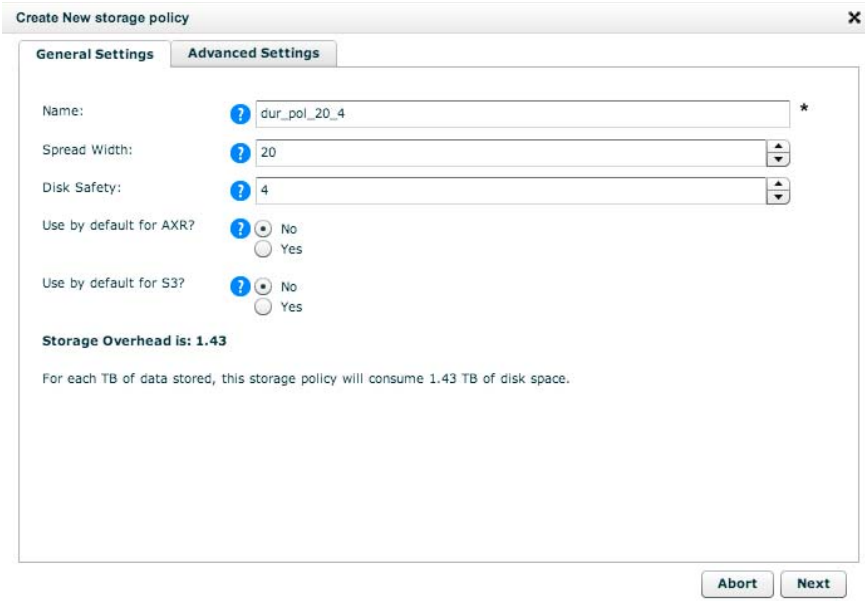

3 Specify a **Name**. It's recommended that you use a name that describes the policy's spread width and disk safety. For example, if the policy's spread width is 20, and the disk safety is 4, you might name the policy **dur_pol_20_4**.

4 Make sure you understand "Spread Width" and "Disk Safety" (see Lattus Terms and Concepts on page 93 for a description of these terms).

In documentation for Lattus, durability policies are written as a ratio of spread width to disk safety (e.g., 20/4).

The table below shows an example of the most common durability policies used with a Lattus 6-node and 20-node systems; and one, two, and three site configurations. Please work with your Pre-Sales representative for exact numbers.

# of Sites	# of Lattus S10 Storage Nodes	Durability Policy	Raw capacity	Usable capacity
1	6	12/4	216	123
2	6	12/7	216	102
3	6	12/5	216	110
1	20	20/4	720	490
2	20	20/11	720	276
3	20	18/7	720	374

- For a general calculation take the RAW/((Spread/(Spread-DiskSafety))*1.143)

5 Select the **Spread Width** and the **Disk Safety**.

- **Multi-Rack and Multi-Geo Considerations for Spread Width**

  The spread width should be evenly divisible by the number of racks or data centers across which the data will be spread. For example, if you wanted to spread data across three sites, you might use a spread width of 18. This would spread your data over six Lattus S10 Storage Nodes/disks at each site (assuming that your Data Balance settings are configured to force data dispersion over systems, racks, and data centers).

- **Multi-Rack and Multi-Geo Considerations for Disk Safety**

  If you're spreading your data across multiple sites, it's recommended that you use a disk safety that protects your data against the loss of one entire site plus one or two additional storage nodes/disks. For example, if you're spreading your data across three sites using a spread width of 18, you might choose a disk safety of 7 or 8.

- **Selecting the Appropriate Durability Policy for Lattus Base Systems**

  a **Recommended Durability Policy for 20-node Lattus Base System**: For a 20-node base Lattus system in a single site, the recommended durability policy is 20/4 (spread width/ disk safety).

b **Recommended Durability Policy for 6-node Lattus Base System**: For a 6-node base Lattus system in a single site, the recommended durability policy is 12/4 (spread width/disk safety).

Note: Due to the limited number of Lattus S10 Storage Nodes on a Lattus 6-node base system, a different spread width and disk safety are recommended. Since there are only 6 nodes, Quantum recommends a Spread Width of 12 and a Disk safety of 4. This will put two pieces of each object on every node (**6** Lattus S10 Storage Nodes x **2** disks per node = **12** disks) and still allow for four disk failures, but only two node failures (**2** disks per node x **2** node failures = **4** disk failures).

---

6 Choose whether you want to use this durability policy as the default for the namespaces you create through the Lattus CMC (which are accessed using the Lattus REST API a.k.a. "AXR") and/or as the default for "buckets" created and accessed using the Amazon Simple Storage Service (S3) API.

a Selecting Yes for Use by default for AXR? will cause this durability policy to be preselected as the default durability policy when you create new namespaces.

Note: You can still select other configured durability policies for your namespaces when you create them.

---

b Selecting Yes for Use by default for S3? will cause this durability policy to be used as the default when creating S3 buckets.

Note: You can still select other configured durability policies for your S3 buckets when you create them. See the Lattus REST API User's Guide for details (http://www.quantum.com/lattusdocs).

If you select **Yes** for **Use by default for S3?**, the **Small File Support?** option will appear. Choose whether or not you want to use the policy for small files (see Lattus Terms and Concepts on page 93 for an explanation of small file support). If you choose **Yes**, provide the file size threshold (default is 512 KiB).

On the Advanced Settings Tab:

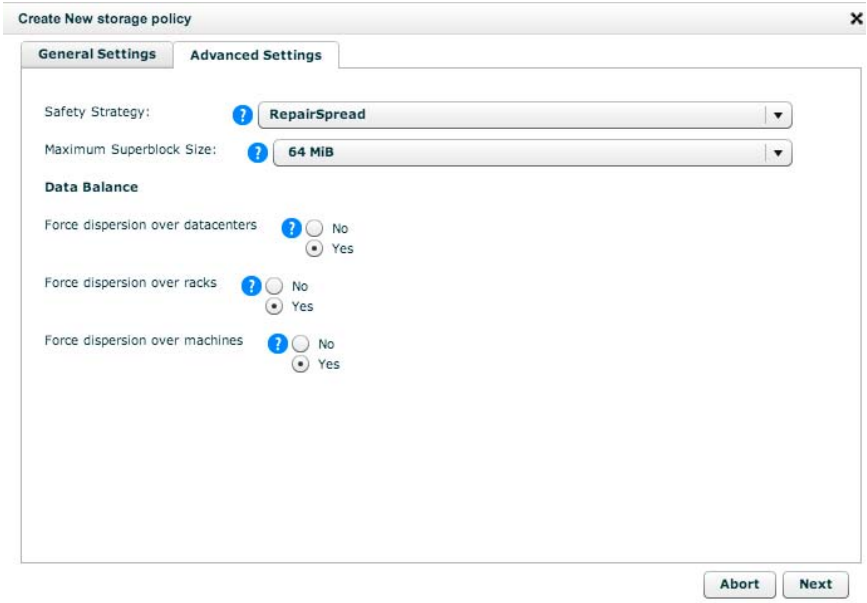

7 Select a Safety Strategy for your durability policy. Quantum always recommends using RepairSpread & DynamicSafety.

Caution: This strategy should be avoided in multi-geo setups with imbalanced latency between sites if the imbalance causes blockstores at a remote site to timeout. In that scenario, performance and data safety will be suboptimal. Quantum recommends avoiding imbalanced latency configurations that might lead to this scenario. However, if it cannot be avoided, then the configuration should be reviewed with Quantum to determine the optimal policy based on your performance and data protection requirements.

(For a list of all the available safety strategies and their descriptions, refer to Lattus Terms and Concepts on page 93).

8 Select the **Maximum Superblock Size**: Lattus supports a minimum superblock size of 2 MiB, and a maximum superblock size of 256 MiB.

Note the following maximum object sizes for the corresponding superblocks:

Superblock Size	Maximum object size
2 MiB	128 GiB
4 MiB	256 GiB
8 MiB	512 GiB
16 MiB	1 TiB
32 MiB	2 TiB
64 MiB*	**4 TiB**
128 Mib	8 TiB
256 MiB	16 TiB

***Quantum recommends using the default superblock size, which is 64 MiB.**

9 **Force dispersion over datacenters**: Select **Yes** if using a multi-geo configuration.

If you select Yes, Lattus will ensure that when it generates a list of disks to store your data upon, your disks are equally spread over these data centers. Choose this option when you want to use Lattus in a multi-geo setup. Ensure that your spread width is a multitude of the number of data centers you are using. For example, if you are spreading data across three data centers, you might select a spread width of 18.

10 **Force dispersion over racks**: Only select **Yes** if you want to have your data spread evenly across your racks in the case of a rack failure. For example, if you have three racks and want to sustain a rack loss, use a suggested durability policy of 18/7. Make sure your spread width is divisible by the number of racks you have.

If you select **Yes**, Lattus will ensure that when Lattus it generates a list of disks to store your data upon, your disks are equally spread over these racks. Choose this option when you want to use Lattus in a multi-rack setup. Ensure that your spread width is a multitude of the number of racks you are using. For example, selecting spreading

over racks doesn't make sense if you are using two racks and one rack is only half filled. Forcing spreading will lead to your environment only being able to get filled to 2/3 of its capacity

11 **Force dispersion over machines**: Always select **Yes**.

12 Click **Next**.

13 In the dialog box, click **OK**.

Creating New MetaStores

A MetaStore is a set of three SSDs spread across three Lattus C10 Controller Nodes. A MetaStore is designed to store the metadata of objects, superblocks, spreads, policies, and namespaces. One or more namespaces can be created on a MetaStore, but a namespace cannot span multiple MetaStores. Additional details about MetaStores can be found at the beginning of this section, as well as the "Lattus Terms and Concepts" section on page 63.

Depending on the number of concurrent object operations or capacity needs, you may need to create multiple MetaStores.

Contact your Quantum sales representative if you need to add MetaStores.

Note: If you plan to create multiple MetaStores, create them one at a time. Do not create them in parallel. Add the next MetaStore after the previous one is successfully installed.

To create a new MetaStore, follow these steps in the Lattus CMC:

1 Go to **Dashboard > Administration > Storage Management > MetaStores**.

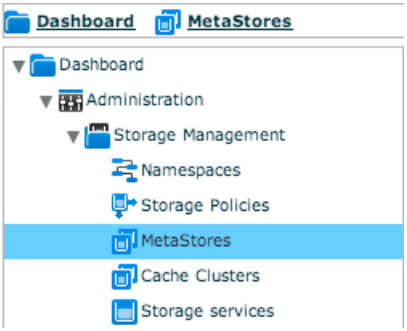

2 In the **Commands** pane, click **Add MetaStore**.

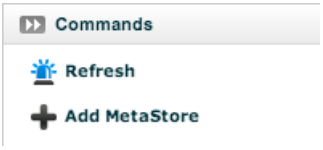

3 In the **Add MetaStore** window, on the **Add MetaStore** tab, specify a name and size for the MetaStore.

4 On the **Disk Configuration** tab, specify the SSDs for the MetaStore. Specify at least three different SSDs on three different Lattus C10 Controller Nodes. For consistency, it's recommended that if your Lattus C10 Controller Nodes contain multiple SSDs, you should select the SSD with the same name (e.g., **sdc**) on each Lattus C10 Controller Node for all members of a single MetaStore.

For example, if you chose to use the disk named **sdc** for a MetaStore named **objectMetaStore** on the first Lattus C10 Controller Node, you should also use the disk named **sdc** on the other two controller nodes. The image below illustrates this example, showing the correct and incorrect way to configure a MetaStore.

Correct

In this example, the MetaStore named **objectMetaStore** is made up of the three SSDs. It's configured correctly to use the SSD named **sdc** on each of the Lattus C10 Controller Nodes.

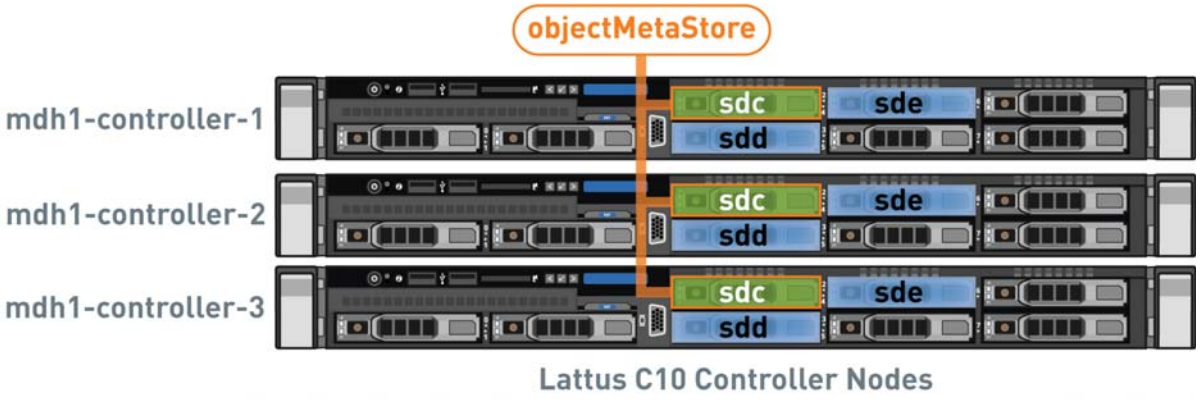

Incorrect

In this example, the MetaStore named **objectMetaStore** is made up of the three SSDs. However, it is configured incorrectly to use an SSD with a different name on each Lattus C10 Controller Node.

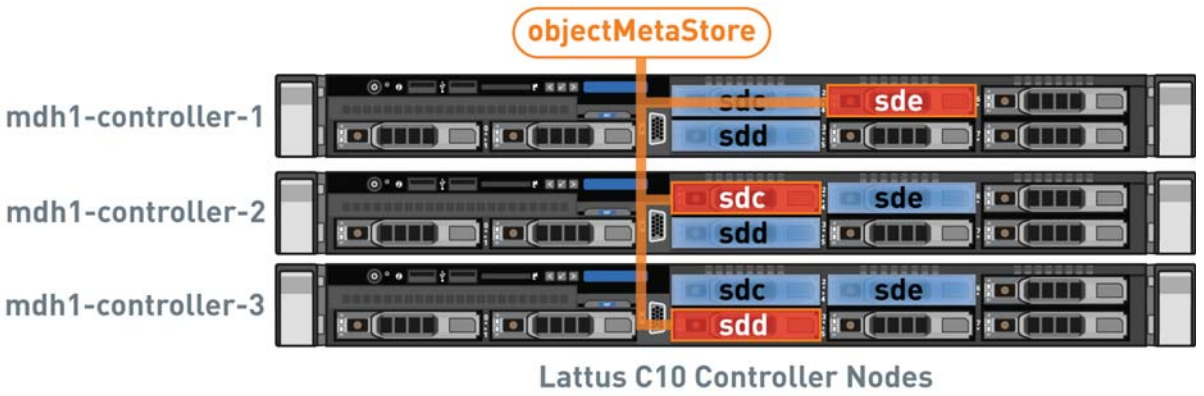

Following is an example of the **MetaStore:[MetaStore Name]** screen for the MetaStore **objectMetaStore** after it's been configured. Here you can see which SSD the MetaStore is using on each of Lattus C10 Controller Nodes.

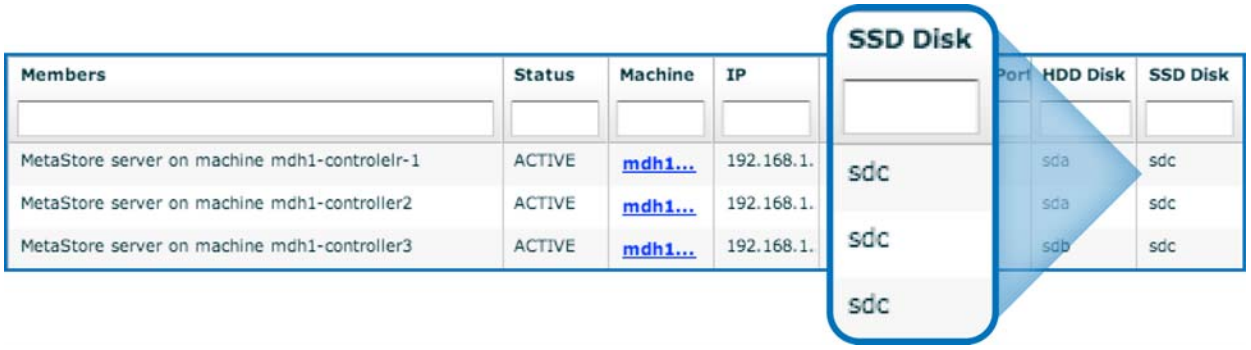

5 After specifying the SSDs for the MetaStore, click **Next**.

6 In the dialog box, click **OK**.

Creating New Namespaces

A namespace is a logical segregation of data. A highly simplified analogy for a namespace would be a folder containing files belonging together. The maximum size of a namespace is limited by the size of the MetaStore. Each namespace will have one associated MetaStore (responsible for storing object and system metadata) and durability policy (a user-configured policy that determines the level of protection for the objects being stored).

To configure a namespace, follow these steps:

1 In the Lattus CMC, go to **Dashboard > Administration > Storage Management > Namespaces**.

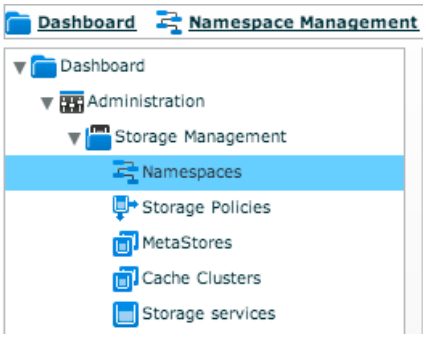

2 In the **Commands** pane, click **Add namespace** to open the **Add namespace** wizard.

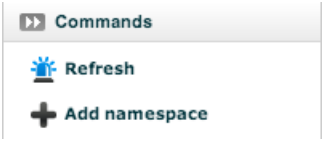

3 From the **Number of namespaces** drop-down menu, select the number of namespaces you want to add and click **Next**.

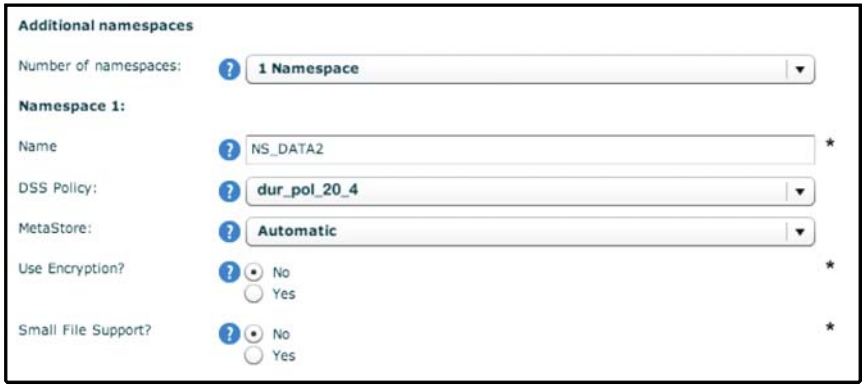

4 For each namespace, set the following:

a **Name**: Enter a name for the namespace. For example, **NS_DATA2**.

b **DSS Policy**: Select the desired durability (storage) policy. For example, **dur_pol_20_4**.

c **MetaStore**: Depending on the MetaStores configured in the system, select the MetaStore you would like to be used for the user data. Selecting **Automatic** from the **Metastore** drop-down would automatically select the MetaStore with the most amount of free space.

d **Use Encryption?**: Choose whether or not you want to have the data encrypted (see Lattus Terms and Concepts on page 93 for details on namespace encryption).

e **Small File Support?**: Choose whether or not you want to use the namespace for small files (see the "Lattus Terms and Concepts" section on page 63 for an explanation of small file support). If you choose **Yes**, provide the file size threshold (default is 512 KiB).

5 Click **Next**.

6 In the dialog box, click **OK**.

Namespace Limitations

Namespaces typically have two limitations:

- Number of objects they can hold
- Number of concurrent operations that can be done

Maximum Number of Objects per Namespace

Every object in a namespace consumes metadata. This metadata is stored in a MetaStore, which lives on a file system on top of an SSD. As such, the maximum number of objects that can be stored is limited by the size of this SSD (SSDs in the Lattus C10 Controller Nodes are 200 GB).

Lattus doesn't support spanning a namespace over multiple MetaStores. In order to identify how many objects can be stored in a namespace, the number of superblocks per object and the size of the object metadata must be determined (your Quantum Professional Services representative can assist you with this). Typically every superblock of an object consumes 200-400 bytes.

When designing for the number of objects, take this limitation in consideration. There are two approaches:

1 When the namespace is exhausted, the application starts using a new namespace tied to a different MetaStore (using different SSDs).

2 Your application is capable of balancing writes into multiple namespaces based upon their fill-level. Initially, when a single MetaStore is sufficient, your application writes to just one of the namespaces. But when your application exceeds a specified limit, it should begin writing to a second namespace. However, it can still use the first namespace depending upon consumption.

Keep in mind that moving an object from one namespace to another not only requires rewriting the metadata but also the data underneath (which drives data retrieval, decoding and encoding the object and finally storing the encoded data on the back-end).

Concurrent Operations

The MetaStore managing object metadata for that namespace is limited in the number of object operations it can handle. These operations are:

- put
- get
- delete
- list
- repair
- monitoring

Note: We recommend to never exceed 1,000 of these operations (put, get, delete, list combined together) per second per MetaStore.

How Superblocks are Stored

Let's assume that you've configured at least one durability policy and one MetaStore, and created a namespace, and you're now writing data to the Lattus object storage.

This section describes at a high level how the superblocks are stored. Storing a superblock basically consists of the following steps:

1 Generating a spread

2 Storing (extra) Check-Blocks and/or full copies to the disks of the spread

3 Validating the result and returning to step 1 if necessary

Generating a Spread

Generating a spread takes into account the following:

- Reusing disks that successfully received data in a previous loop.
- Filtering out all disks that are not ONLINE.
- Filtering out all disks that have failed in a previous loop.

  If the safety strategy for this round is DynamicSafety, it is OK to continue if the resulting spread width is lower than the policy's spread width minus the disk safety (“spread width - disk safety”).

  For all other strategies, the next strategy in the strategy list should be tried if the resulting spread width is lower than the policy's spread width.

Storing the (extra) Check-Blocks

The spread generated in step 1 contains two types of disks:

- **Disks that have data uploaded in previous rounds with success**: For this category nothing needs to be done.
- **Disks that still need their upload to be executed**: For this category the upload needs to be executed

  This upload is executed to all disks in parallel. This upload ends in one of the following situations:

  - All individual uploads have succeeded.
  - Some uploads have succeeded, triggering a grace timeout and canceling the remaining uploads.
  - Some uploads have succeeded, but not enough to trigger the grace timeout, but the global timeout per superblock has canceled the upload.

    The grace timeout gives slower daemons some extra time to finish after the policy's “spread width - disk safety” disks have finished successfully. The default for this timeout is 10 seconds. The next step takes this result and validates it.

Validating the Spread

This step checks the number of disks that have succeeded. If all disks have succeeded, the storing of the superblocks has succeeded. If less have succeeded, the end result depends on the current safety strategy:

Safety Strategy	End Result
RepairSpread	The failed disks are marked as failed and a new round is tried.
DynamicSafety	The superblocks will be stored if at least the policy's "spread width – disk safety" has succeeded, otherwise the storing of the superblock fails.
Stop on Failure/No strategy	The storing of the superblock fails.

Editing Existing Namespaces

Editing a namespace allows you to:

- Select a different durability policy for the namespace.
- Enable or disable encryption for the namespace.
- Enable or disable small file support for the namespace.

Note: You cannot change a namespace's name or MetaStore.

To edit an existing namespace, follow these steps:

1. In the Lattus CMC, **Dashboard > Administration > Storage Management > Namespaces.**

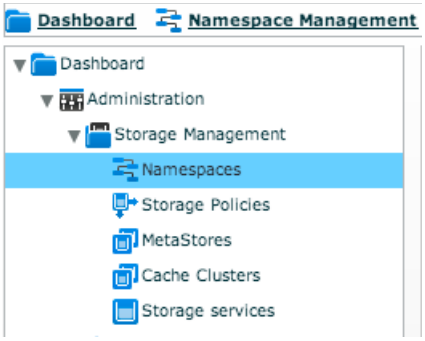

2 Select the namespace you want to modify to enter the **Namespace: system** screen.

3 In the **Commands** pane, click **Edit**.

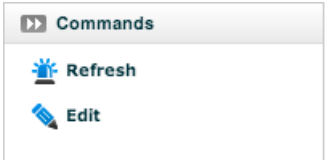

4 Use the options **Edit Namespace** wizard to make your changes, then click **Next** to apply them.

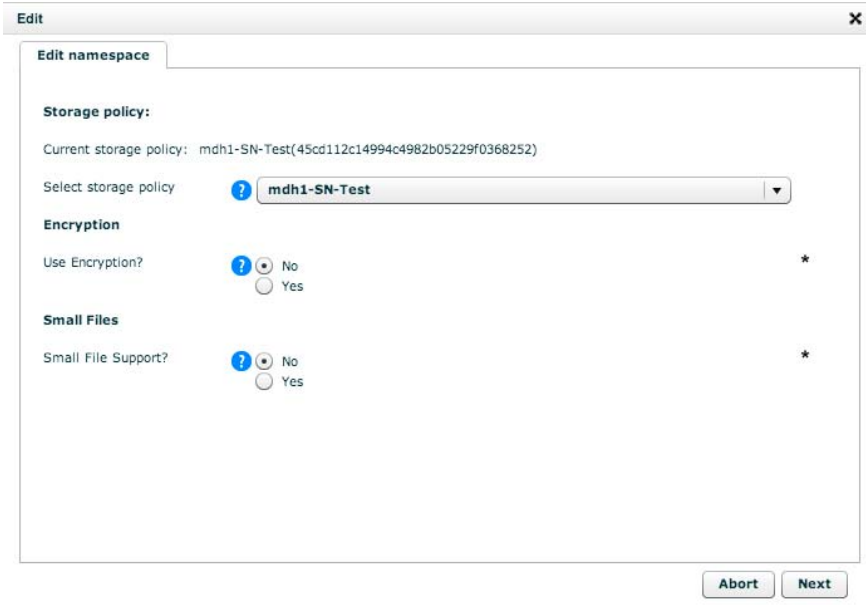
Edit
Edit namespace
Storage policy:
Current storage policy: mdh1-SN-Test(45cd112c14994c4982b05229f0368252)
Select storage policy
mdh1-SN-Test
Encryption
Use Encryption?
No
Yes
*
Small Files
Small File Support?
No
Yes
*
Abort
Next

Chapter 9

Lattus Reference

This chapter covers the following topics:

Lattus Hardware

You can find information about the Lattus A10 Access Node in the Lattus A10 User Essentials. You can find information about the StorNext M660 Metadata Appliance in the StorNext M660 User Essentials.

Lattus C10 Controller Nodes

The Lattus C10 Controller Nodes provide management and control of the Lattus object storage system. A base system contains three controller nodes, and more controller nodes can be added. Controller Node 1 is typically configured as the management node and provides the Lattus CMC GUI interface to the Lattus storage software. The Lattus C10 Controller Nodes provide access to the Lattus object storage from

Lattus object storage-enabled StorNext MDCs, Lattus A10 Access Nodes; as well as applications using HTTP/HTTPS REST, S3, and WebDAV.

Front View

Rear View

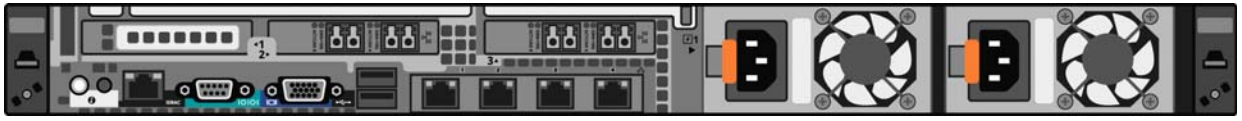

Lattus S10 Storage Nodes

The Lattus S10 Storage Nodes are the main building blocks of the Lattus system. They provide high-density power-efficient “green” storage for the Lattus Object Storage system. Each Lattus S10 Storage Node contains 12 3TB HDDs and software that is specifically designed for Lattus object storage systems. Multiple Lattus S10 Storage Nodes can be spread geographically or located centrally, depending on location, durability and access requirements.

Front View

Rear View

Lattus Rack Switches

The Lattus Rack Switch is the back end network connection between the Lattus C10 Controller Nodes and the Lattus S10 Storage Nodes. The Rack Switch is a 1U, redundant power, rack-mount switch. The switch comes in two configurations:

- **Non-stacking**: Comes with 43 1GbE ports for Lattus S10 Storage Node connectivity; and two dual-port 10-GbE SFP+ modules for Lattus C10 Controller Node connectivity.
- **Stacking**: Comes with 43 1GbE ports for Lattus S10 Storage Node connectivity; one dual-port 10GbE SFP+ modules for Lattus C10 Controller Node connectivity; and one stacking module with two 12GbE stacking ports.

Rules

- A one-rack configuration will typically come with two non-stacking rack switches.
- Two to three rack configurations recommend the use of stacking rack switches, unless either the distance between racks precludes use of stacking rack switches, or the number of controller nodes per rack precludes use of stacking rack switches. Consult your Quantum expert if any questions.
- Greater than three-rack configurations require a system switch and all non-stacking switches.

Notes

- Switches are convertible from non-stacking to stacking and vice versa.
- The stacking cables have a maximum length of four meters, so plan the locations of your racks accordingly. If your racks are further apart, non-stacking rack switches can be used.
- Please consult Quantum Professional Services with questions.

Front View

Rear View

Lattus System Switch

The Lattus System Switch is used as the aggregation switch between racks in a four or greater rack configuration, or configurations that have six or more Lattus C10 Controller Nodes across three racks. The Lattus System Switch connects the Lattus Rack Switches and Lattus C10 Controller Nodes between racks.

Front View

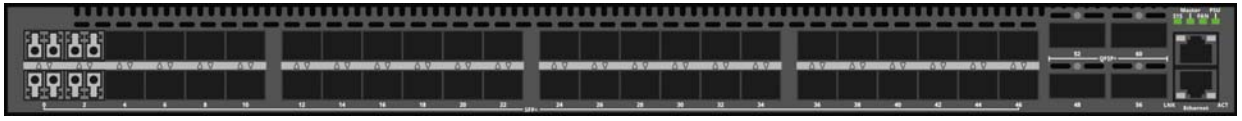

Rear View

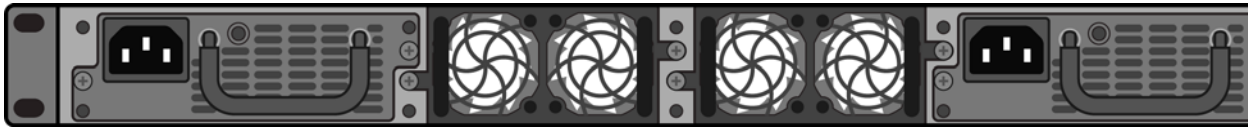

Lattus Interconnect Switch

The Lattus Interconnect Switch connects the Lattus C10 Controller Nodes to external devices or applications. These devices may include StorNext M662 Metadata Appliances (or customer-supplied MDCs) and Lattus A10 Access Nodes, as well as other applications using HTTP/HTTPS REST, S3, or WebDAV.

Front View

Rear View

Lattus Terms and Concepts

Blacklists

Blacklists can relate to a disks, storage nodes, or storage daemons:

When the monitoring agent on each Lattus S10 Storage Node collects disk health information:

1. The monitoring policy on the management node will pull the data collected on each storage node.
2. The monitoring policy on the management node will then populate the blacklist counters and populate the Lattus CMC Dashboard's Blacklists graph for any disks that may have seen errors but not reached their threshold.
3. Disks that continue to have blacklist events will continue to be reported in the Blacklists graph until the disks are decommissioned.

4 The blacklist will also be associated with a specific blockstore, which relates to the disk and storage daemon.

Each blockstore periodically performs a basic read and write operation on its backing store.

Once a blockstore is set on the blacklist, a connection will be attempted every 30 seconds.

- If a connection can be made, Lattus will query the blockstore to determine whether or not it can perform read and write operations.
- If either the read or write operation fails, or both; then the blockstore remains in the blacklist.
- If the blockstore can perform both read and write operations, it becomes un-blacklisted.

In the event a node is put on a blacklist, if the node starts to respond, the un-blacklist process will reset its counters.

Notes

- Blacklist events can be generated from any number of cases. Examples include failed REST events (GET, PUTS), disk smart tests, or any number of latency issues.
- Blacklists are not an indication that a disk, storage node, or storage daemon is in a faulted state, but they should be monitored.
- If the blacklist becomes more severe you should start to investigate the cause.

Blockstore

Blockstore is another term for a storage disk and its associated storage daemon in a Lattus S10 Storage Node.

Check-Blocks

Lattus QSpread technology is an advanced erasure-coding technology. QSpread splits and encodes data objects into multiple Check-Blocks that directly encode redundancy.

These Check-Blocks are placed across Lattus S10 Storage Nodes, using a maximum of one drive per storage node for each spread selected (two drives per storage node in 6-node base systems using a 12/4 durability policy). This will ensure that the unavailability of a storage node does

not affect more than one drive in each spread (or more than two drives in 6-node base systems using a 12/4 durability policy).

QSpread requires only a subset of the Check-Blocks to restore the original data object, as determined by the user-specified durability policy. In the case of the 20/4 policy, QSpread can decode and assemble any object by reading from 16 drives out of the 20 written in the spread.

Client Daemon

The client daemon is a service/process/daemon running on each Lattus C10 Controller Node. The client daemon handles REST read, write, update, and delete requests for data.

The client daemon encodes and decodes data as requested from applications external to the storage system. The Lattus A10 Access Node and StorNext MDCs also communicate with the client daemons, as their requests to the Lattus object storage are also handled through REST.

A single controller can have one or more client daemons running. Each client daemon has the same view of the underlying storage. Each client daemon listens for REST requests on all public networks, but on a unique port to that controller (8080, 8081, etc.). Each client daemon communicates internally on a unique port (23510, 23511, etc.).

Decommissioning Disks

If a disk is in a DEGRADED state and cannot be recovered, it should be decommissioned. When a degraded disk is decommissioned, the following occurs on the Lattus system:

1. The repair manager will run for all objects in each namespace.
2. The repair will attempt to move the data from the decommissioned disk to another capable disk.
3. If the data cannot be moved from the decommissioned disk (e.g., the disk is not readable or the disk has been replaced), a normal repair mechanism will be initiated.
4. The blockstore associated with the disk will be set to an ABANDONED state.
5. The ABANDONED blockstore will no longer be used for storing objects.

Note: On the Lattus CMC Dashboard, the total number of disks in the environment is not decreased, even if a disk is degraded or decommissioned. The Dashboard will not display the number of decommissioned disks in the Lattus system.

Degraded Repair

The Lattus monitoring agent will monitor disk health. In the event a disk exceeds the acceptable blacklist threshold, the disk will be set to a degraded state. Lattus will continue to attempt to use degraded disks in the spread.

Note: On the Lattus CMC Dashboard, the total number of disks in the environment is not decreased, even if a disk is degraded or decommissioned. The Dashboard will display the number of degraded disks in the Lattus system.

Disk Replacement

By decommissioning degraded disks in the Lattus S10 Storage Nodes, you ensure the highest possible durability of the data stored in your Lattus object storage. Quantum will regularly provide replacements for your decommissioned disks so you can replace them.

Disk Safety

Disk Safety specifies the number of concurrent drive failures (on different Lattus S10 Storage Node) that can be handled without any risk of potential data loss. For example, if you store an object in a namespace using a durability policy of 20/4 (spread width/disk safety), you could lose four storage nodes and still be able to get the object back from the remaining 16 nodes. The minimum configurable disk safety is 1 (Quantum does not recommend using a disk safety less than 4), and the maximum is equal to the spread width - 1 (e.g., 20/19). Work with your Quantum Professional Services representative to determine the appropriate settings for your Lattus system.

Durability Policy

A durability policy is a user-configured policy that determines the level of protection for the objects being stored. A single durability policy may

be used for several namespaces, but a single namespace can only use one durability policy at a time.

Using the Lattus CMC, you will need to define the settings for any durability policy you create. The key settings that you will need to understand when configuring any durability policy are **spread width**, **disk safety**, and the **safety strategy**.

In Lattus documentation and training, durability policies will typically be referred to as a ratio of the spread width to the disk safety. For example, a durability policy of 20/4 has a spread width of 20 and a disk safety of 4. Work with your Quantum Professional Services representative to determine the appropriate settings for your Lattus system.

Events

Everything that happens in the Lattus system is considered an "event". Some events are informational, like a service coming up, while others may indicate problems, such as objects having low disk safeties.

Lattus assigns a severity to each type of event. The severity levels are (ranging from least to most severe):

- Information
- Warning
- Error
- Urgent
- Critical
- Unknown

Maintenance Agent

The maintenance agent is a service/process/daemon running on each Lattus S10 Storage Node. It resides on Lattus S10 Storage Nodes by default, but it can be configured for any node in the environment. The maintenance agent is responsible for the self-repairing nature of the Lattus storage back end. The maintenance agent polls storage daemons for objects to be repaired and objects to be deleted. The maintenance agent then instructs the storage daemon responsible for the namespace to conduct the actual deletion process.

Note: The polling for repair work is done every 15 minutes.

The maintenance agent works closely with the storage daemon. It does not need the client daemon to get access to the Lattus back end.

Management Node

The management node is the first Lattus C10 Controller Node that's configured and initialized. The management node provides management of the storage and schedules jobs. The Lattus Cloud Management Center (CMC) will run on the management node.

MetaStore

A MetaStore is a set of three SSDs spread across three Lattus C10 Controller Nodes. A MetaStore is designed to store the metadata of objects, superblocks, spreads, policies, and namespaces. One or more namespaces can be created on a MetaStore, but a namespace cannot span multiple MetaStores.

Listed below are the key details about MetaStores:

- MetaStores are placed on SSDs (solid-state drives) for high input/output operations per second (IOPS).
- Each MetaStore consists of three SSDs on three distinct Lattus C10 Controller Nodes for high availability. The store is considered a “cluster”, each participating controller is considered a “node”, and each service is considered an “instance”.
- To write and retrieve data from a MetaStore, you need a majority of the participating instances available. This means that, since a MetaStore that spans three Lattus C10 Controller Nodes, two instances need to be running.

  Caution: If you are down to one SSD, contact your Quantum sales representative.

- Each instance also has an associated transaction log, called a tlog, which resides on an HDD. You only need one intact copy of the database or tlogs to rebuild the entire MetaStore array, guaranteeing data safety when multiple components fail.
- A MetaStore can have multiple Namespaces, but a Namespace cannot span multiple MetaStores
- A MetaStore can only store a limited number of objects in its Namespaces and is limited by the size of the SSD.

Monitoring Interval

When the monitoring agent detects events, it will continue to check at specified monitoring intervals to determine whether the event is still applicable. If an event is still applicable since last time it was checked, it will continue to be listed as a "live event." If the event is no longer applicable, it will remain a live event for two monitoring intervals and then drop off of the Live Events list (assuming it is still no longer applicable).

Some events do not have a specified monitoring interval. When these types of events are detected, they will remain live for 2,000 seconds after they are no longer applicable.

Namespace

A namespace is a logical segregation of data.

A highly simplified analogy for a namespace would be a folder containing files belonging together. The maximum size of a namespace is limited by the size of the MetaStore. Each namespace will have one associated MetaStore and durability policy.

Namespace Encryption

When creating or editing a namespace, you can choose to encrypt data that is written to the namespace. Lattus uses AES-256-CTR (256-bit Advanced Encryption Standard in CTR mode) to encrypt the data.

The first time you enable namespace encryption, two extra fields will appear for you to create and confirm a password. This password will be used to generate a master key, and the master key will be used to encrypt the keys that are generated to encrypt the namespaces.

Caution: Please contact Quantum Support before enabling encryption for any new or existing namespace. The encryption master key(s) are stored on all three Lattus C10 Controller Nodes. Quantum recommends also keeping a backup of these key(s), which is necessary to avoid permanent and irreversible data-unavailability in the event that all three controllers are lost.

If an object is written to a namespace while encryption is enabled, the object will be encrypted at rest. If an object is written to a namespace while encryption is disabled, the object will not be encrypted.

Enabling or disabling encryption for an existing namespace does not encrypt or decrypt the objects that are already stored in the namespace. However, there are situations where existing objects may become encrypted or decrypted. This happens when existing objects are repaired (e.g., due to decommissioning a degraded disk) after the namespace's encryption settings have changed.

Example:

1. You write objects to a namespace with encryption enabled. The objects are now encrypted.
2. Later, you decide you no longer wish to encrypt data written to the namespace, so you disable encryption. The objects you originally wrote to the namespace remain encrypted. New objects written to the namespace are not encrypted.
3. One of the disks storing Check-Blocks for some of the encrypted objects becomes degraded, so you decommission the disk, which initiates a repair of the objects.
4. As the objects are repaired and their Check-Blocks re-written to other disks, they are written without encryption. Objects that did not require repair will remain encrypted.

Object

Data stored in the Lattus storage system is logically defined as an object and belongs to a specific user-created namespace.

QDynamics

QDynamics are the maintenance processes that are responsible for maintaining the data integrity on the Lattus S10 Storage Nodes. It manages all the housekeeping between the nodes such as data integrity, scrubbing, node management, self-healing, and garbage collection. QDynamics runs on each Lattus S10 Storage Node.

QSpread

QSpread is the algorithm that splits up a given block of data into multiple blocks that can be spread across multiple nodes. This is what enables Lattus to reconstruct the original data out of a sufficient number of Check-Blocks, which is smaller than the total number of Check-Blocks.

Repair (a.k.a. Repair Manager)

Repair, in the context of decommissioned and replaced fault disks, is the repair process that runs on all objects in each namespace.

Once repair manager has been started by its normal 24-hour interval, all objects requiring a repair will be put into a repair task queue for the maintenance agent to run.

Once on the repair task queue, the maintenance agent will task the storage daemons to work on the repair task queue.

Unlike traditional RAID, Lattus' advanced QSpread technology ensures greater levels of data integrity according to the user-specified durability policy. When a drive is decommissioned, Lattus regenerates redundant encodings (i.e., 'Check-Blocks') for all affected objects.This ensures full redundancy of the affected objects is restored to the level defined by the user specified policy.

The time to regenerate redundancy codes is affected by the amount of data and the number of Check-Blocks on the decommissioned drive. To ensure the best data integrity, a minimum average object size of 1.7 MiB - 6.8 MiB per MetaStore is needed to regenerate redundancy codes within acceptable time for a single site configuration (e.g. 3 - 7 hours). If you have any questions about minimum average object size and regenerating redundancy codes we recommend reviewing your configuration and data integrity requirements with Quantum Professional Services.

In a Geo-Spread environment, WAN bandwidth and latency also affect the time required to regenerate redundancy codes. When using the RepairSpread & DynamicSafety safety strategy in Geo-dispersed systems with imbalanced latencies between the data centers (e.g., 20 ms from A to B and 100 ms from A to C), the grace timeout mechanism might result in the blockstores of the data center with the largest latency always timing out. This will make all uploads succeed, but with a lower than expected safety. If the data center with the highest latency is lost, data will be retrievable but if another data center is lost, chances are high that data is not retrievable. Quantum recommends reviewing your Geo-Spread configuration and data integrity requirements with Quantum Professional Services to ensure maximum data protection.

Safety Strategy

A durability policy uses its configured safety strategy to determine if and when Lattus should store data when there is suboptimal disk safety. The available safety strategies and their descriptions are listed below:

- **DynamicSafety**: Use this option if you want your data to be stored as soon as possible, taking into account that it can be stored with a lower actual disk safety. The actual disk safety could even be set to 0, meaning that a subsequent disk failure could lead to data loss.

  Caution: It is not recommended to use this strategy. This will lead to lower durability. It is possible to upload a superblock with safety 0. If after this upload, before any repair has happened, a disk with data from this superblock becomes unavailable, the object is not retrievable anymore.

- **RepairSpread**: Use this option if you always want to store data with the targeted disk safety and have it fail when this cannot be guaranteed. This safety will try to store the object using the generated list of disks. If one or more disks didn't return within the defined timeout, a replacement set of disks will be selected. If that fails, Lattus will retry until it's no longer capable of selecting a list of disks that match your spreading strategy. This strategy provides ultimate reliability but has a performance impact when new disks have to be tried to match the disk safety adding latency and additional data sent over the wire. This strategy is recommended for setups that must have a guaranteed disk safety at all times.

  Caution: Using this strategy can result in errors reported to the client application. Quantum recommends ensuring the client application handles errors gracefully to ensure data protection. Note: StorNext handles errors by reporting RAS alerts, retrying PUTs, as well as potentially disabling paths. This handling ensures maximum data safety.

- **RepairSpread & DynamicSafety**: This safety strategy will try to store your data using a generated list of disks. If the store doesn't succeed at the first attempt within the defined timeout, Lattus will retry once with a new list of disks and if this fails again, Lattus will store the object with a suboptimal disk safety. This is the preferred strategy for systems with a large amount of disks and a spreading

strategy that can be easily matched. Chances of running into a scenario where in a second attempt no matching disks can be identified are slim.

Caution: This strategy should be avoided in multi-geo setups with imbalanced latency between sites if the imbalance causes blockstores at a remote site to timeout. In that scenario, performance and data safety will be suboptimal. Quantum recommends avoiding imbalanced latency configurations that might lead to this scenario. However, if it cannot be avoided, then the configuration should be reviewed with Quantum to determine the optimal policy based on your performance and data protection requirements.

Three-site Setup Examples Using RepairSpread and DynamicSafety Safety Strategy	
Durability Policy (spread width/disk safety)	OK when...
18/6	OK when encoded data is stored on at least 12 disks within the object time-out.
18/7	OK when encoded data is stored on at least 11 disks within the object time-out.
18/8	OK when encoded data is stored on at least 10 disks within the object time-out.

- **Stop on Failure**: Use this option when you don't want to get any performance impact of failing disks or networking issues. Storing the object will only succeed if during the first attempt all data can be stored on the generated list of disks within the defined timeout. If that doesn't work, the store operation will have to be retried from within the application. When a sufficient amount of disks are present the next store operation will avoid disks that failed in the previous run. This option is typically used for performance testing.

Caution: This is not desirable for production systems because single disk/node failures will result in failed uploads.

Small File Support

Enabling small file support is a way to make the storage and retrieval of small files from the Lattus object storage more efficient. When small file support is enabled for a namespace (or enabled for the default durability policy used by S3 buckets), Lattus will check the file size of any object as it is stored in the namespace or bucket. If the file size is equal to or less than the configured small file size threshold (max is 4,096 KiB), Lattus will store a full copy of the object on one disk in the spread, and will use the remaining disks in the spread to store the object as encoded Check-Blocks. As long as the full copy of the object is available, Lattus will retrieve the full copy from the single disk whenever the object is requested by an application. This allows Lattus to retrieve a small file more efficiently by reading it from a single disk rather than having to re-encode it from several disks.

When you enable small file support for a namespace, an additional durability policy is automatically created. This policy will have the same name as the namespace's durability policy (or S3 bucket default policy), but it will include the suffix "_full_copy". For example, if you enable small file support on a namespace that's using a durability policy named "dur_pol_20_4", your Lattus system will automatically create a policy named "dur_pol_20_4_full_copy".

To accommodate storing a full copy of small files, a full copy durability policy's spread width and disk safety are different from the original durability policy. The spread width and disk safety settings for a full copy durability policy are ORIGINAL SPREAD WIDTH - 1/ORIGINAL DISK SAFETY - 1 + OBJECT.

If the full copy cannot be accessed (e.g., the disk storing the full copy has become degraded), Lattus will re-encode the object from the Check-Blocks stored on the other disks in the spread.

Example

Let's say you enable small file support for a namespace that's using a 20/4 durability policy, and leave the Small File Size Threshold set to 512 KiB (default). Then you store a 512 KiB file in the namespace.

Lattus will encode the object and store the Check-Blocks across 19 disks (original spread width - 1) with a disk safety of 3 (original disk safety - 1) and also store a full copy of the object on another disk. Lattus will use the full copy for GET requests as long as the full copy is available. If the full copy is unavailable (e.g., the disk storing it becomes degraded), Lattus will re-encode the object from the

Check-Blocks on the other disks. Lattus needs 16 of the 19 disks to re-encode the object.

If you know you will be storing a lot of small files, you may choose to create a durability policy and namespace specifically for them. The recommended spread width/disk safety settings for the durability policy are 5/4, then enable small file support for the namespace and set the appropriate small file size threshold for your files.

Spread Width

Spread width defines the number of Lattus S10 Storage Node/disks that will be used to store an object in any namespace that uses the selected durability policy.

For example, if a namespace configured on a Lattus system containing 35 Lattus S10 Storage Nodes uses a durability policy with a spread width of 20, any object stored in that namespace will be spread across 20 of the Lattus system's 35 Lattus S10 Storage Nodes, using one disk node per node. The Lattus system determines which nodes/disks will be used. The minimum spread width is 2, and the maximum is 20. Work with your Quantum Professional Services representative to determine the appropriate settings for your Lattus system.

Storage Daemon

The storage daemon is a service/process/daemon running on each Lattus S10 Storage Node. The storage daemon receives requests from the client daemon(s) and maintenance agent(s) and acts as a gateway for requests to disks in the Lattus S10 Storage Nodes. A storage node has by default two storage daemons running. Each daemon is responsible for a set of disks on the node. Every storage daemon listens on all network segments, but listens on a unique port for that Lattus S10 Storage Node (23520, 23521). Each namespace has one storage daemon that acts as an entry point for management information on that namespace. This is referred to as the master storage daemon for that namespace.

Master Storage Daemon

Each namespace has one storage daemon that acts as an entry point for management information on that namespace. This is referred to as the master storage daemon for that namespace. The responsibilities for a master storage daemon for a namespace include troubleshooting information as well as keeping a list of objects that need to be repaired

or deleted. It enumerates objects to be repaired every four hours. A storage daemon can be the master for more than one namespace. The master storage daemon can be the master for more than one namespace.

Superblock

The Lattus encoding algorithm will chunk every object into superblocks that match the configured maximum superblock size. The maximum superblock size will also determine the maximum size of the object. In general, the larger the superblock size, the larger the object. Below are some key points about superblocks:

- **Selecting a bigger superblock size increases performance**: The time (latency) needed to encode a superblock is fairly independent from its size; hence selecting a bigger superblock size will be beneficial for throughput. Larger superblocks also lead to larger Check-Blocks that need to be sent to the storage daemons. Both the network and the storage daemons need fewer operations to transport these and persist them to disk.

  The drawback of using larger superblocks is that more memory is required for encoding. Typically, an amount of memory equal to six times the superblock size is needed per encoding operation. Taking into account that objects are encoded in multiple parallel encoding threads, and multiple objects can be encoded at the same time, the memory requirements can become significant.

- **Selecting a bigger superblock size lowers the metadata overhead**: Objects that are equal to or smaller than the superblock size create only one metadata-entry for the superblock. When there are multiple superblocks for an object, the object needs to keep track of the list of superblocks.

  On average, the additional metadata generated per superblock is 160 bytes. When there are thousands of superblocks, the size of the metadata can become considerably large. Especially in WAN scenarios, this can have an adverse effect on throughput and repair performance.

- **Selecting a smaller superblock size improves the partial read efficiency**: When a partial read of an object is completed, Lattus will decode the entire superblock that entails the partially read data. If you want to read 8 KiB from an object and the superblock size is 64 MiB, Lattus will decode 64 MiB. If the superblock size was set to 2 MiB, Lattus will only decode 2 MiB to read 8 KiB.

- **Do not decrease superblock size to match your object sizes**: If you have objects that are smaller than the superblock size, there will only be a very low penalty for these objects. In order to be able to encode these objects, their size needs to be a multitude of 8 KiB. For example, if your superblock size is 2 MiB and you want to store an object of 13 KiB, Lattus will only add 3 KiB to the object (not 2,035 KiB) in order to encode it. The added padding data (consisting of zeros) will be removed when the object is read again.

Worst Case Overall Disk Safety

Worst Case Overall Disk Safety refers to the status of the object stored with the lowest disk safety. For example, if you have just a single 20/4 durability policy configured (spread width of 20 and a disk safety of 4) and all of the disks used to store the objects are working, the Worst Case Overall Disk Safety will be 4. But if one of the one of the disks used to store an object is degraded, the Worst Case Overall Disk Safety will be 3. (Quantum does not recommend a disk safety less than 4)

If your Lattus system is using more than one durability policy, the optimal disk safety in the Worst Case Overall Disk Safety section of the Dashboard will be the lowest configured disk safety. For example, if you configured a 20/4 durability policy and a 20/11 durability policy, the highest value that could be displayed (optimal disk safety) for the Worst Case Overall Disk Safety would be 4, since the lowest configured disk safety on the system is 4. However, if only the 20/11 durability policy was configured, then the optimal disk safety would be 11; and 11 would be the highest number/number of green dots that could be displayed as the Worst Case Overall Disk Safety.

Lattus Licensing

Lattus A10 Access Node Licensing

Software licensing on the Lattus A10 Access Node includes:

- **Maintenance -** A Maintenance license verifies that your site has purchased Lattus upgrade licenses, and is required for Lattus upgrades. It is also used at run time to verify the Lattus version in the software matches what was purchased.

- **SNFS Wide Area Storage (SNFS-WAS) -** An SNFS Wide Area Storage license enables you to access Wide Area Storage features.

Lattus-M Licenses on the StorNext MDC

Lattus-M Feature Keys enable usage of Lattus S10 Storage Nodes with StorNext as a tier in a policy-managed StorNext Storage Manager archive.

Note: To use Lattus with StorNext, in addition to the purchase of the Lattus-M Feature Keys, the StorNext M662 Metadata Appliance (or customer-supplied StorNext MDC) will need to have its StorNext licensing updated to include:

- **SNSM Wide Area Storage License (SNSM-WAS)**
- **StorNext Storage Manager Licenses** to correspond to the Lattus-M Feature Keys

Lattus Documentation and Training

The following documentation and training resources are available for your Lattus system:

Documentation

The following documentation is available on the Lattus Object Storage Support page: (http://www.quantum.com/lattusdocs).

Lattus

- *Lattus Release Notes* - Provides important information, late-breaking news, and a list of known issues Lattus users should know about.
- *Lattus Site Planning Guide* - Provides component specifications and assists end-users in preparing the site for Lattus installation.
- *Lattus User Essentials* - Provides essential information and instructions about Quantum Lattus Object Storage.

Lattus A10 Access for NFS/CIFS

- *Lattus A10 User's Guide* - This guide provides instructions to assist users in monitoring, maintaining, and performing basic operations on their Lattus A10 Access Node.
- *Lattus A10 User Essentials* - Provides essential information and instructions about the Lattus A10 Access Node.

Lattus-M Integration with StorNext

- *StorNext M660 Release Notes* - Provides important information, late-breaking news, and a list of known issues StorNext M662 Metadata Appliance users should know about.
- *StorNext M660 Site Planning Guide* - Provides an overview of site characteristics, physical requirements, and environmental specifications that are essential for the installation of the StorNext M660 Metadata Appliance.
- *StorNext M660 Hardware Guide* - Describes how to identify key hardware components and provides basic operating instructions for the StorNext M660 Metadata Appliance.
- *StorNext Release Notes* - Provides important information, late-breaking news, and a list of known issues StorNext software users should know about.
- *StorNext User's Guide* - This guide provides information to assist users in performing day-to-day tasks using the StorNext software.

For the latest StorNext Metadata Appliances documentation, go to http://www.quantum.com/snmdcdocs.

For the latest StorNext Software documentation, go to http://www.quantum.com/sndocs.

Training

To access training and support materials related to your Lattus system, such as video demos, visit the Quantum Lattus Object Storage Support page on Quantum.com, and go to the Training tab (http://www.quantum.com/lattushowtos).

Lattus Limitations

The following table describes Lattus limitations you should be aware of when using the system.

Component	Limitation
Maximum Number of storage nodes per management node (environment)	120 When 40 storage daemons fit into a single rack, this implies that the maximum number of supported racks is 3.
Maximum spread width	20
Maximum number of SSDs per controller used for metastores	4
Maximum size of the SSDs used for metastores	240 GB
Maximum number of SSDs per controller used for ReadCache	6
Maximum SuperBlockSize	256 MiB
Maximum number of 10 GbE uplinks per controller (links to the customer application)	2
Maximum number of objects per namespace	This is limited by the size of the SSD storing the metadata for that namespace. Each metadata entry per superblock requires at least 200 bytes. The actual size per metadata entry depends upon the following: • The object name length • The size of your storage pool • The custom-metadata stored per object. A metadata size calculator is available in the *Lattus Installation Guide*. 10% free space on an SSD is reserved to prevent performance degradation as a result of the write amplification process.

Component	Limitation
Maximum number of namespaces per Lattus system	5000
Maximum number of objects per directory when using the WebDav API.	1024
Maximum object size	Lattus supports objects with up to 65,536 superblocks. This results in the following maximum object sizes: • Superblocksize 16 MiB 1 TiB • Superblocksize 32 MiB 2 TiB • Superblocksize 64 MiB 4 TiB • Superblocksize 128 MiB 8 TiB • Superblocksize 256 MiB 16 TiB

Chapter 10

Getting Help

This chapter describes the following tasks:

Locating the System Serial Number

You will need to provide your system serial number for technical support. Lattus system serial numbers begin with the characters **CX** and contain the characters **CKC** in the number - for example, **CX1234CKC01234**.

Here's how to find it:

From the Lattus CMC

1 Select **Dashboard > Administration > Lattus Management > Locations**.

2 In the **Locations** screen, the system's serial number will be listed in the **Name** column.

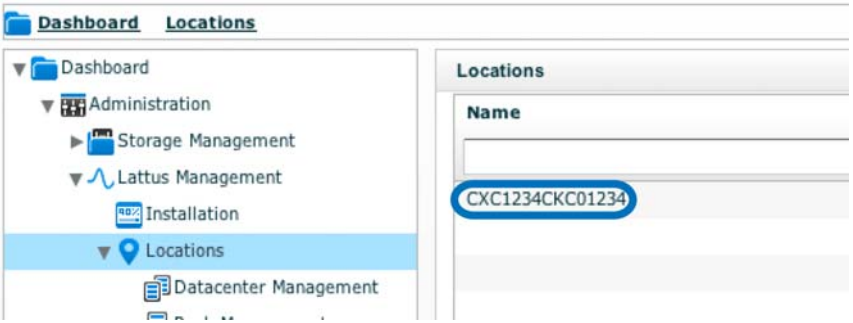

On the Lattus C10 Controller Nodes

1 From the front of any of the three Lattus C10 Controller Nodes, pull out the Service Tag.

2 The system serial number (**SYSTEM SN**) is printed on the white sticker.

Lattus C10 Controller Node

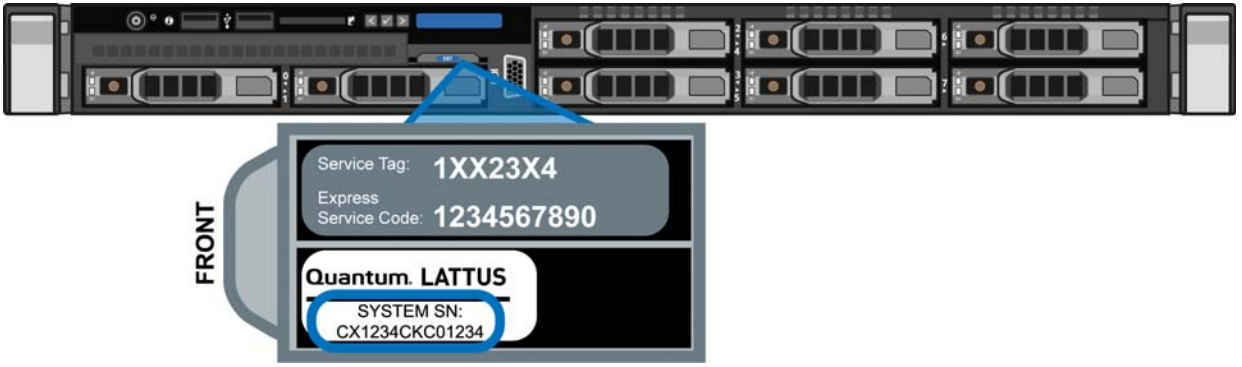

Contacting Quantum Support

You can contact Quantum Support in two ways:

- Access the Online Service Center by opening an Online Service Request at https://onlineservice.quantum.com/OA_HTML/xxibu/jtflogin.jsp
- Call Tech Support. You can find a list of contact numbers for Tech Support at: http://www.quantum.com/ServiceandSupport/Contacts/ProductSelect/Index.aspx