




Mars™ OS 1.0.2

System Administration Guide



NetApp, Inc.
495 East Java Drive
Sunnyvale, CA 94089
U.S.A.

Telephone: +1 (408) 822-6000
Fax: +1 (408) 822-4501
Support telephone: +1 (888) 463-8277
Web: www.netapp.com
Feedback: doccomments@netapp.com

Part number: 215-09727_A0
Friday, February 13, 2015

Contents

About Mars OS	12
Features of Mars OS	13
High performance	13
Industry-leading storage efficiency	13
System monitoring	13
Increased reliability and performance by using nonvolatile memory	13
Understanding clustering	15
Understanding network connections	16
Mars OS management basics	17
Creating the cluster	18
Cluster naming rules	18
Password rules	18
Cluster time	20
Modes of operation	21
Methods for managing the cluster	22
Accessing the CLI through SSH	22
Accessing the CLI through the console port	23
Accessing FlashRay System Manager	23
Browser requirements	23
Screen resolution recommended	24
Using the CLI	25
Using the CLI before creating a cluster	25
Understanding privilege levels	25
Syntax convention in this guide	25

Setting the privilege level	26
Rules for entering commands	26
Command automatic completion	27
Commands typed incorrectly	27
Keyboard shortcuts	27
Linux commands supported	28
Command output format	28
Viewing man pages	28
Displaying help	28
Configuring the cluster after cluster creation	30
Configuring the cluster (CLI)	30
Resetting the admin password (CLI)	31
Reconfiguring the cluster (System Manager)	31
Changing the admin password (System Manager)	33
Password requirements	33
Understanding the security certificate on the cluster	34
Viewing the security certificate	34
Re-generating the security certificate	34
LUN management	35
Understanding LUNs	36
Maximum number of LUNs supported	36
LUN attributes	36
Considerations for specifying the LUN name	36
Considerations for specifying the LUN size	37
LUN state	38
LUN status	38
Understanding LUN mapping	40

4 | Table of Contents

Guidelines for mapping LUNs	40
Choosing a LUN ID for the mapping	40
Understanding LUN maps	40
Creating LUNs	42
Major steps for making a LUN accessible to hosts	42
Creating a LUN (CLI)	42
Creating a LUN (System Manager)	42
Creating multiple LUNs (System Manager)	43
Managing LUNs	45
Mapping LUNs (CLI)	45
Unmapping LUNs (CLI)	46
Mapping and unmapping LUNs (System Manager)	46
Changing the LUN ID for a LUN map (CLI)	47
Changing the LUN ID for a LUN map (System Manager)	47
Showing LUN information (CLI)	47
Showing LUN mappings (CLI)	48
Showing LUN information and mappings (System Manager)	49
Increasing the LUN size (CLI)	49
Increasing the LUN size (System Manager)	49
Taking a LUN offline and bringing it online	49
Renaming a LUN	50
Deleting a LUN (CLI)	51
Deleting a LUN (System Manager)	51
Modifying or creating a comment	52
Clearing the nvfail state	52
LUNs window (System Manager)	53
Displaying LUNs	53

Igroup management	55
Understanding initiator groups (igroups)	56
How an initiator is represented in an initiator group	56
Considerations for creating an initiator group	56
Initiator group naming conventions	57
Creating initiator groups	58
Creating an initiator group (CLI)	58
Creating an initiator group (System Manager)	58
Managing initiator groups	59
Displaying initiator group information (CLI)	59
Displaying initiator group information (System Manager)	59
Modifying an initiator group (CLI)	59
Modifying an initiator group (System Manager)	61
Renaming an initiator group (CLI)	61
Destroying an initiator group (CLI)	62
Deleting an initiator group (System Manager)	62
Creating LUNs to map to an initiator group (System Manager)	62
Space management	64
Different measures of storage space	65
Understanding data compression	67
Effects of compression on write operations	67
Effects of compression on read operations	67
Understanding deduplication	68
Understanding thin provisioning	69
Benefits of thin provisioning	69
Example of a system with more provisioned space than purchased space	69
Considerations when allocating space to LUNs	70

6 | Table of Contents

Monitoring space usage	72
What happens when the cluster is low on space	72
What happens when the cluster runs out of space	72
Reclaiming space	72
Displaying capacity information (System Manager)	73
Fibre Channel Protocol (FCP) management	74
Understanding Fibre Channel Protocol (FCP)	75
Meaning of adapter and port	75
About MPIO (multipath I/O)	75
About ALUA (Asymmetric Logical Unit Access)	75
FC network limits	75
FC switch login limits	76
Hop counts supported	76
Guidelines and recommendations for setting adapter speeds	76
Accessing the FlashRay system and non-NetApp systems from the same host	77
Requirements and recommendations for FC switch configuration and zoning	78
Switch speed	78
NPIV (N_Port ID virtualization)	78
FC port type	78
FC zoning	78
Managing FCP (Fibre Channel Protocol) service	80
Displaying and changing the FCP status	80
Displaying FC adapter information	80
Configuring FC (Fibre Channel) adapters	81
Sample commands for changing the adapter speed	81
Interface management	83
Understanding interfaces	84

Default interfaces	84
When you need interface information	84
Configuring and displaying interfaces	85
Displaying information about interfaces (CLI)	85
Displaying information about initiators	86
Changing the administrative status of an interface	86
Displaying information about interfaces (System Manager)	87
Physical storage management	88
Understanding physical storage	89
Understanding RAID-DP protection	90
Understanding the RAID group on the FlashRay system	90
Understanding data reconstruction	90
What different RAID group statuses mean	90
Why it is important to replace a failed drive	91
Checking RAID group status	93
Checking RAID group status (CLI)	93
Checking RAID group status (System Manager)	94
Managing solid-state drives (SSDs)	96
Understanding SSD names	96
Displaying SSD information (CLI)	96
Displaying SSD information (System Manager)	96
Displaying disk shelf information	99
Disk shelf name format	99
Displaying shelf information (CLI)	99
Displaying shelf information (System Manager)	99
Monitoring the FlashRay system	102
Understanding events	103

8 | Table of Contents

Health monitor and events	103
AutoSupport and events	103
Format of an event message	103
Viewing the event log	103
Creating an event mark in the event log	104
Understanding the health monitor	105
Understanding alert policies	105
How alerts are generated	105
Information in an alert	105
Types of alerts supported	106
How an alert is deleted	106
Displaying alerts	107
Deleting an alert	107
Refreshing the list of alerts	107
Displaying environmental sensor information	108
Threshold-based and discrete sensors	108
Displaying environmental sensors	108
Managing AutoSupport	109
AutoSupport messages sent to NetApp	109
Configuring another AutoSupport destination	109
Reverting AutoSupport to factory settings	111
Triggering AutoSupport messages manually	111
Checking the statuses of AutoSupport messages	111
Understanding performance statistics	112
Objects, instances, and counters	112
Displaying objects, instances, and counters	113
Understanding counter explanations	114

Viewing counter explanations	114
Viewing performance data	115
Viewing performance data for a time period	115
Accessing and monitoring the FlashRay system through the Service Processor (SP)	117
Understanding the Service Processor (SP)	118
What happens when an environmental sensor detects an abnormal condition	118
How the SP is powered	118
Understanding the system event log (SEL) on the SP	118
Understanding the console log on the SP	119
Configuring the SP	120
Command for configuring the SP	120
Configuring IPv4 network settings for the SP	120
Enabling or disabling DHCP for IPv4	120
Configuring IPv6 network settings for the SP	120
Enabling or disabling router-assigned IPv6 addresses	121
Enabling IPv4 or IPv6	121
Disabling IPv4 or IPv6	121
Managing the SP from Mars OS	123
Displaying SP configuration information	123
Rebooting the SP	123
Accessing the SP CLI	124
Number of SP CLI sessions allowed	124
Accessing the SP from the console port and exiting the SP CLI	124
Accessing the SP using SSH	124
Accessing Mars OS through the SP	125
Mars OS authentication and SP authentication	125
Number of concurrent system console sessions started from the SP	125

Concurrent system console sessions	125
Returning from the Mars OS to the SP CLI	126
Using the SP CLI	127
About SP command privilege levels	127
Using the SP help	127
SP commands at the admin privilege level	127
Troubleshooting the storage system by using the SP	131
Troubleshooting an environmental sensor	131
Understanding threshold-based sensors	131
Understanding discrete sensors	132
Rebooting and shutting down	134
Rebooting the cluster	135
Understanding cluster startup	135
Shutting down the cluster	137
Understanding clean and dirty shutdowns	137
Understanding the boot loader	139
Understanding maintenance mode	140
Where to access maintenance mode	140
Entering maintenance mode	140
Using the maintenance mode CLI	141
Exiting maintenance mode	141
Troubleshooting the FlashRay system	142
What happens when the FlashRay system crashes or fails to reboot	143
Listing and deleting core dumps	143
Enabling the marsdiag login for troubleshooting	144
Establishing an sftp connection with the cluster	144
Troubleshooting service disk issues	146

Displaying information about the service disk 146

Health alerts for service disk problems 146

Destroying the cluster 148

 What destroying the cluster means 149

 Destroying the cluster 150

Copyright information 151

Trademark information 152

How to send your comments 153

About Mars OS

This chapter discusses the following topics:

Features of Mars OS	13
Understanding clustering	15
Understanding network connections	16

Features of Mars OS

Mars OS is the software running on the FlashRay system, which stores data on an all-SSD (solid-state drive) shelf.

High performance

Mars OS delivers high performance in the form of consistent and predictable low latency and high throughput. For small random access workloads, throughput is measured in I/O operations per second (IOPS); for sequential and large random workloads, it is measured as bandwidth (for example, GB/s).

Industry-leading storage efficiency

Mars OS offers these features to maximize efficiency:

- Instead of storing data in fixed-size data blocks, Mars OS stores data in variable-sized extents. An extent is a unit of storage space allocation for data in Mars OS. Extents reduce disk fragmentation and enable Mars OS to implement data deduplication and compression more efficiently, thereby giving you more storage space savings.
- Inline deduplication, which is always in effect, eliminates redundant data, guaranteeing that only a unique instance of data is stored in the FlashRay system.
- Inline data compression, which is always in effect, further reduces the physical space required to store data by making extents smaller before data is written to the SSDs.
- Thin provisioning enables Mars OS to maintain a common unallocated storage space that is readily available to any LUNs. A LUN is an application data container you define on the FlashRay system. You can expand a LUN as its storage needs change.

System monitoring

Mars OS continuously monitors the health of the FlashRay system such as the wear level of SSDs and the temperature in the chassis. It takes action by sending messages to technical support, generating alert messages to notify you of the problem, or powering off the system when conditions exceed predetermined safety limits.

Increased reliability and performance by using nonvolatile memory

The NVRAM (nonvolatile memory) module is a battery-backed hardware component on the controller. It enables the FlashRay system to maintain low latency by acknowledging client write requests quickly before data is written to SSDs.

If there is an AC power failure or dirty shutdown, data not yet written to the SSDs is copied (destaged) from the onboard DIMM (dual inline memory module) to the flash memory on the NVRAM module. The destaged contents can be saved in flash memory indefinitely. When AC power resumes, the contents are restaged into the DIMM and then written to the SSDs.

Mars OS also uses the NVRAM to batch write requests into smaller numbers of writes. This reduction in writes can prolong the life of the SSDs.

Related topics

[Space management on page 64](#)

[Monitoring the FlashRay system on page 102](#)

Understanding clustering

A cluster is the logical entity you create before you can use the FlashRay system for storing data. It is also the entity through which you manage your system. In the user interface and in this guide, all operations described are performed at the cluster level unless stated otherwise.

Each FlashRay controller, together with the attached storage, is called a node. Multiple nodes can be attached to one another to form a cluster. This release supports only one node per cluster.

In a future release, nodes of a cluster can be connected to form high availability (HA) groups. When a node in an HA group fails, the other nodes take over its resources to provide continued data service to clients.

Clustering offers these advantages:

- Traditionally when your storage need increases, you upgrade your controller, memory, or storage space on each node. However, there are limits to their expandability. Clustering provides you with a better alternative: it enables you to add nodes to a cluster to accommodate data growth.
- Members of a cluster pool their resources and uniformly distribute data across the cluster. Controller performance and storage utilization are balanced automatically for maximum efficiency, and the SSDs wear uniformly.
- Clustering simplifies storage administration because you manage a cluster as a single entity. When new nodes are added to a cluster, there is no need to update clients to point to the new nodes. The existence of the new nodes is transparent to the clients.

Related topics

[Creating the cluster on page 18](#)

[Configuring the cluster after cluster creation on page 30](#)

Understanding network connections

There are several types of network connections on the FlashRay system.

Types of networks	Description
Data network	This is the network for user data. FC (Fibre Channel) adapters on the controller are connected to the data network through FC switches. The FC ports used for the data network are converged network adapter (CNA) ports, and they are named 0e, 0f, 0g, and 0h.
Management network	All management traffic passes over this Ethernet network. For example, when you use SSH or FlashRay System Manager to connect to the FlashRay system, you communicate with it over the management network. The management port is used for the management network. A wrench icon is marked next to it on the controller.
Cluster interconnect	It is a private, dedicated, redundant network used for communication among HA (high availability) groups within a cluster. This type of network is applicable only when the HA feature is supported in a future release.

Related topics

[Understanding Fibre Channel Protocol \(FCP\) on page 75](#)

[Understanding interfaces on page 84](#)

Mars OS management basics

This chapter discusses the following topics:

Creating the cluster	18
Cluster time	20
Modes of operation	21
Methods for managing the cluster	22
Using the CLI	25
Configuring the cluster after cluster creation	30
Changing the admin password (System Manager)	33
Understanding the security certificate on the cluster	34

Creating the cluster

Before you can use the FlashRay system, you must configure the controller to be a member of a cluster. If the controller is the only or the first node of the cluster, you must create a cluster on it before using it.

It is strongly recommended that you use FlashRay Setup Wizard to create a cluster. Both the Setup Wizard and the Planning Worksheet, which you should fill out before you run the Setup Wizard, are available on the FlashRay page of the NetApp Support Site.

The admin user is the default user on the cluster. You cannot change the name of this account. Managing the cluster either through the CLI or FlashRay System Manager requires that you log in as admin.

This section describes the method to create the cluster with the CLI (command line interface) on the serial console to the point where you can start provisioning storage space.

Before you begin

You must decide on the cluster name and the password for the user named `admin`.

NOTE: The cluster name cannot be changed after cluster creation.

Cluster naming rules

Cluster names must follow these rules:

- Minimum length is two characters; maximum length is 63 characters.
- Legal characters are:
 - A through Z and a through z (case-sensitive)
 - 0 through 9
 - hyphen (-)
- The first and last characters must not be a hyphen.

Password rules

The password for the admin account must follow these rules:

- Minimum length is eight characters; maximum length is 128 characters.
- Legal characters are:
 - A through Z and a through z (case-sensitive)
 - 0 through 9
 - ! @ # \$ % ? ^ ~

Steps

1. On the serial console, log in as admin.
2. Enter the default password:

```
changeme
```

3. Enter the following command to create the cluster:

```
cluster create -name string -password string
```

The `-name` parameter specifies the cluster name. It is mandatory. The `-password` parameter specifies the password for the admin account.

If you prefer not to type a password that is echoed on the screen, you can omit the `-password` parameter. When you are prompted for the password, you can enter it without seeing it displayed on the screen. Then you are prompted to enter the password again for confirmation.

You can include an optional `-node-serialnums` parameter, which specifies the system serial numbers of the controllers to be included in the cluster. Because this release supports only one controller per cluster, this parameter is not needed.

4. After you create the cluster, the cluster has a name and you can log in to it on the serial console as admin. However, you cannot access it through SSH or FlashRay System Manager. To configure its network settings, follow the instructions in [Configuring the cluster after cluster creation on page 30](#).

NOTE: You can execute the `cluster create` command only on a system that has never been set up. If you have used the Setup Wizard or run the `cluster create` command on the controller, entering the `cluster create` command causes the following error message:

```
Cluster already exists
```

Related topics

[Methods for managing the cluster on page 22](#)

[Configuring the cluster after cluster creation on page 30](#)

[What destroying the cluster means on page 149](#)

Cluster time

The cluster uses UTC (Coordinated Universal Time) for its logs. All times displayed in the CLI are in UTC.

FlashRay System Manager displays times in the local timezone used by the browser.

Related topics

[Understanding clustering on page 15](#)

Modes of operation

The FlashRay system operates in one of two modes after it is powered on: operational mode and maintenance mode.

Operational mode

Under normal circumstances, the FlashRay system is in operational mode. In this mode, the cluster can serve data to network clients, and the management port is accessible through SSH or the System Manager. Most tasks described in this guide are carried out in operational mode.

In operational mode, the system prompt at the normal privilege level is as follows:

```
mars>
```

The system prompt at the diag privilege level is as follows:

```
mars (diag)>
```

For information about privilege levels, see [Using the CLI on page 25](#).

Maintenance mode

When the cluster encounters a problem that prevents it from working properly, you can put it in maintenance mode, so that you can troubleshoot the problem. In maintenance mode, data is not served to network clients, and a different set of commands is available than in operational mode. You can access maintenance mode commands only through the system console, not SSH or the System Manager. The system prompt in maintenance mode is as follows:

```
mars - maintenance>
```

Related topics

[Using the CLI on page 25](#)

[Shutting down the cluster on page 137](#)

[Understanding maintenance mode on page 140](#)

Methods for managing the cluster

You manage the FlashRay system as a cluster. For more information about clustering, see [Understanding clustering on page 15](#). The following list describes the methods for managing the cluster:

- FlashRay System Manager, which presents a graphical user interface (GUI) with online help.
- Command line interface (CLI). For detailed information about individual commands, see the *Mars OS Command Reference Guide* or the man pages included in Mars OS.
- Application programming interface (API).

Some system administration tasks require the CLI and some require FlashRay System Manager. For example, to view performance statistics in graphical formats, you must use FlashRay System Manager; to browse log files on the cluster such as the event log, you must use the CLI.

NOTE: Managing the cluster requires that you use the cluster management IP address, as shown in the `cluster management-interface show` command. Do not use the controller's IP address, as shown in the `network port show` command, for SSH connections to the cluster. The controller's IP address is for a single node only, not for the cluster.

Accessing the CLI through SSH

You can access the CLI through SSH at the cluster management IP address (the address you assign to port e0M on the controller when you create the cluster). The following list provides information about SSH connections:

- The SSH client must support the SSH v2 protocol. Mars OS does not support SSH v1.
- If you use a Windows host to access the CLI, you can use any third-party utility that supports SSH v2 (for example, PuTTY).
- The SSH server on the FlashRay system supports connections that use AES (Advanced Encryption Standard) ciphers with 128-bit, 192-bit, or 256-bit encryption keys.

Steps

1. Enter the command on the client as in the following example:

```
ssh admin@10.72.137.28
```

In this example, the command requests an SSH connection from the admin user to the cluster management address, 10.72.137.28.

Alternatively, you can use the cluster name in the `ssh` command if the name has been defined on

your name server.

2. When prompted, enter the admin password you created in the cluster creation procedure.

Accessing the CLI through the console port

Steps

1. Attach an ASCII terminal or a computer running a terminal emulation program such as HyperTerminal to the console port. The following table lists the communication parameter settings for the port.

Parameter	Setting
Baud	9600
Data bit	8
Parity	None
Stop bits	1
Flow control	None

2. Log in as admin with the password you created in the cluster creation procedure.

Accessing FlashRay System Manager

The System Manager is accessible on the management interface through the HTTPS or HTTP protocol. It accepts connections on the default ports, which are port 443 for HTTPS and port 80 for HTTP.

Browser requirements

The following browsers are required:

- Internet Explorer 10 or later, with compatibility mode disabled
- Chrome 37
- Firefox 31 and Firefox 24 (ESR versions)
- Safari 7.0

NOTE: Other browsers might work but are not tested.

Screen resolution recommended

The System Manager is best viewed with a screen resolution of 1,280 x 1,024 pixels. Minimum resolution is 1,240 x 800 pixels.

Steps

1. Point your browser to the cluster management address or the cluster name. You can omit the port number in the URL. For example, if the address is 10.72.137.28, enter one of the following URLs in your browser:

```
https://10.72.137.28
```

```
http://10.72.137.28
```

The FlashRay System Manager login screen appears.

2. If this is the first time you use your computer to access FlashRay System Manager, the browser displays a message similar to the following:

```
The security certificate presented by this website was not issued by a
trusted certificate authority.
```

Accept the certificate and proceed.

NOTE: Different browsers display different warning messages when they are presented with an untrusted security certificate.

3. Log in as admin by using the password you created in the cluster creation procedure.

Related topics

[Understanding interfaces on page 84](#)

[Using the CLI on page 25](#)

Using the CLI

Commands entered on the CLI take effect on the entire cluster unless noted otherwise. For more information about clusters, see [Understanding clustering on page 15](#).

Using the CLI before creating a cluster

You can execute a limited number of commands before you create the cluster (for example, the `cluster create` command to create a cluster and the `system configuration show` command to display controller configuration information). Most commands such as the ones for managing storage space are not meaningful on an uninitialized system.

After you create the cluster, you can use all Mars OS commands except the `cluster create` command.

Understanding privilege levels

Commands are arranged into privilege levels as described in the following table.

Privilege level	Types of commands
normal	Commands for common or routine tasks, such as those for managing network adapters and provisioning storage for applications. Most commands are at this level. It is the default privilege level.
diag (diagnostic)	Commands with potentially disruptive results or with output requiring advanced knowledge to interpret. Use them only with guidance from technical support to diagnose and fix problems.

Syntax convention in this guide

In this guide, a vertical bar (|) in a command syntax separates mutually exclusive values for a parameter. Do not confuse this character with the vertical bar to pipe a command output to the `more` or `grep` command. For more information about redirecting command output, see [Linux commands supported on page 28](#).

When referencing an instance of an object (for example, a disk shelf), some commands accept a name, a UID (unique identifier), or a UUID (universally unique identifier). UIDs and UUIDs are assigned by Mars OS and can be displayed. You can use them when writing programs for manipulating instances of objects on the FlashRay system. When you enter a command interactively, unless a UID or UUID is mandatory, it is more convenient to type the name because it is shorter and mnemonic.

For commands that accept a name, a UID, or UUID, this guide includes only the `-name` parameter in the syntax.

Setting the privilege level

Steps

1. Enter one of the following commands:

- `set -privilege normal`
- `set -privilege diag`

If you change the privilege level to `diag`, it changes back to the normal level the next time you start a CLI session.

Rules for entering commands

Most commands include one or more parameters. The parameters require you to specify a value for them. A few rules exist for specifying values:

- A value can be a number, a Boolean specifier, a selection from an enumerated list of predefined values, or a text string. The following examples illustrate different types of parameter values:
 - `event log show -lines 10`
 - `fcp interface modify -name fcl -status-admin up`
 - `lun create -name lun1 -size 3g`
 - `network fcp adapter modify -name 0h -speed 8G`
- If a parameter accepts a Boolean value (true or false), you can omit the value, and the parameter is set to true. For example, the following commands have the same effect:
 - `system service-processor network modify -ipv4-enable`
 - `system service-processor network modify -ipv4-enable true`
- Some parameters can accept a comma-separated list of two or more values.
- Whenever you specify text containing a space, you must enclose the entity in quotation marks as in the following example:

```
event mark -comment "Beginning of performance test"
```
- The CLI interprets a question mark (?) as the command to display help information. (The question mark is not echoed on the screen.) As a result, when you want a question mark in a command to be interpreted literally, you must enclose it in quotation marks.
- You can type commands and parameters in uppercase or lowercase. The letters are converted to

lowercase.

Parameter values are case-sensitive. For example, LUN names `lun1` and `LUN1` represent different LUNs. The exceptions are the units for storage space and the hexadecimal values in addresses or WWPNs (World Wide Port Names). For example, if you type `10G` and `10g`, Mars OS considers them to be the same amount of storage space. If you type `11:22:33:44:55:66:77:aa` and `11:22:33:44:55:66:77:AA` for a WWPN, Mars OS considers them to be the same WWPN.

Command automatic completion

If you do not want to enter the entire command, you can press the Space or Tab key to have Mars OS complete the command. If the partial command entered is not unique, the possible commands are displayed.

For example, if you type `system config` and press the Tab or Space key, Mars OS finishes typing the entire command as follows:

```
system configuration show
```

If you type `igroup` and press the Tab or Space key, Mars OS displays the possible commands that you might want to enter: `add`, `create`, `destroy`, `modify`, `remove`, `rename`, and `show`.

Commands typed incorrectly

If you mistype a command, you cannot use the Tab, Space, or Enter key to complete the command. The terminal bell rings to alert you to the mistake. You must correct the typing error before trying to execute the command.

For example, if you type `ssystem`, you cannot press any key for automatic completion or the Space key to continue typing the subcommand or parameter name. You must correctly type `system` before typing the rest of the command and pressing Enter.

Keyboard shortcuts

The following list describes how to use keyboard shortcuts on the CLI:

- Ctrl-A moves the cursor to the beginning of the line.
- Ctrl-E moves the cursor to the end of the line.
- Left and right arrow keys move the cursor left and right.
- Up arrow key retypes the previous command.
- Backspace key deletes characters.

For a complete list of movement keys, enter the `help` command.

Linux commands supported

Mars OS supports these Linux commands on the CLI: `history`, `grep`, and `more`. For example, you can enter the `history` command to see the most recent commands you entered.

You can also pipe Mars OS command output to the `more` or `grep` command. The following examples show how these commands are used:

- This command displays the output of the `system configuration show` command one screen at a time, and you can press the Space key to display the next screen:

```
system configuration show | more
```

- This command displays the `lun show` command output to include only lines containing the string `lun1`:

```
lun show | grep lun1
```

NOTE: The `help` and `history` commands are exceptions. The output of these commands cannot be redirected to the `more` or `grep` command.

Command output format

Instead of viewing command output in standard format, you can optionally include the `-format` parameter to display or transmit it in another format. The supported format is JSON (JavaScript Object Notation). For example, the following command displays LUN information in JSON:

```
lun show -format json
```

The JSON parameters are consistent with the ones used in the REST API, which is the programming interface to the cluster.

Viewing man pages

Man pages provide detailed information about commands at the normal and diagnostic (`diag`) levels. The following is an example of displaying a man page for the `lun create` command:

```
man lun create
```

You can also display the man page for a category of commands. For example, you can enter `man lun` to display a man page about the subcommands under `lun`.

Displaying help

You can display help information about the CLI in the following ways:

- Typing a question mark (?) displays a list of commands available at your current privilege level.
- Typing a command followed by a question mark (?) displays help about the command. The information depends on when you type the question mark. For example, if you type it after `lun create`, a list of parameters is displayed. If you type it after `lun create -name`, only an explanation for the `-name` parameter is displayed.
- Entering the `help` command displays information about how to use different keys for command automatic completion and for moving the cursor.

Related topics

[Methods for managing the cluster on page 22](#)

Configuring the cluster after cluster creation

This section describes:

- Configuring the cluster so that its management interface is accessible on the network after you create the cluster by using the `cluster create` command.
- Using FlashRay System Manager to modify the configuration of your cluster after you create the cluster by using FlashRay Setup Wizard.

NOTE: You cannot rename your cluster. Changing the name requires that you destroy the cluster and create it again with FlashRay Setup Wizard or the `cluster create` command. You can destroy a cluster only by using the CLI, not the System Manager.

Configuring the cluster (CLI)

To specify...	Use this command
DNS domain for the cluster or DNS server addresses	<pre>network dns-server modify -domain <i>string</i> -ipv4-address-list <i>address</i>,...</pre> The <code>-domain</code> parameter specifies the domain name. The <code>-ipv4-address-list</code> parameter specifies IPv4 addresses of one or two domain name servers.
Cluster management interface settings (IP address, default gateway address, and netmask)	<pre>cluster managment-interface modify -ip-address <i>string</i> -gateway <i>string</i> -netmask <i>string</i></pre> All parameters are mandatory.
Admin password if you know the existing password	<pre>security login modify</pre> The command prompts you for the existing password and then the new password. The two passwords must be different. For password requirements, see Creating the cluster on page 18.

After you configure the cluster, it is strongly recommended that you perform these tasks:

- Configuring the service processor (SP). For information about the SP, see [Understanding the Service Processor \(SP\) on page 118](#).
- Enabling the transmission of AutoSupport messages to NetApp. For information about AutoSupport, see [Managing AutoSupport on page 109](#).

Resetting the admin password (CLI)

If you forget the admin password, you can reset it.

Before you begin

You can reset the admin password only when the cluster is in maintenance mode. Also, the command for resetting the password must be entered on the controller's serial console. You cannot reset the password through the `system console` command on the Service Processor (SP).

Steps

1. On the console, enter the following command to boot the cluster to maintenance mode:

```
cluster reboot -maintenance
```

2. Enter `y` when the following confirmation message is displayed:

```
Are you sure you want to reboot the cluster?(y/n)
```

3. At the maintenance mode prompt, enter the following command:

```
security adminlogin reset
```

If currently you have a console session through the `system console` command on the SP, the session is disabled. It is re-enabled after this command is finished.

4. You are prompted to press Ctrl-C to change the password or wait five seconds to terminate the `security adminlogin reset` command.

Press Ctrl-C to continue.

5. When prompted, enter the new admin password.
6. Re-enter the new password for confirmation. The password is reset and you can use it after the next reboot.

7. Enter the following command to boot the cluster to operational mode:

```
continue-boot
```

Reconfiguring the cluster (System Manager)

Steps

1. Go to **Manager > CLUSTER > Basic Settings**.
2. Click **Edit**.
3. Enter the new values for the cluster settings. Then click **Apply**.

NOTE: Some types of information in this window such as Serial Number, OS Version, Protocol Available, and Number of Nodes are for display only and cannot be modified. This version of Mars OS supports only the Fibre Channel protocol and one node per cluster.

4. If you want to change the admin password, see [Changing the admin password \(System Manager\) on page 33](#).

NOTE: If you modify the IP address, the System Manager asks if you want to reconnect by using the new address. If you choose not to reconnect immediately, the change takes effect the next time you launch the System Manager by pointing your browser to the new IP address.

Related topics

[Creating the cluster on page 18](#)

[What destroying the cluster means on page 149](#)

[Configuring the SP on page 120](#)

[Understanding maintenance mode on page 140](#)

Changing the admin password (System Manager)

You can use FlashRay System Manager to change the admin password that you initially set when you created the cluster.

NOTE: This procedure assumes that you have the existing password. If you forget your password, you cannot use FlashRay System Manager to create a new one. You must use the CLI on the serial console to access maintenance mode to create a password.

Password requirements

When you change the password, make sure the new password meets these requirements:

- The password is case sensitive.
- The password must have eight to 128 characters.
- The password can contain these characters:
 - Letters from A through Z and a through z.
 - Digits from 0 through 9.
 - These special characters: ! @ # \$ % ? ^ ~

Also, the new password must not be the same as the old password.

Steps

1. In the top menu bar, click the **admin** button to display the pull-down menu.
2. Click **Change Password**.
3. In the dialog box, enter the old password and the new password.
4. Enter the new password again to confirm it.
5. Click **OK**.

Related topics

[Configuring the cluster after cluster creation on page 30](#)

Understanding the security certificate on the cluster

Mars OS generates a self-signed security certificate during the cluster creation process. This certificate must be validated by the web browser before you can use the System Manager to manage the cluster. For you to use the Mars OS REST API to access the cluster, the API client must accept the certificate for the SSL session.

On the CLI, you can view various fields in the certificate. Also, you can re-generate the certificate when it is about to expire or change the information in the certificate such as the cluster location. The certificate generated by Mars OS is valid for 365 days.

Viewing the security certificate

Steps

1. To view the security certificate, enter the following command:

```
security certificate show
```

Re-generating the security certificate

Steps

1. To re-generate the security certificate, enter the following command:

```
security certificate modify -parameter value
```

Each parameter is a field in the certificate. Changing the value of a field generates a new certificate.

To see a parameter list, enter the following command:

```
security certificate modify ?
```

Related topics

[Methods for managing the cluster on page 22](#)

LUN management

This chapter discusses the following topics:

Understanding LUNs	36
Understanding LUN mapping	40
Creating LUNs	42
Managing LUNs	45
LUNs window (System Manager)	53

Understanding LUNs

A LUN is a logical unit (a data container) that appears as a disk to SAN (Storage Area Network) hosts. You must create a LUN on the cluster before you can use the cluster to store data. Data in a LUN is distributed uniformly across all SSDs in the cluster. Therefore you can consider a LUN a virtual disk.

Maximum number of LUNs supported

You can create up to 250 LUNs on a cluster.

LUN attributes

Before you create a LUN, you need to make decisions on these LUN attributes:

- LUN name
- LUN size

NOTE: There is another attribute, maximum extent size, which is for performance tuning. You should keep its default setting unless you are instructed by technical support to change it.

Considerations for specifying the LUN name

It is best to choose a name that reflects how the LUN is used (for example, a name related to the application that uses the LUN or the type of data it stores). LUN names must follow these rules:

- Maximum length is 255 characters.
- Legal characters are:
 - A through Z and a through z (case-sensitive)
 - 0 through 9
 - hyphen (-)
 - underscore (_)
 - braces ({})
 - period (.)
 - colon (:)
- Spaces are not allowed.

Considerations for specifying the LUN size

The LUN size you specify represents the maximum space the LUN might occupy. It is the space presented to the client application using the LUN. Due to thin provisioning, Mars OS does not immediately allocate the requested space to the LUN; space is allocated as the application writes to the LUN. For example, if you specify the LUN size to be 1 TiB, Mars OS allows the client application to use up to 1 TiB of space provided that space is available. However, actual space allocation happens only when the application writes data to the SSDs.

To allocate space efficiently, it is important to understand concepts discussed in [Different measures of storage space on page 65](#).

NOTE: You can increase, but not decrease, the size of a LUN after its creation.

Units used for LUN sizes (CLI)

When entering commands with a LUN size, use one of the following units:

- k (kibibytes)
- m (mebibytes)
- g (gibibytes)
- t (tebibytes)

To specify in bytes, omit the unit. For example, to specify a LUN size of 512 bytes, type 512 in the command.

Maximum provisioned space

The total provisioned space on a cluster is 150 TiB. When you add up all the LUN sizes on your cluster, the sum must not exceed this limit. An error message is displayed if you try to create a LUN that would cause the sum of LUN sizes to be greater than 150 TiB.

Maximum and minimum LUN size

The maximum size supported is 64 TiB.

NOTE: Make sure that the host operating system supports the LUN size you specify.

The minimum size supported is 512 bytes. If you specify a value less than 512 bytes, Mars OS automatically rounds it up to 512 bytes.

Fractional LUN sizes

Decimal numbers are allowed when you specify a LUN size. Mars OS converts a fractional quantity to a whole number of bytes when creating or resizing a LUN. For example, if you specify 1.1234m (1.1234 MiB), the resulting LUN size is 1,177,970 bytes. The number of bytes is then rounded up to the nearest multiple of 512.

The CLI does not limit the number of decimal places in the LUN size. Mars OS calculates the byte quantity based on the exact number you enter on the CLI. However, if you use the FlashRay Setup Wizard or System Manager to create or resize a LUN, the size specified is limited to three decimal places. For example, 1.1115 GiB is rounded up to 1.112 GiB, and 1.1114 GiB is rounded down to 1.111 GiB. After the number is rounded, Mars OS converts it to the number of bytes.

LUN state

The LUN state, which can be online or offline, indicates whether the LUN is accessible to client applications.

A LUN must be online to be accessible to client applications. After you create a LUN, it is online by default. You can map or unmap a LUN when it is online or offline. For information about mapping and unmapping LUNs, see [Understanding LUN mapping on page 40](#).

You can put a LUN in the offline state to prevent the LUN from being accessed by client applications. Mars OS can also put a LUN in the offline state when it changes the LUN status to no-space or corrupt.

LUN status

The LUN status is set by Mars OS to inform you of the condition of the LUN, which might have an effect on the LUN's online or offline state. Taking corrective actions can clear a LUN's status.

Normal status

When there are no problems with the LUN, the LUN status is normal.

Nvfail status

The nvfail status is set by Mars OS if the cluster detects a problem at reboot that prevents data from being written from the NVRAM to the SSDs. Client applications cannot read from or write to a LUN with this status, regardless of the LUN state. As a result, the LUN is protected against possible corruption that could result from an application's use of stale data.

When you see a LUN with nvfail status, you must take corrective actions to restore the client application data to a known good state because data in some write requests sent before the NVRAM failure might have been lost. The exact corrective steps depend on the client application. In general, the steps are as follows:

1. Stop the application so that it does not try to access the LUN.
2. Clear the nvfail status. When the LUN is in the online state, you can check data consistency and correct any problems, if found. For information about clearing the nvfail status, see [Clearing the nvfail state on page 52](#).

NOTE: To clear the nvfail status, you must use the `lun modify` command on the CLI. The System Manager does not provide a way to clear it.

3. Restart the application. It can now access the LUN, which contains the restored data.

No-space status

When an attempt to write to a LUN fails because the cluster is out of space, Mars OS sets the LUN status to no-space and the LUN state to offline. See [Monitoring space usage on page 72](#) for information about how the level of available space on the cluster affects a LUN.

Corrupt status

When Mars OS detects data corruption in a LUN, it sets the LUN status to corrupt and the LUN state to offline. Call technical support when you see a LUN with this status.

Related topics

[Creating LUNs on page 42](#)

[Managing LUNs on page 45](#)

[Understanding LUN mapping on page 40](#)

Understanding LUN mapping

Each LUN has to be mapped to an initiator group (igroup) before host initiators can access data in the LUN. If you have never created an initiator group, see [Understanding initiator groups \(igroups\) on page 56](#) for more information before proceeding.

Guidelines for mapping LUNs

After creating initiator groups, you need to decide which LUNs should be mapped to the initiator groups. Follow these guidelines for mapping LUNs to initiator groups:

- You can map a LUN to an initiator group only once. You cannot map a LUN to the same igroup multiple times. For example, if you have already mapped lun1 to igroup1, you cannot map lun1 to igroup1 again.
- You can map a LUN to up to 4,096 initiator groups.
- You can map a LUN only once to a specific initiator through the initiator group. For example, if the WWPN 11:22:33:44:55:66:77:88 is in igroup1 and igroup2, you cannot map lun1 to both igroup1 and igroup2.
- Unique LUN IDs must be used for multiple mappings within the same initiator group.

Choosing a LUN ID for the mapping

The LUN ID is a number used by initiators in an igroup to identify a LUN. Therefore, a LUN ID must be unique within an igroup. For example, if two LUNs are mapped to the same igroup, the two mappings must use different LUN IDs.

A LUN ID is not an attribute of a LUN. It is associated with a LUN mapping. For example, the LUN ID used for mapping lun1 to igroup1 can be 1 while the LUN ID used for mapping the same LUN (lun1) to igroup2 is 2.

You can assign a LUN ID when mapping a LUN to an igroup or accept the default assigned by Mars OS. The LUN ID can range from 0 to 4,095.

NOTE: If you change the LUN ID when a client application is accessing the cluster, the I/O operation might fail. It is recommended that you keep the LUN ID after you create the LUN. If you must change it, do so only when no client applications are using the cluster.

Understanding LUN maps

After you map a LUN to an initiator group, Mars OS creates a LUN map, which is a set of values that define the relationship among the initiator group, the LUN, and the LUN ID. The following are examples of LUN maps created after you map lun1 and lun2:

- {igroup1, lun1, 1}
- {igroup1, lun2, 2}
- {igroup2, lun1, 1}

Related topics

[Managing initiator groups on page 59](#)

[Understanding LUNs on page 36](#)

[Creating LUNs on page 42](#)

[Managing LUNs on page 45](#)

Creating LUNs

Before you follow the instructions in this section to create a LUN, be sure to read [Understanding LUNs on page 36](#) to understand what decisions you need to make in advance.

Major steps for making a LUN accessible to hosts

These are the major steps required to enable hosts to access data in a LUN:

1. Create a LUN.
2. Create an initiator group.
3. Map a LUN to an initiator group.

Creating a LUN (CLI)

Steps

1. Enter the following command:

```
lun create -name name -size size
```

You can omit the unit in *size*, or specify k (kibibytes), m (mebibytes), g (gibibytes), or t (tebibytes).

If *size* is less than 512, Mars OS rounds the value up to 512 bytes.

The following examples show how to create LUNs of various sizes:

- `lun create -name lun1 -size 500`

A LUN named `lun1` is created. Because its specified size is 500 bytes, Mars OS rounds up the LUN size to 512 bytes.

- `lun create -name lun2 -size 550`

A LUN named `lun2` is created. Because its specified size is 550 bytes, Mars OS rounds up the LUN size to 1 KiB.

- `lun create -size 1g -name database1`

A LUN named `database1` is created. Its size is 1 GiB.

Creating a LUN (System Manager)

Steps

1. Go to the **Manager > STORAGE > All Storage > LUNs** window.
2. Click the **Create LUNs** button.

3. In the **Create LUNs** dialog box, select 1 as the number of LUNs to be created.
4. Enter the LUN name and size. Then click **Next**.
5. Select the initiator groups to be mapped to the LUN in one of the following ways:
 - Click the checkbox next to Available to select all initiator groups.
 - Click the checkbox next to each initiator group name.
 - Enter a string in the search box to display initiator group names with the string, and then select all or some of the initiator groups.
6. Click  to move the initiator groups to the Selected section.
 To move an initiator group back to the Available section, click the checkbox next to the initiator group name in the Selected section and click . To move all initiator groups back to the Available section, click .
7. If you want to choose the LUN ID for mapping the newly created LUN to the selected initiator group, enter the LUN ID. By default, the LUN ID is automatically assigned.
8. Click **Next**.
9. Make sure that the values in the confirmation dialog box are correct. If not, click **Back** to make corrections. Otherwise, click **Finish**.
10. Click **Close**.

Creating multiple LUNs (System Manager)

Steps

1. Go to the **Manager > STORAGE > All Storage > LUNs** window.
2. Click the **Create LUNs** button.
3. In the **Create LUNs** dialog box, select the number of LUNs to be created.
4. If you want to manually name the LUN, go to step 5. If you want the LUNs to be named automatically and to have the same size, follow these steps:
 - a. Select **Auto-name and size**.
 - b. Enter a prefix for the LUN name. For example, if you use eng as the prefix, all LUN names start with eng.
 - c. Select a suffix for the LUN name to number the LUNs. The suffix can be one to four digits. For example, it can be 1, 01, 001, or 0001.

- d. Enter the number to start the suffix from. For example, if you select a two-digit suffix and start with 0, the first LUN is numbered 00, the second is numbered 01, and so on.
 - e. Enter the LUN size.
 - f. Click **Next**. Then go to step 8.
5. Select **Manually name and size**.
 6. Enter the name and size for each LUN.
 7. Click **Next**.
 8. Select the initiator groups to be mapped to the LUN in one of the following ways:
 - Click the checkbox next to Available to select all initiator groups.
 - Click the checkbox next to each initiator group name.
 - Enter a string in the search box to display initiator group names with the string, and then select all or some of the initiator groups.
 9. Click  to move the initiator groups to the Selected section.

To move an initiator group back to the Available section, click the checkbox next to the initiator group name in the Selected section and click . To move all initiator groups back to the Available section, click .
 10. If you want to choose the LUN ID for mapping a newly created LUN to the selected initiator group, enter the LUN ID. By default, the LUN ID is automatically assigned.
 11. Click **Next**.
 12. Make sure that the values in the confirmation dialog box are correct. If not, click **Back** to make corrections. Otherwise, click **Finish**.
 13. Click **Close**.

Related topics

[Understanding initiator groups \(igroups\) on page 56](#)

[Creating initiator groups on page 58](#)

[Understanding LUN mapping on page 40](#)

Managing LUNs

Before you follow the instructions to manage a LUN, be sure to read [Understanding LUNs on page 36](#) to understand what decisions you need to make in advance.

For each existing LUN, you can perform the following tasks:

- Map and unmap
- Change the LUN ID for a LUN map
- Show
- Resize

NOTE: You can increase, but not decrease, the size of a LUN.

- Take offline and bring online
- Rename
- Delete
- Create or modify a comment
- Clear nvfail state

Mapping LUNs (CLI)

Steps

1. To map a LUN to an initiator group, enter either of these commands:

- `lun map -igroup string -name string -lun-id integer`
- `lun mapping create -igroup string -name string -lun-id integer`

For example, you can use either of the these commands to map a LUN to an initiator group:

- `lun map -igroup engineering -name database1 -lun-id 1`
- `lun mapping create -igroup engineering -name database1 -lun-id 1`

The LUN named database1 is mapped to an initiator group named engineering. For this mapping, the LUN ID is 1. After you enter this command, hosts in engineering can access database1, and Mars OS creates this LUN map:

```
{engineering, database1, 1}
```

Unmapping LUNs (CLI)

Steps

1. Enter one of the following commands:

- `lun unmap -igroup string -name string`
- `lun mapping delete -igroup string -name string`

For example, you can use either of the following commands to unmap a LUN:

- `lun unmap -igroup engineering -name database1`
- `lun mapping delete -igroup engineering -name database1`

The LUN named database1 is unmapped from an initiator group named engineering. The hosts in engineering can no longer access database1.

Mapping and unmapping LUNs (System Manager)

Steps

1. Go to the **Manager > STORAGE > All Storage > LUNs** window.
2. Click a LUN name.
3. Click the **Edit Properties** button. In the pull-down menu, select **Mappings**. The **Edit Mappings** dialog box is displayed.
4. In the **Edit Mappings** dialog box, follow these steps:
 - a. The Available section displays initiator groups that you can map to the LUN. To reduce the length of the list, you can enter a string under Initiator Groups to display only names containing the string.
 - b. In the Available section, click the initiator group name to select it. To select multiple initiator groups, click the checkbox next to each name. To select all initiator groups, click the checkbox next to Available.
 - c. Click  to move the initiator groups to the Selected section.
 - d. In the Selected section, you can optionally specify the LUN ID for the mapping.
 - e. If you want to unmap an initiator group, click the checkbox next to its name in the Selected section and click . If you want to unmap all initiator groups, click .
 - f. Click **OK**.

Changing the LUN ID for a LUN map (CLI)

Steps

1. Unmap the LUN.
2. Map the LUN by using the new LUN ID.

Changing the LUN ID for a LUN map (System Manager)

Steps

1. Go to the **Manager > STORAGE > All Storage > LUNs** window.
2. Click a LUN name.
3. Click the **Edit Properties** button. In the pull-down menu, select **Mappings**. The **Edit Mappings** dialog box is displayed.
4. In the Selected section, type the new LUN ID and click **OK**.

Showing LUN information (CLI)

Steps

1. To show information about LUNs, enter the following command:

```
lun show [ -names name,... ] [ -v ]
```

The `-names` parameter specifies the LUNs to be displayed. Without the parameter, all LUNs are displayed.

Without the `-v` option, the command displays the LUN name, size, state, status, serial number and mapped initiator groups. The `-v` option displays a more verbose output but omits the initiator group names. The following table describes the command output with the `-v` option.

Fields in the lun show -v command	Description
LUN Name	Name of the LUN specified when you created the LUN
Size	LUN size rounded to two decimal places for display only. The actual number of bytes in the LUN is displayed in parentheses.
Max Extent Size	Maximum size of an extent used by the LUN. It affects the performance of the LUN. Leave it at default unless you are instructed by technical support to modify it.

Fields in the lun show -v command	Description
Serial Number	Serial number of the LUN assigned by Mars OS when the LUN was created. It is 12 characters long and unique within the cluster. The characters can be letters from A through Z and a through z, digits from 0 through 9, forward slash (/), and hyphen (-). It can be changed only on an offline LUN. However, do not change it unless you are instructed by technical support to do so.
LUN UUID	Universally unique identifier for the LUN, which is a 32-character hexadecimal number.
State	It can be online or offline.
Status	It can be normal, nvfail, no-space, or corrupt.
Comment	A comment you can add to an existing LUN. Each LUN is created with an empty comment. You can add or modify a comment only with the CLI, not the System Manager.
Block size	It is always 512 bytes. It cannot be modified.
Creation time	Time when the LUN was created. It is displayed in UTC format.
Mapped	Whether the LUN is mapped or unmapped. To see the igroups mapped to the LUN, use the <code>lun show</code> command without the <code>-v</code> option or the <code>lun mapping show</code> command.

Showing LUN mappings (CLI)

Steps

1. To display information about how LUNs are mapped to initiator groups, enter the following command:

```
lun mapping show [ -igroups igroup-name,... ] [ -names lun-name,... ]
```

If you omit the parameters, all LUN maps on the cluster are displayed. Otherwise, only the LUN maps for the specified initiator groups or LUNs are displayed. The `-igroups` and `-names` parameters are mutually exclusive. The information displayed for each LUN map is the initiator group name, LUN name, and the LUN ID for the mapping.

Showing LUN information and mappings (System Manager)

Steps

1. Go to the **Manager > STORAGE > All Storage > LUNs** window.
2. See [LUNs window \(System Manager\) on page 53](#) for information about the contents of this window.

Increasing the LUN size (CLI)

Steps

1. To resize a LUN, enter the following command:

```
lun resize -name name -size size
```

Increasing the LUN size (System Manager)

Steps

1. Go to the **Manager > STORAGE > All Storage > LUNs** window.
2. Select a LUN.
3. Click **Edit Properties**. In the pull-down menu, select **Size**.
4. In the dialog box, enter the LUN size and then click **Close**.

Taking a LUN offline and bringing it online

All LUNs by default are online. If you do not want any host to access a particular LUN, you can take it offline without affecting the mapping relationship between the LUN and the initiator groups.

Taking a LUN offline (CLI)

Steps

1. Verify that no applications are using the LUN.
2. Enter the following command:

```
lun modify -name name -state offline
```

The following is an example:

```
lun modify -name database1 -state offline
```

The LUN named database1 is offline and inaccessible to any host.

Bringing a LUN online (CLI)

Steps

1. Enter the following command to bring a LUN online:

```
lun modify -name name -state online
```

The following is an example:

```
lun modify -name database1 -state online
```

The LUN named database1 is brought online and is accessible to initiator groups that are mapped to it.

Taking a LUN offline and bringing a LUN online (System Manager)

Steps

1. Verify that no applications are using the LUN.
2. Go to the **Manager > STORAGE > All Storage > LUNs** window.
3. Select a LUN.
4. Click **Edit Properties**.
5. In the pull-down menu, select **State**. Then select **Online** or **Offline**.

Renaming a LUN

You can rename a LUN at any time. Renaming a LUN does not affect the mapping between the LUN and initiator groups. Also, the renamed LUN has the same LUN ID as before.

Renaming a LUN (CLI)

Steps

1. Enter the following command:

```
lun rename -name string -new-name string
```

For example, the following command changes the name from lun1 to lun2:

```
lun rename -name lun1 -new-name lun2
```

Renaming a LUN (System Manager)

Steps

1. Go to the **Manager > STORAGE > All Storage > LUNs** window.
2. Click **Edit Properties**. In the pull-down menu, select **Name**.
3. In the dialog box, enter the new name and click **OK**.

Deleting a LUN (CLI)

Considerations for a forced deletion

- Mars OS checks to see if the LUN you try to delete is online and mapped. However, you can force a deletion by including the `-force` parameter in the `lun delete` command.
- If a LUN is not mapped to any initiator groups, you can delete the LUN without the `-force` parameter whether or not it is offline.

Steps

1. To delete a LUN that is mapped but offline or that is unmapped, enter the following command:

```
lun delete -name name
```

To force a deletion of a LUN that is mapped and online, enter the following command:

```
lun delete -name name -force
```

The following examples show how LUNs are deleted:

- `lun delete -name database1`
- `lun delete -name database2 -force`

Deleting a LUN (System Manager)

Steps

1. Go to the **Manager > STORAGE > All Storage > LUNs** window.
2. Verify that the LUN is not online and mapped. You cannot delete an online LUN that is currently mapped to an igroup.
3. Select the LUN to delete.
4. Click **Delete**.

A confirmation dialog box is displayed, warning that deleting a LUN permanently removes all data in the LUN. To confirm, click the checkbox in the dialog box and then click **OK**.

Modifying or creating a comment

You can write a comment for a LUN whether it is online or offline.

Steps

1. Enter the following command to modify or create a comment:

```
lun modify -name name -comment "string"
```

The string can be up to 255 characters. The following example creates a comment for the LUN named database1:

```
lun modify -name database1 -comment "Database created for testing purposes"
```

Clearing the nvfail state

Steps

1. Enter the following command to clear the nvfail state:

```
lun modify -name string -nvfail clear
```

This command clear the nvfail state of the specified LUN. After the nvfail state is cleared, the LUN is online.

Related topics

[Creating LUNs on page 42](#)

[Understanding LUNs on page 36](#)

[Understanding LUN mapping on page 40](#)

LUNs window (System Manager)

In the **Manager > STORAGE > All Storage > LUNs** window, you can provision storage space and configure the cluster so that network hosts can access the provisioned space.

It is important for you to read the following topics before using the LUNs window:

- [Understanding LUNs on page 36](#)
- [Understanding LUN mapping on page 40](#)
- [Understanding initiator groups \(igroups\) on page 56](#)

The **LUNs** window provides two tabs: **LUNs** and **Initiator Groups**. By default, the **LUNs** tab is selected, which enables you to perform the following tasks:

- Creating a LUN
- Editing LUN properties (such as name, size, mapping, and state)
- Deleting a LUN

Displaying LUNs

The **LUNs** window displays a table showing LUNs currently on the cluster. The following list describes different ways to display the table contents:

- Entering a string in the **Search** field displays the LUNs with names that contain the string. The search string is case-insensitive. For example, searching for ABC and searching for abc give you the same result.
- Clicking the column heading sorts the information. For example, if you click the **Name** heading, the information is sorted alphabetically by the LUN names.

The following table explains the information in the LUN table.

Column	Description
Name	It is the LUN name. You can click the box preceding the name to select the LUN. After you select it, you can choose to edit its properties or delete it.
Size	It is the LUN size. Because of rounding, the LUN size displayed might be different from the value you specified when you created or resized the

Column	Description
	LUN.
State	It shows both the LUN state and status. The following text is an example: Offline, no space In this example, the LUN state is offline and status is no-space. That is, the cluster runs out of space and as a result, Mars OS takes the LUN off line. If the LUN status is normal, only the LUN state is displayed in this column. See Understanding LUNs on page 36 for information about LUN state and status.
Serial number	It is a unique, 12-byte, ASCII string generated by the storage system. Many multipathing software packages use this serial number to identify redundant paths to the same LUN. You cannot change the serial number in the System Manager.
Mapped to	This column shows the names of the initiator groups to which the LUN is mapped. For each mapping, the LUN ID is displayed after the initiator group name. For more information about LUN IDs, see Understanding LUN mapping on page 40.

Related topics

[Creating initiator groups on page 58](#)

[Managing initiator groups on page 59](#)

[Understanding LUNs on page 36](#)

Igroup management

This chapter discusses the following topics:

- Understanding initiator groups (igroups) 56
- Creating initiator groups 58
- Managing initiator groups 59

Understanding initiator groups (igroups)

An initiator group is a group of software initiators or host bus adapters (HBAs) that function as SCSI initiators. In this guide, an HBA or a host initiator is referred to as initiator.

Through initiator groups, you control which initiators have access to specific LUNs. Access to a particular LUN is restricted only to initiator groups that are mapped to it.

How an initiator is represented in an initiator group

Each initiator is represented by a World Wide Port Name (WWPN). The Fibre Channel protocol (FCP) uses WWPNs as unique addresses to identify initiators. An initiator group consists of a list of WWPNs, and initiators with these WWPNs can access the LUNs mapped to the initiator group. For more information about FCP, see [Understanding Fibre Channel Protocol \(FCP\) on page 75](#).

Considerations for creating an initiator group

The following are general considerations for creating an initiator group:

- You can create an initiator group before or after you create a LUN. You can add initiators to or remove initiators from an initiator group at any time.
- All initiators in the same initiator group must run the same operating system (OS).
- An initiator can belong to multiple initiator groups, provided that these requirements are met:
 - An initiator cannot belong to multiple initiator groups that have different OS types. For example, initiator A cannot be in a Linux initiator group and a Windows initiator group at the same time.
 - An initiator cannot be in multiple initiator groups that are mapped to the same LUN. For example, you should not add initiator1 to igroup1 and igroup2, which are both mapped to the LUN named eng-database. For more information about mapping LUNs to initiator groups, see [Understanding LUN mapping on page 40](#).

The following table lists the limits related to initiator group and initiators.

Maximum number of ...	Number
Initiator groups	4,096
Initiators per initiator group	256

Maximum number of ...	Number
LUN maps per initiator group	250
LUN maps per initiator	250
Initiator groups a LUN can be mapped to	4,096

Initiator group naming conventions

The following list describes the naming conventions for initiator groups:

- An initiator group name consists of one to 96 characters.
- Legal characters are:
 - A through Z and a through z (case-sensitive)
 - 0 through 9
 - hyphen (-)
 - underscore (_)
 - braces ({})
 - period (.)
- Spaces are not allowed.

Related topics

[Creating initiator groups on page 58](#)

[Managing initiator groups on page 59](#)

Creating initiator groups

Before you follow the instructions in this section to create an initiator group, be sure to read [Understanding initiator groups \(igroups\) on page 56](#) to understand how initiator groups work.

Creating an initiator group (CLI)

Steps

1. Enter the following command:

```
igroup create -name string -initiators string,... -ostype os-type
```

The string for the `-name` parameter is the initiator group name. The string for the `-initiators` parameter is the WWPN of the initiator. The `-ostype` parameter specifies the operating system running on the initiator.

The following example creates an initiator group named `my_igroup` with two initiators that run the Linux operating system:

```
igroup create -name my_igroup -initiators  
00:11:22:33:44:55:66:77,77:66:55:44:33:22:11:00 -ostype linux
```

Creating an initiator group (System Manager)

Steps

1. Go to the **Manager > All Storage > LUNs** window.
2. Select the **Initiator Groups** tab.
3. Click **Create Initiator Group**.
4. In the dialog box, type the initiator group name and select the operating system type.
5. For each initiator in the initiator group, type the WWPN and click **Add**.
6. Click **OK**.

Related topics

[Understanding initiator groups \(igroups\) on page 56](#)

[Managing initiator groups on page 59](#)

Managing initiator groups

For each existing initiator group, you can perform the following tasks:

- Show
- Modify
- Rename
- Destroy

NOTE: The CLI uses the term "destroy," and the System Manager uses the term "delete." They refer to the same operation.

- Create LUNs for mapping to the initiator group

NOTE: If you use the System Manager, you can create and map LUNs while managing the initiator groups. If you use the CLI, you can map LUNs only when you manage the LUNs, not when you manage the initiator groups.

Displaying initiator group information (CLI)

Steps

1. Enter the following command:

```
igroup show [ -igroups string, ... ]
```

If you do not use the `-igroups` parameter to specify the initiator group names, information for all initiator groups is displayed.

Displaying initiator group information (System Manager)

Steps

1. Go to the **Manager > STORAGE > All Storage > LUNs** window.
2. Select the **Initiator Groups** tab. Each entry in the initiator groups table displays the initiator group name, WWPNs of the initiators in the initiator group, the OS type, the names of the LUNs mapped to the initiator group, and the LUN ID for each mapping.

Modifying an initiator group (CLI)

You can modify the following characteristics of an initiator group:

- membership (for example, by adding an initiator)
- operating system type (OS type)

Adding or removing initiators

Steps

1. Enter the following commands to add initiators to and remove initiators from an initiator group:

```
igroup add -initiators string, ... -name string  
igroup remove -initiators string, ... -name string
```

In these commands, the `-initiators` parameter specifies the WWPN of the initiator being added or removed, and the `-name` parameter specifies the initiator group name.

The following examples add and remove two initiators:

```
igroup add -initiators 00:11:22:33:44:55:66:77,77:66:55:44:33:22:11:00 -name  
my_igroup  
igroup remove -initiators 00:11:22:33:44:55:66:77,77:66:55:44:33:22:11:00 -  
name my_igroup
```

Changing the operating system type (OS type)

You can change the OS type for one initiator group or for multiple initiator groups at the same time. Modifying multiple initiator groups at the same time is necessary when an initiator belongs to multiple initiator groups. Consider the following example:

If an initiator belongs to both `igroup1` and `igroup2`, you cannot change only the OS type of `igroup1`. You must change the OS type for both `igroup1` and `igroup2` at the same time. This is because an initiator cannot simultaneously belong to multiple initiator groups of different OS types.

NOTE: Changing the OS type for multiple initiator groups at the same time is possible only on the CLI. You cannot accomplish this task by using the System Manager.

Steps

1. Enter the following command to change the OS type of an initiator group:

```
igroup modify -names string,... -ostype string
```

For example, enter the following command to change the OS type to Linux for `igroup1` and `igroup2`:

```
igroup modify -names igroup1,igroup2 -ostype linux
```

Modifying an initiator group (System Manager)

Steps

1. Go to the **Manager > STORAGE > All Storage > LUNs** window.
2. Click the **Initiator Groups** tab. The initiator groups are displayed.

Entering a string in the **Search** field displays the initiator groups with names that contain the string. The search string is case-insensitive. For example, searching for ABC and searching for abc give you the same result.

3. Select the initiator group to be modified.
4. Click **Edit Properties** to display a pull-down menu. The following table describes how to modify different attributes of the initiator group:

To change...	Follow these steps
Initiator group name	1. Choose Name in the pull-down menu. 2. In the Rename Initiator Group dialog box, type the new name. Then click OK.
Initiators	1. Choose Initiators in the pull-down menu. 2. In the Change initiators dialog box, click the cross to remove an initiator, or type the WWPN of an initiator and click Add. 3. After you finish removing and adding initiators, click Close.
OS type	1. Choose OS Type in the pull-down menu. 2. In the Change Initiator Group OS Type dialog box, choose an OS type in the pull-down menu. Click OK.

Renaming an initiator group (CLI)

Steps

1. Enter the following command:

```
igroup rename -new-name string -name string
```

Destroying an initiator group (CLI)

By default, you can destroy only those initiator groups not mapped to any LUNs. However, you can force the operation by using the `-force` parameter.

Steps

1. Enter the following command:

```
igroup destroy [ -force ] -name string
```

If the `-force` parameter is specified, the command destroys the igroup even though it is mapped.

If the `-force` parameter is omitted, the command does not destroy the igroup if it is mapped.

The following example destroys an initiator group named `my_igroup` even when it is mapped:

```
igroup destroy -force -name my_igroup
```

Deleting an initiator group (System Manager)

Steps

1. Go to the **Manager > STORAGE > All Storage > LUNs** window.
2. Click the **Initiator Groups** tab.
3. Verify that there are no LUNs mapped to the initiator group to be deleted.
4. Click the name of the initiator group to select it.
5. Click **Delete**.

Creating LUNs to map to an initiator group (System Manager)

Steps

1. Go to the **Manager > All Storage > LUNs** window.
2. Click the **Initiator Groups** tab.
3. Select the initiator group to which newly created LUNs will map.
4. Click **Create LUNs**.
5. Enter the values in the Create LUNs dialog box. For information about this dialog box, see [Creating LUNs on page 42](#). After you click **OK**, the newly created LUNs are mapped to the selected initiator group. The mappings use the automatically assigned LUN IDs.

Related topics

[Understanding initiator groups \(igroups\) on page 56](#)

[Creating initiator groups on page 58](#)

Space management

This chapter covers the following topics:

Different measures of storage space	65
Understanding data compression	67
Understanding deduplication	68
Understanding thin provisioning	69
Monitoring space usage	72

Different measures of storage space

Understanding how different measures of storage space are calculated helps you interpret capacity statistics and plan for space provisioning. The following table explains different measures of storage space.

Measure of storage space	Meaning	Example (on a cluster with 24 drives)
Purchased space	Total raw capacity, which is equivalent to the sum of the capacity numbers printed on the labels of the installed SSDs. While the drive labels use the decimal prefix for byte quantities (for example, GB), purchased space uses the binary prefix (for example, GiB). One GB is equal to 10^9 (1,000,000,000) bytes. One GiB is equal to 2^{30} (1,073,741,824) bytes.	If 480-GB is printed on each drive label, purchased space per drive is 447 GiB, and the total purchased space is equal to: $24 \times 447 \text{ GiB} = 10.477 \text{ TiB}$
Maximum storable space	Physical storage space that Mars OS can use for storing user data. It is smaller than the purchased space due to space reserved for system use.	After space is reserved for system use, maximum storable space is approximately 7.15 TiB.
Provisioned space	Maximum amount of space that can possibly be consumed by client applications based on your request. The provisioned space for each LUN is the LUN size specified when you create or resize the LUN.	Suppose you create five LUNs and specify the size of each as 4 TiB. The provisioned space is calculated as follows: $4 \times 5 \text{ TiB} = 20 \text{ TiB}$
Written space	Amount of space needed for write operations from the client application's point of view. It is the sum of user data and metadata. The amount of metadata varies with the type of workload.	Suppose an application tries to write 10 TiB of data to the LUNs, and there is 150 GiB of metadata for that write operation. Written space is approximately 10.15 TiB. That

Measure of storage space	Meaning	Example (on a cluster with 24 drives)
		is, without compression and deduplication, 10.15 TiB would be needed for that write operation.
Stored space	Amount of space actually consumed to store the data written by the client application. It is typically smaller than the written space because of compression and deduplication.	Because of compression and deduplication, the amount of space consumed by the 10.15 TiB of data being written by the client application is actually 5 TiB. Therefore, stored space is 5 TiB.
Estimated writable space	Estimated amount of space that clients can use to store data, based on maximum storable space and current space savings that result from compression and deduplication. It is equal to maximum storable space multiplied by space savings. To calculate space savings, divide written space by stored space.	Space savings is calculated as follows: $10.15 \text{ TiB} / 5 \text{ TiB} = 2.03$ Therefore, estimated writable space is calculated as follows: $7.2 \text{ TiB} \times 2.03 = 14.62 \text{ TiB}$

Related topics

[Understanding physical storage on page 89](#)

Understanding data compression

Data compression is a Mars OS feature that enables you to store more data in the same amount of space than without data compression. This feature is always enabled.

Effects of compression on write operations

Mars OS performs data compression as an inline operation. That is, as data is sent to the cluster, it is compressed in memory before being written to the SSDs. This reduces both the number of write operations and the amount of data written.

Effects of compression on read operations

When a read request comes in, Mars OS reads only the compressed extents that contain the requested data. (An extent is the unit of storage space allocation for data in Mars OS.) This optimizes the amount of I/O being used to service the request. Only the required compressed data is decompressed in memory.

Related topics

[Different measures of storage space on page 65](#)

Understanding deduplication

Mars OS uses inline deduplication to maximize storage efficiency and reduce wear on the SSDs. Before data is written to the storage system, Mars OS checks to see if any of the data extents to be stored are identical to others already stored on the system. (An extent is the unit of storage space allocation for data in Mars OS). If so, it does not write the duplicate data to the SSDs. This process of eliminating duplicate data reduces the amount of physical storage space required for applications.

To determine how much space Mars OS saves due to deduplication, see the **Dashboard** window of FlashRay System Manager.

Related topics

[Different measures of storage space on page 65](#)

Understanding thin provisioning

Mars OS does not provision a fixed amount of storage space to each LUN; instead, it allocates space dynamically. This method of allocating space is called thin provisioning. It enables Mars OS to present you with more logical storage than the underlying physical storage on the system.

Benefits of thin provisioning

The following list describes the benefits of thin provisioning.

Minimum wasted storage space

When you create a LUN, you typically specify a LUN size that can accommodate future growth, not the space needed at present. It is rare for a LUN to be immediately filled to capacity upon its creation. A portion of the requested space (that is, the LUN size) might remain unused for a long time. Without thin provisioning, a storage system would have to dedicate the requested amount of space to the LUN, causing the unused space to go wasted, maybe for a long time. With thin provisioning, only space required for the application to write to the LUN is given to the LUN. Space requested but not consumed is available for all other LUNs to use.

Ability to provision more space than physical space

Because applications seldom use up the requested space on all LUNs at a given time, it is possible for Mars OS to create LUNs whose combined size exceeds the amount of physical storage. You can provision more space for LUNs with the same number of drives than if thin provisioning was not in effect, which means a reduction of up front and operational costs.

Flexibility in LUN sizing

Because storage space is not bound to LUNs forever, you can easily expand a LUN if it needs more storage than its original size. Mars OS simply allocates to it the space unused by the other LUNs as its storage need increases. The ability to allocate shared free space eliminates the need for purchasing physical storage whenever data in a LUN outgrows the original LUN size. This results in more flexibility in LUN configuration, as well as lower costs associated with operating and managing additional hardware.

Example of a system with more provisioned space than purchased space

For simplicity, this example does not take into consideration the space savings due to compression and deduplication. It only illustrates how Mars OS can present more logical storage to applications than the physical storage through thin provisioning.

Suppose you purchase 10 TiB of storage space for the system. After taking into account space reserved for system use, the cluster has a maximum storable space of 7 TiB. If you create five LUNs of 2 TiB each, there is 2 TiB of space on each LUN to which the application can store data. Together these LUNs represent a combined provisioned space of 10 TiB.

Mars OS allows you to provision 10 TiB of space on this cluster with only 7 TiB of maximum storable space because it assumes that at any given time, there is less than 7 TiB of data written to the LUNs.

The following table shows how space is provisioned on this cluster.

LUN name	LUN size (provisioned space)	Written space	Purchased space	Maximum storable space	Free space for any LUN to use
lun1	2 TiB	1 TiB			
lun2	2 TiB	1 TiB			
lun3	2 TiB	1 TiB			
lun4	2 TiB	1 TiB			
lun5	2 TiB	1 TiB			
Total	10 TiB	5 TiB	10 TiB	7 TiB	2 TiB

If the LUNs consume 1 TiB each, the unused 2 TiB is free for any existing or newly created LUN to use.

Suppose the application for lun1 begins to consume space rapidly, you can resize lun1 from 2 TiB to 3 TiB without having to add any new SSDs. Even if the application for lun1 needs to fill the LUN to capacity, it can use the shared free space that is not consumed by the other LUNs.

Considerations when allocating space to LUNs

Because of space savings from compression and deduplication, the space that clients can use to store data (the estimated writable space) is greater than the maximum storage space. For the formula that determines the estimated writable space, see [Different measures of storage space on page 65](#).

Before provisioning storage space, go to the **Capacity** pane of the System Manager **Dashboard** to view the estimated writable space and thin provisioning leverage.

Thin provisioning leverage is the ratio of provisioned space to estimated writable space. For example, if you provision a total of 20 TiB for five LUNs of equal size and the estimated writable space is 18.8 TiB, the thin provisioning leverage is 20 to 18.8, which is approximately 1.06.

In this example, the leverage is low and the cluster has a low risk of running out of space. It happens only when each LUN is almost at capacity.

Consider another example in which you provision a total of 40 TiB for the same LUNs. If the estimated writable space remains at 18.8 TiB, the thin provisioning leverage is 40:18.8, which is approximately 2.12. In this case, when the LUNs are filled to approximately half capacity, the cluster runs out of space.

The more highly leveraged your cluster is, the more likely it will run out of space when an application tries to write to a LUN. It is important to monitor storage usage to ensure that your provisioning strategy efficiently uses the storage space and yet minimizes the risk of write failures due to overcommitting.

Related topics

[Different measures of storage space on page 65](#)

[Monitoring space usage on page 72](#)

Monitoring space usage

To monitor space usage on the FlashRay system, view the capacity statistics displayed in the System Manager.

NOTE: You can also view capacity statistics on the CLI by using the `stats show -object capacity` command. However, the CLI is not the recommended method because the System Manager offers the same information in a format that is easier to read.

What happens when the cluster is low on space

When the cluster is low on space, an event and console message are generated as follows:

```
MARS is in low space condition with number physical capacity
```

The number in the actual message is the amount of remaining physical space on the cluster that can be used for storing client application data.

Client applications can continue reading from and writing to the cluster, but it is recommended that you remove unnecessary LUNs to make more space available.

What happens when the cluster runs out of space

The following list describes what happens when the cluster does not have enough space to perform a write operation:

- The following event and console message are generated:

```
MARS has overflowed its physical storable capacity
```

- The LUN to which the client application sends the write request is taken offline, and the status of the LUN is changed to no-space. The following event is generated to reflect the change of the LUN state:

```
LUN serialnum taken offline due to lack of space in the system.
```

For information about events, see [Understanding events on page 103](#).

Reclaiming space

If you see an offline LUN due to a space shortage, make space available on the cluster before putting the LUN back online. You can reclaim space by deleting LUNs that you do not need. After you delete the LUNs, check the event log or the console to determine whether the amount of available space is at a safe level. When the cluster is no longer in danger of running out of space, the following event message is generated, which also appears on the console:

```
MARS is recovered from its earlier overflow space condition.
```

This message indicates that you can put the offline LUN back online.

For more information about how to delete a LUN and change the state of a LUN, see [Managing LUNs on page 45](#).

Displaying capacity information (System Manager)

Steps

1. Go to the **Dashboard**.
2. Make sure that the **Performance** pane is in the default view. You can see capacity statistics in the **Capacity** pane. For information about terms and statistics in this pane, see the FlashRay System Manager online help.

Related topics

[Different measures of storage space on page 65](#)

[Understanding thin provisioning on page 69](#)

[Understanding performance statistics on page 112](#)

Fibre Channel Protocol (FCP) management

This chapter discusses the following topics:

Understanding Fibre Channel Protocol (FCP)	75
Requirements and recommendations for FC switch configuration and zoning	78
Managing FCP (Fibre Channel Protocol) service	80
Configuring FC (Fibre Channel) adapters	81

Understanding Fibre Channel Protocol (FCP)

The FlashRay system is a component of the FC fabric, which is a switched network topology. FCP is the protocol used to transfer data between the FlashRay system and network hosts.

Components of an FC network include targets, initiators, and switches. FC adapters in the FlashRay system are targets, and hosts accessing the FlashRay system are initiators.

Each initiator or target registers with the Fabric Name Server when it is connected to an FC switch. Each is identified by a WWNN (World Wide Node Name) and a WWPN (World Wide Port Name).

Meaning of adapter and port

Mars OS uses the terms adapter and port interchangeably. Both terms refer to the CNA (converged network adapter) connector on the back of the controller to which you attach the FC cable.

The adapters are named 0e, 0f, 0g, and 0h. You cannot change an adapter name. You need these names when you display or change adapter settings.

About MPIO (multipath I/O)

It is recommended that you have a minimum of two paths for each host to reach a LUN to eliminate single points of failure. For example, you can connect two initiator ports to a switch, which in turn is connected to two FC ports on the FlashRay controller. In this way, the host can access the cluster by using any of the four paths.

For a host to access a LUN through more than one path, MPIO software is required on the host. The software presents a single disk to the host operating system for all paths to a LUN. Without it, the operating system could treat each path as a separate disk, which can lead to data corruption.

About ALUA (Asymmetric Logical Unit Access)

ALUA is a method used to determine the optimized and non-optimized paths between the initiator and the LUN. The optimized paths are the most direct route from your initiator to the LUN. The non-optimized paths provide additional routes to the LUN if the optimized paths are not available.

Mars OS supports ALUA, and it is enabled at all times. Because the current release supports only one node per cluster, all paths to the cluster are optimized.

FC network limits

The following table lists the FC network parameters and their limits.

Parameter	Limit
Paths from host to a LUN	Windows: 32 Linux: 2,048 devices (each path to a LUN is a device) ESX: 256 (local drives and CD-ROMs count against this value)
Number of initiator logins per FC port	256
Number of initiator logins per controller	1,024

FC switch login limits

Each FC interface logs into an FC switch port. Each switch port has a limit on the number of logins from target ports. The total number of logins from a target port on the FlashRay system equals the number of interfaces plus one login from the FC adapter with which the interfaces are associated. In this release, because there is one interface per adapter, the total number of logins from a target port is two. Be sure not to exceed the switch vendor's login limit.

For the meaning of FC interfaces, see [Understanding interfaces on page 84](#).

Hop counts supported

The maximum supported FC hop count between a host and storage system depends on switch supplier. The hop count is defined as the number of switches in the path between the initiator and target. Cisco also refers to this value as the diameter of the SAN fabric.

Both Brocade and Cisco switches support a hop count of 7.

Guidelines and recommendations for setting adapter speeds

The following list offers guidelines and recommendations for setting the speeds of the FC adapters on the FlashRay system to maximize efficiency:

- The FC adapters on the FlashRay system are 16-Gb/s. The supported speed settings are 16 Gb/s, 8 Gb/s, 4 Gb/s, and auto (automatic negotiation).
- Set the adapter speed to match the speed of the device to which it connects instead of using automatic negotiation, which might take longer to reconnect after an interruption.
- Set all adapters used by a given host to the same speed.

NOTE: You can change the adapter speed only when the adapter is down. Adapter configuration must be done through the CLI, not FlashRay System Manager.

Accessing the FlashRay system and non-NetApp systems from the same host

A host can access the FlashRay system and other vendors' storage systems as long as certain requirements are met. The requirements are described in the Interoperability Matrix Tool (IMT) on the NetApp support site.

Related topics

[Managing FCP \(Fibre Channel Protocol\) service on page 80](#)

[Configuring FC \(Fibre Channel\) adapters on page 81](#)

[Understanding interfaces on page 84](#)

Requirements and recommendations for FC switch configuration and zoning

To increase your FlashRay system's availability and efficiency, make sure you meet the requirements and follow the recommendations when configuring FC switches and implementing FC zoning.

Switch speed

A fixed link speed setting works best, especially for large fabrics, because it provides the best performance for fabric rebuilds. In large fabrics, fixed link speeds can create significant time savings.

Although automatic negotiation provides the greatest flexibility, it does not always perform as expected. Also, it adds time to the overall fabric-build sequence because the FC port has to automatically negotiate.

NPIV (N_Port ID virtualization)

Mars OS uses NPIV to present FC targets to a fabric. Make sure to use an NPIV-capable switch and to enable NPIV on it.

FC port type

It is recommended that you set the switch port topology to F (point-to-point).

FC zoning

Follow these recommendations when you implement zoning:

- Implement zoning if four or more hosts are connected to a SAN.
- Although World Wide Node Name zoning is possible with some switch vendors, use World Wide Port Name zoning.
- Limit the zone size while still maintaining manageability. Multiple zones can overlap to limit size. Ideally, a zone is defined for each host or host cluster.
- Use single-initiator zoning to eliminate crosstalk between initiator HBAs.

NOTE: If the zoning method requires you to specify the WWPN of the target port, you must specify the WWPN of the interface, not the WWNN of the physical port. For information about interfaces, see [Understanding interfaces on page 84](#).

World Wide Name-based zoning

Zoning based on World Wide Name (WWN) specifies the WWN of the members to be included within the zone. Some switch vendors allow you to specify either World Wide Node Names or World Wide Port Names. WWPN zoning is recommended because it provides flexibility—access is not determined by where the device is physically connected to the fabric. You can move a cable from one port to another without reconfiguring zones.

For more information on interfaces, see [Understanding interfaces on page 84](#).

Related topics

[Understanding Fibre Channel Protocol \(FCP\) on page 75](#)

Managing FCP (Fibre Channel Protocol) service

By default, FCP is always running and all FC targets are enabled. If you want to isolate the cluster from the FC fabric (for example, for troubleshooting), you can disable the FCP service on all targets on the cluster and prevent the discovery of the cluster.

You can display and change the FCP status, and display adapter information only with the CLI, not FlashRay System Manager.

Displaying and changing the FCP status

Steps

1. To display the FCP status, enter the following command:

```
fcv show
```

The Status column of the output is up (which means FCP service is started) or down (which means FCP service is stopped).

2. To stop FCP service, enter the following command:

```
fcv stop
```

3. To start FCP, enter the following command:

```
fcv start
```

Displaying FC adapter information

You can display information about all FC adapters or selected adapters.

Steps

1. Enter either of the following commands, which provide the same result:

```
fcv adapter show [ -names adapter_name,... ] [ -v ]
```

```
network fcv adapter show [ -names adapter_name,... ] [ -v ]
```

If the `-names` parameter is omitted, information for all adapters is displayed. Examples of *adapter_name* are 0e and 0f. The `-v` option displays command output in verbose mode.

Related topics

[Understanding Fibre Channel Protocol \(FCP\) on page 75](#)

Configuring FC (Fibre Channel) adapters

By default, all FC adapters are configured up with the speed set to auto. You can change the state or the speed of each FC adapter.

You can configure the FC adapter only with the CLI, not FlashRay System Manager.

Before you begin

Configuring an adapter requires that the interface associated with the adapter be offline. To display or change the state of an interface, see [Configuring and displaying interfaces on page 85](#).

Steps

1. If you want to configure the adapter up or down, enter either of following commands, which produce the same result:

```
fcpx adapter modify -name adapter_name -state [ up | down ]
network fcpx adapter modify -name adapter_name -state [ up | down ]
```

When the command output describes the state of an adapter or interface, the terms *up* and *online* have the same meaning, and the terms *down* and *offline* have the same meaning.

2. If you want to change the adapter speed, make sure that the adapter is configured down. Then enter either of the following commands, which produce the same result:

```
fcpx adapter modify -name adapter_name -speed [ 4g | 8g | 16g | auto ]
network fcpx adapter modify -name adapter_name -speed [ 4g | 8g | 16g | auto ]
```

Sample commands for changing the adapter speed

The following sequence of commands configures the interface and adapter down, changes the adapter speed, and then configures the interface and adapter up:

```
fcpx interface modify -name fc1 -status-admin down
fcpx adapter modify -name 0e -state down
fcpx adapter modify -name 0e -speed 4g
fcpx interface modify -name fc1 -status-admin up
fcpx adapter modify -name 0e -state up
```

The following command verifies that the new speed is in effect:

```
fcpx adapter show -names 0e -v
```

82 | Fibre Channel Protocol (FCP) management

```
Node: mf8k-56a_01
Adapter: 0e
Physical Protocol: fcp
Maximum Speed: 16
Driver Status: ONLINE
Driver Substatus: ADAPTER_UP
Port Address: 270000
Firmware Revision: 6.10.0
Hardware Revision: 2
Data Link Rate (Gbit): 8
Fabric Established: True
Connection Established: ptp
Mediatype: ptp
Configured Speed: 4G
Admin Status: up
Switch Port: mars-9148-sw1 : 1/19
```

NOTE: Do not confuse the `fcp adapter show` (or its equivalent, `network fcp adapter show`) command with the `network port show` command. The `network port show` command shows the IP information of the e0M port, not FC adapter information.

Related topics

[Understanding interfaces on page 84](#)

Interface management

This chapter discusses the following topics:

- Understanding interfaces 84
- Configuring and displaying interfaces 85

Understanding interfaces

An interface is a virtual network port associated with a physical adapter such as an FC adapter. By default, each FC adapter on the FlashRay system is configured with an interface.

Each interface has an assigned WWPN (World Wide Port Name), and all interfaces have the same WWNN (World Wide Node Name). If your FC zoning method requires the WWPN of the target port in the zone, use the WWPN of the interface.

Default interfaces

The names for the default interfaces created by Mars OS are fc1, fc2, fc3, and fc4, which correspond to adapters named 0e, 0f, 0g, and 0h.

All interfaces are created by Mars OS. You cannot create, destroy, or change the names of the interfaces.

When you need interface information

You need to display interface information and manage interfaces in these situations:

- When you set up FC zones, you need to specify the interface's WWPN (World Wide Port Name).
- When you want to determine which or how many initiators are logged in to the FC adapter.
- Before you configure an FC adapter down (for example, to modify the adapter speed), you must take the interface associated with the FC adapter offline.
- When you check connectivity between an interface and a Fibre Channel (FC) port.

NOTE: You can only display interface information in the System Manager. Use the CLI to manage interfaces.

Related topics

[Understanding Fibre Channel Protocol \(FCP\) on page 75](#)

[Configuring FC \(Fibre Channel\) adapters on page 81](#)

[Configuring and displaying interfaces on page 85](#)

Configuring and displaying interfaces

You can display interface information with either the CLI or FlashRay System Manager. However, you can configure interfaces only with the CLI.

Displaying information about interfaces (CLI)

Steps

1. Enter the following command to display interface information:

```
fcpl interface show [ -names interface_name,... ]
```

The `-names` parameter specifies the interface names such as `fc1` and `fc2`. If you omit the parameter, the command displays information about all interfaces. The following table explains the information in the command output.

Field in the command output	Meaning
Name	The interface name, such as <code>fc1</code> .
Adapter	The adapter name, such as <code>0e</code> .
Node	Name of the node in the cluster.
RTPID	Relative target port ID, which is used to identify different paths to a LUN from an initiator that uses multipath software.
Admin	Administrative status of the interface, which can be <code>up</code> (the interface is online) or <code>down</code> (the interface is offline).
Oper	Operational status of the interface, which is the status of the adapter associated with the interface. It can be <code>up</code> (the adapter is online), <code>down</code> (the adapter is offline), or <code>link-down</code> (network cable is not attached to the adapter). The interface can be configured up or down

Field in the command output	Meaning
	even when the network cable is not attached to the adapter. The adapter speed is 0 when the operational status is <code>link-down</code> .
Speed	Speed of the adapter.
WWNN	WWNN of the adapter.
WWPN	WWPN of the interface.

Displaying information about initiators

Steps

1. To see information about initiators currently accessing the cluster through the interfaces, enter the following command:

```
fcplib initiator show [ -names interface_name,... ]
```

If you omit interface names, the command displays information about initiators that access all interfaces.

Changing the administrative status of an interface

Steps

1. Enter the following command:

```
fcplib interface modify -name interface_name -admin-status up | down
```

The following example changes the administrative status of the `fc1` interface to down, which means that the interface is taken offline:

```
fcplib interface modify -name fc1 -status-admin down
```

The following example changes the administrative status of the `fc1` interface to up, which means that the interface is online:

```
fcplib interface modify -name fc1 -status-admin up
```

Displaying information about interfaces (System Manager)

Steps

1. Go to the **Manager** window.
2. Click **CLUSTER > Networks** in the navigation pane.

Related topics

[Understanding interfaces on page 84](#)

Physical storage management

This chapter discusses the following topics:

Understanding physical storage	89
Understanding RAID-DP protection	90
Checking RAID group status	93
Managing solid-state drives (SSDs)	96
Displaying disk shelf information	99

Understanding physical storage

Solid-state disks (SSDs) provide the basic units of physical storage for the FlashRay system.

Physical storage space is equivalent to the storage space you purchase. See [Different measures of storage space on page 65](#) for information about how to determine the amount of space you can use for storing data.

The FlashRay system uses all SSDs in the cluster to store user data; there are no spare drives or dedicated RAID parity drives.

NOTE: Both the CLI (command line interface) and FlashRay System Manager use the term "disk" to refer to an SSD. However, there are no spinning disks in the disk shelves or in the system chassis.

Related topics

[Managing solid-state drives \(SSDs\) on page 96](#)

Understanding RAID-DP protection

Mars OS implements RAID-DP (dual-parity RAID 6) to ensure data availability even when two drives fail simultaneously in a RAID group.

NOTE: While the cluster can continue to serve data when drives are missing in the RAID group, all drives must be operational before you can create the cluster through the Setup Wizard or the `cluster create` command. If drives are missing in the RAID group before a cluster is created, the FlashRay controller boots to maintenance mode.

Understanding the RAID group on the FlashRay system

Parity data on the cluster is distributed across all drives in the RAID group. No dedicated parity drives are used, but the space occupied by the parity data is equivalent to the capacity of two drives. Mars OS supports one RAID group per cluster.

Data is written across all SSDs in the RAID group. Under normal circumstances, the Mars OS file system has reserved space that can be used for data reconstruction after an SSD failure.

Understanding data reconstruction

When one out of 24 SSDs fails (or is removed) in the RAID group, Mars OS uses parity information on the surviving SSDs to reconstruct the data that is originally on the missing SSD. Data reconstructed in this process consumes the file system reserved space.

After data reconstruction is complete, the RAID group regains normal status. However, the RAID group operates without any file system reserved space, and as a result, it is not possible for another data reconstruction to take place if an additional SSD fails.

If you replace the failed SSD, Mars OS gradually rewrites data to all SSDs including the newly installed one. After all data is rewritten, the file system reserved space is reclaimed.

What different RAID group statuses mean

The status of a RAID group can be normal, degraded, and double degraded. Understanding how a RAID group transitions from one status to another helps you manage the SSDs. Both the `storage raid show` command and the **Manager > CLUSTER > Health** window in System Manager display the RAID group status and corrective action.

NOTE: The terms RAID group status and RAID group state have the same meaning in this guide, the CLI, and the System Manager.

Normal status

A RAID group is normal in the following scenarios:

- All SSDs are operational and in ready state. There is file system reserved space that can be used for data reconstruction.
- One SSD fails and data reconstruction starts on the remaining 23 SSDs. After reconstruction is complete, the RAID group functions normally. However, there is no file system reserved space that can be used for data reconstruction.

Degraded status

A RAID group is degraded in the following scenarios:

- One SSD fails. Data reconstruction is in progress on the remaining 23 SSDs.
- One SSD fails and the remaining 23 SSDs are in ready state. After data reconstruction is complete, a second SSD fails. Because no file system reserved space for data reconstruction is available, data reconstruction cannot take place.
- Two SSDs fail simultaneously. The RAID group's status becomes double degraded and data reconstruction starts. After reconstruction is complete, the status is changed to degraded.

Double degraded status

When a RAID group is double degraded, the cluster cannot function if another drive fails. It is critical that you replace the failed SSDs as soon as possible to reduce the risk of data loss. The following list describes how a RAID group becomes double degraded:

- One SSD fails and the remaining 23 SSDs are in ready state. After data reconstruction is complete, a second SSD fails. Now the status of the RAID group is degraded. If a third SSD fails when the RAID group is degraded, the RAID group's status becomes double degraded.
- One SSD fails and the remaining 23 SSDs are in ready state. After data reconstruction is complete, two SSDs fail simultaneously.
- Two SSDs fail simultaneously. The RAID group's status becomes double degraded and data reconstruction starts. After reconstruction is complete, the status is changed to degraded. If a third SSD fails at this time, the status becomes double degraded.

NOTE: In these scenarios, the cluster can continue to function with 21 SSDs because the three failed SSDs do not stop functioning at the same time. If three SSDs fail simultaneously, the cluster crashes. Therefore, do not remove three functioning SSDs from the disk shelf at the same time.

Why it is important to replace a failed drive

Even though client applications can continue using the FlashRay system with a degraded or double degraded RAID group, it is recommended that you replace a failed drive as soon as possible for these reasons:

- If the RAID group is currently double degraded, you risk losing accessibility to the entire cluster when an additional drive fails.
- Performance might be adversely affected if there are missing drives in the RAID group.
- When all drives are functioning normally, Mars OS uses the file system reserved space to reduce write amplification. Write amplification is a phenomenon associated with SSDs that means more data is actually written than the amount of data in the write request. When the file system reserved space is consumed for data reconstruction, Mars OS cannot reduce write amplification as efficiently as it would when all SSDs in the RAID group are in ready state. As a result, the SSDs wear out more quickly.

Related topics

[Checking RAID group status on page 93](#)

[Managing solid-state drives \(SSDs\) on page 96](#)

Checking RAID group status

Checking RAID group status (CLI)

Steps

1. Enter the `storage raid show` command to display the status of the RAID group and the corrective action.

If the RAID group is normal with no failed drives, the following output is displayed:

```
-----
Status   Corrective Action
-----
Normal   None, all disks functional.
-----
```

If the RAID group has one failed drive, the following output is displayed:

```
-----
Status   Corrective Action
-----
Degraded Replace the 1 failed disk!
-----
```

If the RAID group has two failed drives, the following output is displayed:

```
-----
Status           Corrective Action
-----
Double Degraded  Replace the 2 failed disks!
-----
```

When the RAID group is in degraded or double degraded mode, the Status column might show the progress of data reconstruction as follows:

```
Degraded, Reconstruction (10%), Estimated remaining time : 188 sec
```

The actual remaining time might be different from the estimated remaining time, depending on the current traffic load on your system and other factors.

If the column does not show the progress, data reconstruction is complete.

2. To determine which drive needs replacement, enter the following command and look for the drive with the faulted state:

```
storage disk show
```

Checking RAID group status (System Manager)

Steps

1. Go to the **Manager > CLUSTER > Health** window.

The following table describes information displayed in the **Manager > CLUSTER > Health** window.

NOTE: In the navigation pane of the **Manager** window, a badge appears next to **Health** if a cluster component is in critical condition. The number in the badge indicates the number of components whose health icon is a red cross (✖) in the **Manager > CLUSTER > Health** window. For example, **1** indicates that one component needs immediate attention.

Column	Description
! (Health)	Health of the component. This release of the System Manager shows only the RAID group health. The following list describes the possible icons in the column: • A green dot (●) indicates that the RAID group is healthy. All 24 SSDs are operational. • A yellow triangle (▲) indicates that the RAID group needs your attention or that the RAID group does not provide maximum performance. The RAID group status can be either normal or degraded. The following list describes the typical scenarios in which this icon is displayed: ◦ The RAID group is degraded because of a single-SSD failure. For the meaning of degraded status, see Understanding RAID-DP protection on page 90. ◦ Data reconstruction is complete after a single-SSD failure. The RAID group status is normal. No corrective action is needed. • A red cross (✖) indicates that the RAID group needs your immediate attention to prevent service disruption. The following list describes the typical scenarios in which this icon is displayed:

Column	Description
	◦ The status of the RAID group is double degraded and data reconstruction is in progress. For the meaning of double degraded status, see Understanding RAID-DP protection on page 90. ◦ Two SSDs have failed and data reconstruction is complete. The status of the RAID group is degraded.
Component	The FlashRay component for which health information is displayed. Future releases of the System Manager will show the health information about components other than RAID.
State	RAID group status, which can be normal, degraded, or double degraded.
Details	Detailed information about the component. For example, if data is being reconstructed in the RAID group, this column indicates that data reconstruction is in progress. It also shows the estimated remaining time before reconstruction is complete. The actual remaining time might be different from the estimated remaining time, depending on the current traffic load on your system and other factors.

Related topics

[Understanding RAID-DP protection on page 90](#)

[Managing solid-state drives \(SSDs\) on page 96](#)

Managing solid-state drives (SSDs)

You can display SSD information and manage SSDs with the CLI. FlashRay System Manager only enables you to display SSD information.

Understanding SSD names

The SSD name is in the following form:

<SAS port name>.<shelf ID>.<bay>.

Each disk shelf is attached to four SAS ports on the FlashRay controller, and each SSD is assigned a SAS port. The SAS port name consists of a number followed by a letter. The built-in SAS ports on the controller are 0a, 0b, 0c, and 0d. This part of the SSD name is dynamically assigned by Mars OS, and you cannot change it.

For example, if the SSD name is 0a.0.1, the SSD is installed in bay 1, which is the second slot from the left, on the shelf with ID 0. The SAS port slot assigned to the disk is 0a.

SAS port numbers, disk shelf IDs, and bay numbers start at 0.

Displaying SSD information (CLI)

Steps

1. Enter the following command to display information about SSDs on the storage system:

```
storage disk show
```

Displaying SSD information (System Manager)

Steps

1. Go to the **Manager** window.
2. In the Navigation pane, click **HARDWARE > Inventory**. The **Inventory** window displays SSD information by default.

The following table describes the information displayed in the **Manager > HARDWARE > Inventory > SSDs** window.

Column	Description
! (Health)	Health of the SSD. The following list describes the possible icons in

Column	Description
	the column: • A green dot () means that the SSD is healthy. • A yellow triangle () means that the SSD needs your attention (for example, it is powered off). • A red cross () means that the SSD is in critical condition and you should replace it immediately to prevent service disruption. This column reflects the worse of the State column and Wear column. For example, if an SSD has failed but its wear level is 0, the State column shows Faulted and the Wear column shows . Because a Faulted state is critical, the Health column shows .
Name	SSD name.
Shelf	Shelf ID, which can be the default shelf ID or the one you assign by using the disk shelf ID button on the disk shelf. See the disk shelf documentation for more information for changing the shelf ID.
Bay	Bay numbers on the disk shelf, which go from left to right.
State	SSD state. The following list explains the possible states: • Ready means the SSD is operational. • Faulted means the SSD failed and is not operational. • Powered off means that the SSD has been

Column	Description
	turned off. If an SSD is in this state, no information about this SSD is displayed.
Serial Number	SSD serial numbers, which are assigned and cannot be changed.
Firmware	SSD firmware version. Firmware can be updated only through the CLI (command line interface).
Size	Physical capacity of the SSD as shown on the label on the drive.
Wear	Wear level expressed as a percentage. The lower the percentage, the longer the SSD's remaining life is. This column uses icons (🟢, 🟡, and 🔴) in the same way as the Health column.

Related topics

[Displaying disk shelf information on page 99](#)

[Understanding physical storage on page 89](#)

Displaying disk shelf information

You can display general and fault information about a disk shelf.

Disk shelf name format

You can display information about a specified shelf. The shelf name is in the following format:

<SAS port name>.shelf<shelf ID>

The SAS port name is assigned by Mars OS. It is 0a when multiple paths are available between the disk shelf and the controller. If only one path is available, the SAS port name reflects the port used (for example, 0c). The shelf ID is the ID you set on the shelf. The following is an example of a shelf name for a shelf with ID 0:

0a.shelf0

Displaying shelf information (CLI)

Command	Description
<code>storage shelf show [- name <i>string</i>]</code>	Displays general information about the specified shelf or all shelves such as the firmware version and serial number.
<code>storage shelf fault show [-name <i>string</i>]</code>	Displays information about elements on the shelf such as SSDs and power supplies.

Displaying shelf information (System Manager)

Steps

1. Go to the **Manager > HARDWARE > Inventory** window.
2. Click **Shelves**.

The following table describes the information displayed in the **Manager > HARDWARE > Inventory > Shelves** window.

Column	Description
! (Health)	Shows the general health of the shelf. The following list describes the possible icons in this

Column	Description
	column: • A green dot () means that the disk shelf is healthy. • A yellow triangle () means that you should monitor the disk shelf because a component on it might need replacement soon. • A red cross () means that the disk shelf is not operational.
Name	The disk shelf name.
State	When the shelf is operational, the state is online.
Serial Number	Each shelf is assigned a serial number by NetApp.
Unique ID	Each shelf is assigned a unique ID (for example, 5:00a:098003:48c73c).
Module A FW	Firmware version on module A of the disk shelf.
Module B FW	Firmware version on module B of the disk shelf.
Temp	Whether the temperature is within the acceptable range.
Fans	Whether the fans operate normally.
PSU	Whether the power supply operates normally.
Voltage	Whether the voltage is within the acceptable range.

Related topics

[Managing solid-state drives \(SSDs\) on page 96](#)

[Understanding physical storage on page 89](#)

Monitoring the FlashRay system

This chapter discusses the following topics:

Understanding events	103
Understanding the health monitor	105
Displaying environmental sensor information	108
Managing AutoSupport	109
Understanding performance statistics	112

Understanding events

An event generated by Mars OS is an occurrence in the FlashRay system that might be of interest, whether or not it requires corrective action. Different components of the system can trigger events. For example, an NVRAM battery temperature rising beyond its limit and a firmware upgrade on an adapter are considered events.

While you can decide whether to view information about events, you cannot enable or disable events, or configure Mars OS to generate different types of event information.

Health monitor and events

The event log is a record of all the events that occur in the system. The health monitor gets information from events, interprets the information, and generates alert messages to notify you of potential problems. Not all events are interpreted by the health monitor in the current release. To see event information not interpreted by the health monitor, you can view the event log directly.

AutoSupport and events

Mars OS defines which events trigger AutoSupport messages. For example, a system reboot causes an AutoSupport message to be generated. You cannot define when AutoSupport messages are sent based on event types, severity levels, and so on.

Format of an event message

All event messages are saved to the event log. Event messages of critical severity level and above are also displayed on the console. Each message consists of the time stamp, cluster name, event name, severity, and description.

NOTE: The severity levels are pre-defined. You cannot change their definitions. The severity levels help technical support to categorize information about your FlashRay system. Viewing the severity level alone is not sufficient to determine whether you need to take corrective action.

Viewing the event log

Steps

1. To view the event log, enter the following command:

```
event log show -lines integer
```

The integer is the number of lines that you want to display in the command output. If you omit the number of lines, the most recent 10 events are displayed.

Creating an event mark in the event log

You can create an event mark in the event log. An event mark is a manually created event that serves as a bookmark in the event log. It enables you to quickly locate a particular place in the event log. For example, you can create an event mark after you start a performance test. You can view the event log later to determine whether there are events generated after the event mark that can be attributed to the performance test.

Steps

1. Enter the following command to create an event mark:

```
event mark -comment string
```

The string in the command is the description of the event. Enclose the string in quotation marks if it includes spaces. For example:

```
event mark -comment "Performance test started"
```

An event named `evmgr.event.mark` is created with the notice severity level. When you view the event log in the future, you can search for the string to locate this event.

Related topics

[Understanding the health monitor on page 105](#)

[Managing AutoSupport on page 109](#)

Understanding the health monitor

The health monitor is a Mars OS daemon that correlates and aggregates conditions across all cluster components. Based on these conditions, the health monitor generates messages to alert you to potential problems or to provide steps to correct a situation.

Using the health monitor is a more efficient way of solving and preventing system problems than viewing the event log directly. From the event log, you can learn that a particular event happened (for example, a power supply on the controller was turned off). However, viewing this event in isolation does not enable you to determine whether a corrective action is needed. Because the health monitor gathers input from all components in the system, it can give you information and suggestions with more clarity than the event log does.

Understanding alert policies

An alert policy is a rule that determines when an alert is created and deleted. You cannot modify the policy, but you can disable or enable a policy at the diag privilege level. The alert policy status is not persistent across reboots. If you disable a policy, it becomes re-enabled after your system reboots.

NOTE: Do not disable an alert policy unless you are instructed by technical support to do so. If you disable a policy, the alert normally generated based on that policy is no longer generated.

How alerts are generated

The health monitor gets input from events generated by various FlashRay subsystems. The health monitor analyzes the input to determine whether an alert requiring your attention has occurred, by taking into account the number of event occurrences, the timing, and other events. An alert might mean that a problem exists or is developing in the system.

The health monitor generates an alert to notify you of a condition. An event is created to reflect that the health monitor has generated an alert.

Information in an alert

The following table describes the information in an alert.

Field in an alert	Description
policy	Name of the policy.
managed class	Class of objects being monitored by the policy.

Field in an alert	Description
managed_id	Identifier of object being monitored by the policy.
description	Description of the problem that triggers the alert.
corrective_action	Action recommended by the health monitor for solving an existing problem or avoiding future problems.
severity	It shows how urgent the issue is. Alert severity levels are different from event severity levels.
creation_time	Time when the alert is created.
tags	Logical groupings for sorting.

NOTE: The policy, managed class, and managed id combine to uniquely identify an alert instance.

Types of alerts supported

The health monitor does not alert you to all possible system problems or incidents. The following list shows examples of the alerts:

- Controller problems such as incorrect CPU configuration and incorrect memory size.
- PCI slot problems such as incorrect installation of PCI devices in the system.
- NVRAM problems such as an excessive number of bad blocks.
- Service disk problems such as an excessive number of relocated sectors and a high SSD wear level.
- Core file notifications, which inform you that core dumps exist on the cluster.
- Environmental conditions such as a high temperature in the chassis.

How an alert is deleted

An alert persists until the underlying condition for the alert has cleared. For example, after you take the corrective action, the health monitor automatically deletes the alert.

You can delete an alert after the underlying problem has been solved, without waiting for the health monitor to do so after it polls the FlashRay subsystems. You can also refresh the list of alerts, which causes the health monitor to poll the FlashRay subsystems and update the alert list to reflect the current condition of the cluster.

Alerts are not persistent across reboots.

Displaying alerts

Steps

1. Enter the following command:

```
system health alert show
```

Deleting an alert

Before you begin

Gather information from the following fields in the `system health alert show` command output:

- `policy`
- `managed class`
- `managed id`

Values in these fields are necessary to uniquely identify an alert to be deleted.

Steps

1. Enter the following command to delete an alert:

```
system health alert delete -policy-id string -managed-entity-class string -  
managed-entity-id string
```

The `-policy-id` parameter is the policy. The `-managed-entity-class` parameter specifies the managed class. The `-managed-entity-id` parameter specifies the managed id.

If the string contains a space, enclose the string with quotation marks (for example, "PCI slots").

Refreshing the list of alerts

Steps

1. Enter the following command to refresh the list of alerts:

```
system health alert refresh
```

Displaying environmental sensor information

You can display information gathered by environmental sensors in the system. For example, you can display information from sensors for fans and power supplies.

Threshold-based and discrete sensors

The threshold-based sensors provide numerical values. The discrete sensors provide values that represent the states of the sensors. For example, the PSU1 Temp sensor is threshold-based, which shows the temperature readings; the PSU1 Status sensor is discrete, which shows the states of the sensor. Mars OS displays the following information for these sensors:

PSU1 Status	normal	ON
PSU1 Temp	normal	23 C

Displaying environmental sensors

Steps

1. Enter the following command:

```
system environment show
```

The output contains a large number of entries. It is recommended that you redirect the output to the `more` or `grep` command so that you can easily find the relevant information. For example, the following example displays only the information about fans:

```
system environment show | grep fan
```

Related topics

[Understanding the health monitor on page 105](#)

Managing AutoSupport

AutoSupport is a Mars OS feature that enables your FlashRay system to periodically send the system's health-related information to technical support and other specified destinations. Information in AutoSupport messages is used by technical support to help you solve problems or prevent future problems.

AutoSupport messages sent to NetApp

Transmission of AutoSupport messages is disabled by default. You can enable it when you create your cluster by using the FlashRay Setup Wizard. You can also enable it by using the CLI on an existing cluster.

AutoSupport information is collected whether or not the transmission of AutoSupport messages is enabled. The most recent 250 messages are stored on the cluster. The older messages are deleted.

If enabled, AutoSupport messages, which are compressed files, are sent by the HTTPS protocol to NetApp, the default destination. You cannot customize or remove the default destination. For example, you cannot change the protocol used for the transmission.

You cannot customize AutoSupport message contents sent to NetApp. Each message includes the event log, current system configuration, and performance data. For more information about events, see [Understanding events on page 103](#).

Enabling or disabling AutoSupport messages to NetApp

Steps

1. Enter the following command to see whether AutoSupport messages are sent to NetApp:

```
autosupport callhome show
```

2. Enter the following command to enable or disable AutoSupport messages to NetApp:

```
autosupport callhome enable | disable
```

Configuring another AutoSupport destination

You can create another destination for receiving AutoSupport messages. When you add a destination, you specify the URI to which the messages are sent and the protocol (HTTP or HTTPS) used for the transmission.

Displaying AutoSupport destination information

Steps

1. Enter the following command to display the current destinations of AutoSupport messages:

```
autosupport destination show
```

The command always displays the settings for sending AutoSupport messages to NetApp. It also displays the settings for other destinations if you have created other destinations.

Adding a destination

Steps

1. Enter the following command:

```
autosupport destination add -uri string -enabled [ true | false ]
```

All parameters except for the `-enabled` parameter are required. By default, AutoSupport messages are enabled for the destination being added. If you specify `false` for the `-enabled` parameter, AutoSupport messages are not sent to this destination even though it is added.

Removing a destination

Steps

1. Enter the following command to display the uuid of the destination to be removed:

```
autosupport destination show
```

2. Enter the following command to remove the destination:

```
autosupport destination remove -uuid uuid
```

Changing the settings for a particular destination

Steps

1. Enter the following command to display the uuid of the destination to be modified:

```
autosupport destination show
```

2. Enter the following command to change a destination's setting:

```
autosupport destination modify -uuid uuid -uri string -enabled [ true | false ]
```

Reverting AutoSupport to factory settings

You can revert AutoSupport to factory settings, for example, to remove all the manually added destinations. Resetting AutoSupport also disables the transmission of AutoSupport messages to NetApp if you have enabled it.

Steps

1. Enter the following command to revert to the default AutoSupport configuration:

```
autosupport destination factory-reset
```

Triggering AutoSupport messages manually

You can trigger an AutoSupport message manually to test if AutoSupport messages are sent properly to destinations for which AutoSupport is enabled.

Steps

1. Enter the following command:

```
autosupport invoke
```

Checking the statuses of AutoSupport messages

You can check the statuses of AutoSupport messages, for example, to verify whether AutoSupport destination is set up correctly or to see a history of messages sent.

Steps

1. Enter the following command:

```
autosupport history show
```

Related topics

[Understanding events on page 103](#)

Understanding performance statistics

Mars OS generates system statistics for use in monitoring cluster performance. Performance data is available both at the command line interface (CLI) and through the FlashRay System Manager Dashboard window.

The Dashboard window provides a graphical user interface that aggregates performance statistics and displays that data as real-time performance, capacity, and health information for the cluster. For more information, see the FlashRay System Manager online help.

The CLI enables you to view performance data for specific system components, or objects.

Objects, instances, and counters

Components of the FlashRay system are identified as objects, with some objects having multiple instances. Mars OS uses counters to track the performance of an object's instances, and of the object as a whole. Information in object counters comprises the performance data shown in the CLI output and System Manager Dashboard.

An object is any of the following:

- Physical entity, such as disks, shelves, nodes, and ports
- Logical entity, such as LUNs and clusters
- Protocol, such as Fibre Channel Protocol (FCP)

Each object has zero or more instances. For example, a LUN named lun1 is an instance of the LUN object, and each SSD in a shelf is an instance of the disk object.

A counter is a predefined performance metric that provides data about an object or an instance of an object. For example:

- Capacity of the SSDs
- Average latency for a LUN
- Percentage wear of an SSD

The following table shows the relationship between an object and its instances and counters. In this table, the `disk` object has two instances: `0b.0.7` and `0b.0.8`. The object's counters provide data about each of these instances. The table shows three of the object's counters: `capacity`, `total_reads`, and `total_writes`.

Object	Instances	Counters
disk	0b.0.7	capacity
		total_reads
		total_writes
	0b.0.8	capacity
		total_reads
		total_writes

Displaying objects, instances, and counters

To display performance data, Mars OS requires an object name. With an object name, you can display the counters and instances for that object.

NOTE: Counter values are not aggregated and apply individually to each object instance. If an object has no instances, counters apply to the object itself.

Steps

1. Enter the following command to display a list of objects and their descriptions:

```
stats list
```

2. Enter the following command to display all counters for an object:

```
stats list -object string
```

In this command, *string* is the object name. For example, to display a list of LUN counters along with counter descriptions, enter the following command:

```
stats list -object lun
```

3. Enter the following command to display all instances of an object:

```
stats list -instances string
```

In this command, *string* is the object name for which you want to list the instances. For example, to display names and descriptions for all LUNs in the cluster, enter the following command:

```
stats list -instances lun
```

If an object has no instances, none are displayed.

Understanding counter explanations

The `stats` command displays detailed counter data for each object.

The explanation for each counter includes the counter name, type, description, and scope. The type specifies the value type of the counter. The description provides a brief explanation of the counter. The scope indicates whether the counter value applies to a single instance, such as a disk or LUN, or globally to the FlashRay system.

The following table describes the counter types.

Counter type	Description
Average	Cumulative value divided by another base value
Calculated	Counter value calculated from other system data
Histogram	Groups of related counters
Rate	Counter value based on the amount of data over time
Raw	Counter value for which no processing is needed (the raw value is used)
String	Counter value is text
Delta	Difference between a previous counter value and the current value

Viewing counter explanations

Steps

1. To view the explanations for all counters for an object, enter the following command:

```
stats explain -object string
```

For example, to view the explanations for all disk counters, specify the disk object:

```
stats explain -object disk
```

2. To view a specific counter explanation, enter an object name followed by a specific counter name, or a comma separated list of counters. For example, to get the explanation of `total_reads` and `total_writes` for the disk object, enter this command:

```
stats explain -object disk -counters total_reads,total_writes
```

Viewing performance data

You can view counters only for a single object at a time, though you can have multiple instances or counters for that object. By default, the `stats show` command displays data for all instances of an object in the cluster.

Steps

1. To view the parameters available for displaying system statistics, enter the following command:

```
stats show ?
```

2. To view default performance data for an object, enter the name of the object that you want to track:

```
stats show -object string
```

For example, to display performance data for all LUNs in the cluster, enter the following command:

```
stats show -object lun
```

3. To view specific performance data, enter the object name, along with any additional parameter settings that you want to apply. For example, to view the number of read and write operations for the LUN named `lun01`, enter the following command:

```
stats show -object lun -instances lun01 -counters r_ops,w_ops
```

Viewing performance data for a time period

You can view more than one data sample at a time. You can collect more than one sample from the same object at the same time. You can also specify multiple counters in the same sample.

Steps

1. To view performance data for a sample period, enter the `stats start` command along with the specific object that you want to track. In the command, specify any scope, instance, or counter values that you want to view. For example, to start tracking read and write operations for the LUN named `lun01`, enter the following command:

```
stats start -object lun -instances lun01 -counters r_ops,w_ops
```

The command generates a unique identifier that identifies this sample. For example:

```
ID: 7bd9e6c0-05f4-11e4-ad65-00a09856d62e
```

2. To stop the sample period, enter the `stats stop` command followed by the sample ID:

```
stats stop -id 7bd9e6c0-05f4-11e4-ad65-00a09856d62e
```

The command displays information about read and write operations collected by Mars OS for lun01 during the sample period.

Related topics

[Methods for managing the cluster on page 22](#)

Accessing and monitoring the FlashRay system through the Service Processor (SP)

This chapter discusses the following topics:

Understanding the Service Processor (SP)	118
Configuring the SP	120
Managing the SP from Mars OS	123
Accessing the SP CLI	124
Accessing Mars OS through the SP	125
Using the SP CLI	127
Troubleshooting the storage system by using the SP	131

Understanding the Service Processor (SP)

The SP is a hardware component on the FlashRay controller that serves as a remote management device. It serves the following purposes:

- Monitoring environmental conditions to ensure system availability and prevent data corruption. For example:
 - The SP, in collaboration with Mars OS, monitors the temperature within the chassis and sets the system fan speed to the lowest level for the system electronics to function properly. This minimizes both system power consumption and noise.
 - The SP monitors (but does not control) temperatures, voltages, currents, and fans on the power supplies.
- Automatically restarting a hung controller
- Giving you access to the FlashRay system regardless of the state of Mars OS
- Managing and updating FRU (field replaceable units) firmware

The FlashRay system can operate only when the SP functions properly. You cannot use the system without the SP.

What happens when an environmental sensor detects an abnormal condition

The SP alerts you to abnormal conditions by logging major events and associated forensics for later retrieval and analysis.

If the condition causes the FlashRay system to shut down, you can access the SP and issue commands at the SP prompt to power on the system.

How the SP is powered

The SP enables you to diagnose, shut down, power cycle, or reboot the controller, regardless of the controller's state. This is possible because the SP is powered by a standby voltage, which is available as long as there is input power to at least one power supply on the controller.

Understanding the system event log (SEL) on the SP

The SP has a nonvolatile memory buffer that stores up to 4,000 events in a system event log (SEL) to help you diagnose issues. Use the SP's `events` command to view a list of messages.

The SEL contains the following data:

- Hardware events detected by the SP such as sensor status about power supplies, voltage, or other components
- Errors detected by the SP such as a communication error, a fan failure, or a memory or CPU error

Understanding the console log on the SP

Regardless of whether a user is logged in to the console, the SP stores console messages in the console log. The console log persists as long as the SP has power from either of the controller's power supplies. Because the SP operates with standby power, it remains available even when the controller is power cycled or turned off.

Related topics

[Configuring the SP on page 120](#)

[Managing the SP from Mars OS on page 123](#)

Configuring the SP

If you want to access the SP over the network instead of through the console port, you must first configure the network settings of the SP. You can configure the SP to use IPv4 or IPv6 addresses (or both).

You must use the Mars OS CLI, not FlashRay System Manager, to configure the SP.

Command for configuring the SP

The following command configures the network settings of the SP:

```
system service-processor network modify
```

Configuring IPv4 network settings for the SP

The following command configures the IPv4 network settings:

```
system service-processor network modify -ip-address string -gateway string -  
netmask string
```

The SP's IP address, gateway address, and netmask must be configured for the SP to use IPv4 connections. If some of these parameters are already configured, you only need to enter the one you want to configure or modify.

In the following example, the netmask and gateway have already been set, and the command is used to modify the IP address only:

```
system service-processor network modify -ip-address x.x.x.x
```

In this example, *x.x.x.x* is an IPv4 address.

If the gateway or netmask is not configured, this command fails. You must enter the command with all unconfigured parameters.

Enabling or disabling DHCP for IPv4

DHCP is enabled by default. If you do not specify an IP address for the SP, an address is assigned through DHCP. The following command enables or disables DHCP for IPv4:

```
system service-processor network modify -dhcp-enable true | false
```

Configuring IPv6 network settings for the SP

The following command configures the IPv6 network settings:

```
system service-processor network modify -ipv6-address string -ipv6-gateway  
string -ipv6-address-prefix-length string
```

The `-ipv6-address-prefix-length` parameter specifies the number of high-order bits in the IPv6 address that define the network.

The SP's IP address, gateway address, and address prefix length must be configured for the SP to use IPv6 connections. If some of these parameters are already configured, you only need to enter the one you want to configure or modify.

In the following example, the address prefix length and gateway are already in effect, and the command is used to modify the IP address only:

```
system service-processor network modify -ipv6-address x:x:x:x:x:x:x
```

In this example, `x:x:x:x:x:x:x` is an IPv6 address.

If the IPv6 gateway or IPv6 address prefix length is not configured, this command fails. You must re-enter the command with all unconfigured parameters.

Enabling or disabling router-assigned IPv6 addresses

The following command enables the SP to acquire a router-assigned IPv6 address:

```
system service-processor network modify -ipv6-router-address-enable true
```

If you want to use static addressing, enter the following command to prevent the SP from obtaining a router-assigned address:

```
system service-processor network modify -ipv6-router-address-enable false
```

Enabling IPv4 or IPv6

To enable the IP version used by the SP, simply configure the IP addresses for that version. The IP version is enabled automatically. For example, after you configure an IPv4 address, IPv4 is automatically enabled.

Disabling IPv4 or IPv6

NOTE: Disabling an IP version deletes all configuration for that IP version. For example, if you disable IPv6, the IPv6 addressing information such as the IPv6 address, IPv6 address prefix length, and so on, is deleted. If you want to re-enable IPv6, you must use the `system service-processor network modify` command to configure the parameters again.

The following commands disable the IP version used by the SP:

- `system service-processor network modify -ipv4-enable false`
- `system service-processor network modify -ipv6-enable false`

When you display the network configuration of the SP by using the `system service-processor network show` command, the addresses of the disabled IP version are shown as a dash.

122 | Accessing and monitoring the FlashRay system through the Service Processor (SP)

Related topics

[Managing the SP from Mars OS on page 123](#)

Managing the SP from Mars OS

In addition to configuring the SP's network settings, you can use other Mars OS commands to manage the SP. For information about configuring network settings, see [Configuring the SP on page 120](#).

Displaying SP configuration information

Steps

1. Enter the following command to display the uptime, IP addresses, MAC address, and firmware version:

```
system service-processor show [ -force ]
```

Enter the following command to display more detailed network information such as default gateway address and whether the link on the SP network interface is up:

```
system service-processor network show [ -force ]
```

The information displayed by these commands is cached information. It might not be up-to-date. To display the current information, use the `-force` parameter. Displaying current information takes slightly longer than displaying cached information.

The cache is refreshed every hour or when one of the following events takes place:

- when you enter the `system service-processor modify` command
- when you enter the `system service-processor network modify` command
- when you display SP information with the `-force` parameter

Rebooting the SP

You can reboot the SP using the `system service-processor` command.

Steps

1. Enter the following command:

```
system service-processor reboot
```

Related topics

[Configuring the SP on page 120](#)

[Using the SP CLI on page 127](#)

Accessing the SP CLI

You can access the SP CLI from the FlashRay serial port or over the network if you configure it according to instructions in [Configuring the SP on page 120](#). After you access the SP, you can use its CLI to issue commands.

Number of SP CLI sessions allowed

Only one SP CLI session can be established at any given time, regardless of whether the session is initiated from the console port or through SSH. For example, if an SP CLI session is active for another user, your attempt to connect to the SP results in these messages:

```
User admin has an active CLI session.
```

```
Would you like to disconnect that session, and start yours [y/n]?
```

If you type Y, the other user's session is terminated.

Accessing the SP from the console port and exiting the SP CLI

Steps

1. Attach a console as you normally would for entering commands on the FlashRay console port.
2. At the Mars OS login prompt, press Ctrl-G. The `SP>` prompt appears.
3. From the `SP>` prompt, you can press Ctrl-D, or use the `exit` command, to exit the SP CLI and return to the Mars OS prompt.

Accessing the SP using SSH

Steps

1. As in the case of managing Mars OS, you connect to the SP through the FlashRay system's management port, which is e0M. The address to use in the SSH connection request is the SP's IP address. For example, if the SP's IP address is 192.168.56.2, enter the following command:

```
ssh admin@192.168.56.2
```
2. When you are prompted for the password, enter the admin password.

Related topics

[Using the SP CLI on page 127](#)

Accessing Mars OS through the SP

The SP stays operational regardless of the state of the FlashRay system. When Mars OS is down, if you need to manage the system, you must do so by using the SP CLI.

After you invoke the Mars OS system console through the SP, you can enter any Mars OS commands on the CLI as you normally would.

Mars OS authentication and SP authentication

Mars OS and the SP use the same login credential. For the SP, the user name is admin, and the password is the admin password. If you use the `cluster destroy` command to destroy the cluster, the admin password is reset to the default (changeme) for both Mars OS and the SP.

Number of concurrent system console sessions started from the SP

Only one system console connection to Mars OS through the SP is allowed at a time. For example, if a user is currently logged in to the Mars OS system console through the SP, you can still start an SP CLI session, but your attempt to start a system console session through the SP results in the following messages:

```
User admin has an active console session.
```

```
Would you like to disconnect that session, and start yours [y/n]?
```

Typing Y terminates the other user's system console session.

Concurrent system console sessions

Although only one system console session through the SP is allowed, it is possible for two console sessions to exist at the same time. For example, when one user accesses the Mars OS CLI on the console port, another user is allowed to establish a system console session through the SP. Both users can access the Mars OS CLI simultaneously, and the input and output are mirrored across these sessions.

Steps

1. At the SP> prompt, enter the following command:

```
system console
```

The Mars OS prompt appears, allowing you to enter any Mars OS commands. At this point, you stop communicating with the SP CLI, which means that another user can establish an SP session to access the SP CLI.

Returning from the Mars OS to the SP CLI

Steps

1. From the Mars OS prompt, press Ctrl-D to return to the SP session.

If another user initiated an SP CLI session when you were using the Mars OS system console, a message is displayed to ask if you want to terminate that user's SP CLI session. This is because only one SP CLI session can be in use at any time.

Related topics

[Understanding the Service Processor \(SP\) on page 118](#)

Using the SP CLI

About SP command privilege levels

SP commands are organized by privilege levels, which are admin, advanced, and diagnostic. The admin commands are relevant to more common use cases. Contact technical support if you want to use commands other than those at the admin level.

Using the SP help

You can display a list of SP commands by entering `?` or `help` at the `SP>` prompt.

You can display information about a particular SP command by entering `?` or `help` followed by the command.

NOTE: The SP CLI is different from the Mars OS CLI in terms of command structure and general command syntax. Use the SP help to familiarize yourself with the CLI before entering the SP commands.

SP commands at the admin privilege level

The following table describes the SP commands at the admin level. Some commands might cause a disconnection from the SP or a dirty shutdown of the FlashRay system. These commands display a confirmation message before executing. The `-f` option in these commands suppresses the confirmation message.

NOTE: The `rsa` command is displayed in the `help` output but is not supported by the FlashRay system.

Command	Description
<code>date</code>	Show the date.
<code>events all</code>	Show all events in the system event log.
<code>events info</code>	Show information about the system event log.
<code>events newest <i>number</i></code>	Show the newest <i>number</i> system events.
<code>events oldest <i>number</i></code>	Show the oldest <i>number</i> system events.
<code>events search <i>string</i></code>	Show system events containing the search string.

Command	Description
	Searches are case sensitive.
<code>help [command]</code>	Show command help.
<code>priv set [admin advanced diag]</code>	Set privilege level for SP commands. Use the <code>-q</code> option to suppress warning messages.
<code>priv show</code>	Show current privilege level.
<code>quit</code>	Disconnect from the SP.
<code>sp devices-version</code>	Show the revision of programmable devices.
<code>sp log history show</code>	Show the SP log archives, their sizes, and creation times. Mars OS does not support access to SP logs.
<code>sp reboot</code>	Reboot the SP.
<code>sp status [-v] [-d]</code>	Show status information about the SP such as firmware version and network settings. Output is in verbose mode with the <code>-v</code> option or debug mode with the <code>-d</code> option.
<code>sp update URL [-f]</code>	Update SP firmware. URL must have fewer than 200 characters.
<code>sp uptime</code>	Show SP uptime information.
<code>sp verbose show</code>	Show whether verbose mode is enabled or disabled.
<code>sp verbose set 0 1</code>	Disable (0) or enable (1) verbose mode.
<code>sp identify led [number]</code>	Turn on the system attention LED for the specified number of seconds. This command fails if the LED is already on.
<code>system battery show</code>	Show the NVRAM battery status.

Command	Description
<code>system console</code>	Connect to the system console.
<code>system core</code>	Dump the system core and reset. This command causes a dirty shutdown.
<code>system fru list</code>	Show the FRU (field replaceable units) list.
<code>system fru show <i>fru_id</i></code>	Show specified FRU's information such as part number and serial number.
<code>system log [-a] [-p <i>page_size</i>]] [-l <i>log_number</i>]</code>	Show the system console log. Display all console logs with the <code>-a</code> option. List the number lines to display per page with the <code>-p</code> option. View specific logs with the <code>-l</code> option.
<code>system power cycle [-f]</code>	Power the controller off, then on.
<code>system power on</code>	Power the controller on.
<code>system power off [-f]</code>	Power the controller off.
<code>system power status</code>	Show whether the controller is powered on or off.
<code>system reset backup [-f]</code>	Reset the controller using the backup BIOS firmware.
<code>system reset current [-f]</code>	Reset the controller using the current BIOS firmware.
<code>system reset primary [-f]</code>	Reset the controller using the primary BIOS firmware.
<code>system sensors show</code>	Show all environmental sensors' statuses and thresholds.
<code>system sensors show normal critical</code>	Show statuses of environmental sensors in normal or critical condition.
<code>system sensors lastcritical</code>	Show history statuses of environmental sensors in critical condition.

Command	Description
<code>system sensors filter <i>string</i></code>	Show sensor information that contains the specified string.
<code>system sensors get <i>sensor_name</i></code>	Show information about the specified environmental sensor.
<code>version</code>	Show SP version.
<code>exit</code>	Disconnect from the SP.

NOTE: The `system` command for turning the power off when Mars OS is running causes a dirty shutdown.

Related topics

[Accessing the SP CLI on page 124](#)

Troubleshooting the storage system by using the SP

If the storage system becomes unresponsive, you can use the SP to troubleshoot the system. This section describes the commonly used troubleshooting procedures and SP commands.

Troubleshooting an environmental sensor

You can use the SP to determine whether an environmental sensor has reached an abnormal condition. There are two types of sensors: threshold-based sensors and discrete sensors. You can display information about both types of sensors at the SP prompt.

Steps

1. To display information about all sensors, enter one of the following commands:

```
system sensors
```

```
system sensors show
```

To display information about a specific sensor, enter the following command:

```
system sensors get sensor_name
```

Understanding threshold-based sensors

Threshold-based sensors take periodic readings of various system components. The SP compares the reading of a threshold-based sensor against its preset threshold limit that defines a component's acceptable operating conditions. Based on the reading, the SP displays the sensor state to help you monitor the condition of the component.

Examples of threshold-based sensors include sensors for the system temperature, voltage, current, and fan speeds.

Threshold-based sensors have the following thresholds:

- lower critical (LCR)
- lower noncritical (LNC)
- upper noncritical (UNC)
- upper critical (UCR)

A sensor reading between LNC and LCR or between UNC and UCR means that the component is showing signs of problem and a system failure might occur. Therefore, you should plan for component service soon.

A sensor reading below LCR or above UCR means that the component is malfunctioning and a system failure is about to occur. The component needs immediate attention.

Output of system sensors command for threshold-based sensors

The `Current` column in the output shows the reading of a sensor. When the reading crosses a threshold, the sensor reports a problem, indicating its severity in the `Status` column. For example, when a sensor exceeds a limit for a particular threshold, the sensor's status in the output changes from `ok` to either `nc` (noncritical) or `cr` (critical), and an event message is logged in the SEL (system event log).

Some threshold-based sensors do not list limits for all four thresholds. The missing thresholds show `na` as their limits for these sensors in the command output. This means there is no limit or severity concern for the given threshold, and the SP does not monitor the sensor for that threshold.

Understanding discrete sensors

Discrete sensors do not have thresholds. Examples of discrete sensors include sensors for the fan, power supply unit (PSU) fault, and system fault.

The information in the `Current` column in the `system sensors` command output is ignored by the SP.

To interpret the value in the `Status` column of the `system sensors` command output, enter the following command:

```
system sensors get sensor_name
```

NOTE: Only the `System_FW_Status` sensor is used for troubleshooting the FlashRay system. Other discrete sensors such as `System_Watchdog`, `PSU1-Input_Type`, and `PSU2-Input_Type` are not supported.

Troubleshooting with the System_FW_Status sensor

The `System_FW_Status` sensor is used for troubleshoot FlashRay firmware issues. The `Status` column of the `system sensors` command output for this sensor is in the form of `0xAA`, where `AA` is a value that indicates the status of the firmware. The following table explains the possible values in the `Status` column:

Possible values in the <code>Status</code> column	Description
00	System software has cleanly shut down.
01	Memory initialization in progress.

Possible values in the <i>Status</i> column	Description
03	Memory Initialization done.
05	User has entered Setup.
06	BIOS has reset SP.
1f	BIOS is starting up.
20	LOADER is running.
21	LOADER is programming the Primary BIOS firmware, DO not power down the system.
22	LOADER is programming the alternate/backup BIOS firmware, Do not power down the system.
2f	Storage OS is running i.e. Bootloader has transferred control to Storage OS.

Related topics

[Using the SP CLI on page 127](#)

Rebooting and shutting down

This chapter covers the following topics:

Rebooting the cluster	135
Shutting down the cluster	137
Understanding the boot loader	139
Understanding maintenance mode	140

Rebooting the cluster

Rebooting the cluster causes a clean shutdown followed by a cluster startup. The following list describes typical scenarios when you need to reboot the cluster:

- After a Mars OS upgrade, you must reboot the cluster so that it can use the new version.
- If the cluster is in operational mode and you need to clear system problems (for example, in response to alerts generated by the health monitor), you can reboot the cluster to maintenance mode to gain access to troubleshooting tools.

Understanding cluster startup

Because rebooting the cluster performs a shutdown first, which terminates SSH connections to the cluster, use the serial console to initiate the reboot.

After you reboot the cluster to operational mode, the cluster prompts you for the cluster login name and password. After you log in, the `mars>` prompt appears. However, this does not mean that the cluster has finished starting up. If you enter a Mars OS command, you might see the following message:

```
Mars OS is not ready to execute this command.
```

To see the progress of the cluster startup, enter the following command:

```
cluster startup show
```

The command displays a message similar to the following if the startup is not complete:

```
Status: startup is running, started at Wed Jul 23 19:27:35 PDT 2014, duration  
thus far is 36 sec
```

The command displays a message similar to the following if the startup is complete:

```
Status: startup completed successfully at Wed Jul 23 19:28:27 PDT 2014
```

Steps

1. Enter the following command:

```
cluster reboot [ -force ] [ -maintenance ]
```

If the `-force` parameter is specified, the command reboots the cluster immediately without displaying a confirmation message. If you omit the `-force` parameter, the command displays a confirmation message.

If the `-maintenance` parameter is specified, the command boots all nodes of the cluster to maintenance mode. If you omit the parameter, the command boots the cluster to Mars OS operational mode.

You are notified of the low space condition only if you use the CLI. Also, if you see a LUN with no-space status in the State column of the **Manager > STORAGE > All Storage > LUNs** window, you must use the CLI to check event or console messages when performing the procedure to reclaim space.

Related topics

[Shutting down the cluster on page 137](#)

[Modes of operation on page 21](#)

[Understanding maintenance mode on page 140](#)

Shutting down the cluster

Shutting down the cluster stops all network activities on the cluster. Client applications can no longer perform input and output operations on the cluster. In addition, all SSH connections to the management interface and all System Manager sessions are terminated.

Typically you shut down the cluster in these scenarios:

- When you replace or add a field replaceable unit (FRU) that is not hot swappable (for example, the NVRAM battery).
- When removing an SSD would cause the cluster to crash because the RAID group is double degraded. See [Understanding RAID-DP protection on page 90](#) for information about determining the status of the RAID group.

After the cluster is shut down, the FlashRay system displays the boot loader prompt. At this prompt, you can boot the system to operational mode by following instructions in [Understanding the boot loader on page 139](#).

Understanding clean and dirty shutdowns

The `shutdown` or `reboot` command causes Mars OS to flush data from memory to the SSDs. Shutdowns caused by these commands are considered clean.

A dirty shutdown is caused by one of the following events:

- Mars OS crashes.
- The cluster loses power when it is in operational mode because:
 - there is a power outage
 - you turn off the power supplies on the controller
 - you power off the controller through the Service Processor (SP)
- You use the `system reset` command at the SP prompt.

When the cluster reboots after a dirty shutdown, Mars OS processes the data in the NVRAM to bring the cluster to the state before the shutdown so that no user data is lost. The reboot process is longer than that after a clean shutdown, and occasionally it takes multiple reboots to put the cluster back in operational mode.

Steps

Because shutting down the cluster terminates SSH connections to it, use the serial console to initiate the shutdown.

1. Enter the following command to shut down the cluster:

```
cluster shutdown [ -power-off ] [ -force ]
```

If the `-force` parameter is specified, the cluster is shut down immediately without displaying a confirmation message. If you omit the `-force` parameter, the command displays a confirmation message.

During the shutdown, some hardware configuration information is displayed on the console, which is followed by these messages and prompt:

```
Initializing Devices.  
Default boot device name = hd0.0  
LOADER-A*>
```

If the `-power-off` parameter is specified, the command also powers off all nodes in the cluster after the shutdown.

2. To boot the FlashRay system to operational mode, follow the instructions in [Understanding the boot loader on page 139](#).

Related topics

[Modes of operation on page 21](#)

[Understanding maintenance mode on page 140](#)

[Understanding the boot loader on page 139](#)

Understanding the boot loader

Each FlashRay system is shipped with a boot loader stored in the flash memory on the controller. It provides the user interface to the FlashRay system's BIOS (basic I/O system). The boot loader prompt is as follows:

```
LOADER-A*>
```

The letter A in the prompt refers to the slot in which the controller is installed.

In this guide, "boot loader" is shortened to "loader."

The loader is not accessible through an SSH connection to the cluster management interface or through the System Manager. You can see the loader prompt only through the console.

Typically you see the loader prompt after you enter the `cluster shutdown` command (for example, to prepare for the installation of a field replaceable unit).

When you are at the loader prompt and you need to boot the FlashRay system to operational mode, enter the following command:

```
boot_mars
```

NOTE: Loader commands are different from Mars OS commands in operational or maintenance mode. Do not use commands other than `boot_mars` at the loader prompt unless you are instructed by technical support to do so.

Related topics

[Modes of operation on page 21](#)

Understanding maintenance mode

The following list describes the reasons for using maintenance mode:

- When a node fails to boot, putting it in maintenance mode enables you to execute commands for gathering information and troubleshooting problems.
- After the FlashRay system crashes, technical support might advise you to boot the system to maintenance mode and start the consistency checker.

NOTE: Run the consistency checker only when you are instructed by technical support to do so.

- If you forget the password for the admin account, you must use a command in maintenance mode to re-create the admin password as described in [Resetting the admin password \(CLI\) on page 31](#).

Some, but not all, operational mode commands are available in maintenance mode.

NOTE: With the exception of the `security adminlogin reset` command and some commands for showing status information, commands in maintenance mode are for troubleshooting the FlashRay system, which enables the system to continue booting after a system crash. Using these commands requires advanced knowledge about Mars OS. Contact technical support if you need to use maintenance mode commands to solve system problems.

Where to access maintenance mode

Maintenance mode is not accessible through an SSH connection to the cluster management port or through the System Manager.

You can obtain access to maintenance mode through the console either by attaching a terminal to the controller's console port or by using the Service Processor's `system console` command.

Entering maintenance mode

From the system console, you can use one of the following methods to enter maintenance mode.

When the system is booting

Steps

1. When you see the following message in the boot sequence, interrupt the booting by pressing Ctrl-C:

```
#####
^C to start into maintenance mode.
#####
```

When the system is in operational mode

Steps

1. Enter the `cluster reboot -maintenance` command.
2. Enter `y` at the following prompt:

```
Are you sure you want to reboot the cluster?(y/n)
```

Using the maintenance mode CLI

The CLI in maintenance mode works in the same way as in operational mode. Also, these commands have the same function and syntax in both maintenance mode and operational mode: `help`, `history`, `man`, and `set`.

Commands are organized into normal and diag privilege levels as in operational mode. For detailed information about commands in maintenance mode, see the man pages.

In maintenance mode, commands might take longer to execute than in operational mode. The following message is displayed when the command is being processed:

```
Command is in progress.
```

Exiting maintenance mode

After you finish using the commands in maintenance mode, you can return to operational mode by entering the `continue-boot` command.

Related topics

[Modes of operation on page 21](#)

[Shutting down the cluster on page 137](#)

Troubleshooting the FlashRay system

This chapter covers the following topics:

What happens when the FlashRay system crashes or fails to reboot	143
Troubleshooting service disk issues	146

What happens when the FlashRay system crashes or fails to reboot

When problems cause the FlashRay system to crash or when the system fails to reboot, these events take place:

- The FlashRay system creates a core dump. It is a file containing information about the state of the system and information that helps technical support diagnose the problem.
- In some situations, when system problems prevent Mars OS from rebooting, the system automatically enters maintenance mode. In other situations, the system tries to reboot until the Mars OS process can run successfully. To break the reboot loop, press Ctrl-C when the following message is displayed:

```
#####
^C to start into maintenance mode.
#####
```

This interrupts the boot sequence, and Mars OS enters maintenance mode.

Whether the system enters maintenance mode automatically or you use Ctrl-C to interrupt the boot sequence, call technical support for information about how to use maintenance mode commands to troubleshoot. After fixing the problem, you can boot the system to operational mode.

Listing and deleting core dumps

You can list and delete core dumps either in operational mode or maintenance mode. Delete core dumps only if you are told to do so by technical support.

Steps

1. Enter the following command to list core dumps:

```
system coredump show
```

2. If technical support determines that a core dump is no longer needed for problem resolution, you can delete it to free up space by using the following command:

```
system coredump delete -corename string
```

Enabling the marsdiag login for troubleshooting

Technical support might give you troubleshooting instructions that require you to enable the marsdiag login. For example, to be able to transfer a core dump from your system to NetApp, you must establish a `sftp` (SSH file transfer protocol) session with your FlashRay system by logging in as the marsdiag user.

You can enable the marsdiag login in one of two ways:

- at the diag privilege level in operational mode
- at the normal privilege level in maintenance mode

By default, marsdiag login is disabled. At each system reboot or shutdown, the login is disabled and the password is unset.

NOTE: The marsdiag login is used only for `sftp` connections. For managing the FlashRay system, you must log in as admin. For example, if you try to start an SSH connection by using the marsdiag login, this error message is displayed:

This service allows sftp connections only.

Steps

1. Enter the following command to change the privilege level if Mars OS is in operational mode:

```
set -privilege diag
```

2. Enter the following command to create the marsdiag password and enable the marsdiag login:

```
security marsdiag enable -password string
```

The password requirements are the same as those for the admin login.

Establishing an sftp connection with the cluster

After you enable the marsdiag login, in a new terminal window, start an `sftp` connection with the cluster by using marsdiag as the login name and the marsdiag password. After you log in, you can access files on the cluster according to instructions from technical support.

For example, you can enter the following command to establish an `sftp` connection with a cluster that has a management IP address 172.10.11.2:

```
sftp marsdiag@172.10.11.2
```

You can use a host name instead of an IP address if the address can be resolved by a name server.

Related topics

[Modes of operation on page 21](#)

[Shutting down the cluster on page 137](#)

Troubleshooting service disk issues

The service disk is a 480-GB SSD on a PCIe card installed in the FlashRay system. It stores vital system information, system logs, and core dumps. The data stored is persistent across reboots. Because the service disk must function properly for the FlashRay system to boot and operate, a daemon called `sdmon` continuously runs to monitor the health of the service disk.

Displaying information about the service disk

The `sdmon` daemon polls the service disk hourly to retrieve SMART (self-monitoring, analysis, and reporting technology) data from it. You can use the `storage service-disk smart show` command to display the SMART data, which is similar to the following example:

Service disk name: /dev/sda

Attribute	Value	Threshold	Status
Relocated Sector (%)	0	95	ok
Relocated Sectors	0	4986	ok
Power on hours (%)	1	100	ok
Power on hours	2234	500000	ok
Wear Level (%)	1	100	ok
Wear Level (Program/Erase Cycles)	76	18950	ok
Spare Blocks Used (%)	6	95	ok
Spare Blocks Used	591	4986	ok
Temperature (C)	26	100	ok
Sectors Written (%)	42	100	ok
Sectors Written	1795141340	N/A	ok

Health alerts for service disk problems

If the polling result shows that a SMART threshold is exceeded, the daemon generates an event. Some, but not all, of these events cause the health monitor to display alert messages to notify you that the disk needs replacement. For replacement instructions, see the *FlashRay Controller Installation and Service Guide*. The following list shows the problem descriptions and corrective actions in the alert messages:

- The service disk's relocated sector count has exceeded the recommended threshold. Replace the service disk now.

- The service disk's wear-leveling has exceeded the recommended threshold. Replace the service disk now.
- The service disk's power-on hours has exceeded the recommended threshold. Replace the service disk now.
- The service disk's temperature has exceeded the recommended threshold. Replace the service disk now.
- The service disk's used-reserved-block count has exceeded the recommended threshold. Replace the service disk now.
- The service disk's sectors-written count has exceeded the recommended threshold. Replace the service disk now.

NOTE: The SMART thresholds are pre-set by the SSD vendor. They cannot be modified in the field.

In addition, the health monitor generates alerts to notify you of high space utilization on the service disk. The following list shows the problem descriptions and corrective actions in the alert messages:

- The service disk's crash storage is not mounted. Contact NetApp technical support.
- The service disk's crash storage has x% remaining. Remove core files from the crash directory.
- The service disk's crash storage has x% remaining. Further crashes may not dump core. Remove core files from the crash directory.

In the actual messages, x% is the percentage of space remaining on your service disk.

Related topics

[Understanding events on page 103](#)

[Understanding the health monitor on page 105](#)

Destroying the cluster

This chapter discusses the following topics:

What destroying the cluster means	149
Destroying the cluster	150

What destroying the cluster means

You can destroy your cluster by using the CLI. Destroying the cluster has the following effects:

- All user data and configuration information except the service processor (SP) configuration are destroyed. The system is reverted to the factory configuration and the admin password is reset to the factory default.
- The cluster is shut down and the controller displays the boot loader prompt.

Destroying the cluster does not:

- Zero the SSDs
- Remove the current version of Mars OS from the system.
- Delete system logs and core dumps.
- Revert the SP configuration to factory settings. The previously configured network settings of the SP remain in effect.

NOTE: When you run the FlashRay Setup Wizard to create a cluster again, leaving the field for the SP's IP address blank does not cause DHCP to assign a new address to the SP. The previously assigned address continues to be in effect until you change it.

Related topics

[Creating the cluster on page 18](#)

[Understanding the boot loader on page 139](#)

Destroying the cluster

Because destroying the cluster terminates all SSH connections to it, use the CLI on the serial console to destroy the cluster. In this way, you remain connected to the controller after the operation.

NOTE: Be sure you understand the effects of destroying the cluster described in [What destroying the cluster means on page 149](#) before you start the following procedure. The effects are not reversible, and you cannot recover the data lost due to the procedure.

Steps

1. On the serial console, enter the following command:

```
cluster destroy [ -force ]
```

If the `-force` parameter is specified, the command destroys the cluster immediately without displaying a confirmation message. If you omit the parameter, the command displays a confirmation message.

Related topics

[Creating the cluster on page 18](#)

[Understanding the boot loader on page 139](#)

Copyright information

Copyright © 1994–2015 NetApp, Inc. All rights reserved. Printed in the U.S.

No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark information

NetApp, the NetApp logo, Network Appliance, the Network Appliance logo, Akorri, ApplianceWatch, ASUP, AutoSupport, BalancePoint, BalancePoint Predictor, Bycast, Campaign Express, ComplianceClock, Cryptainer, CryptoShred, CyberSnap, Data Center Fitness, Data ONTAP, DataFabric, DataFort, Decru, Decru DataFort, DenseStak, Engenio, Engenio logo, E-Stack, ExpressPod, FAServer, FastStak, FilerView, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexSuite, FlexVol, FPolicy, GetSuccessful, gFiler, Go further, faster, Imagine Virtually Anything, Lifetime Key Management, LockVault, Mars, Manage ONTAP, MetroCluster, MultiStore, NearStore, NetCache, NOW (NetApp on the Web), Onaro, OnCommand, ONTAPI, OpenKey, PerformanceStak, RAID-DP, ReplicatorX, SANscreen, SANshare, SANtricity, SecureAdmin, SecureShare, Select, Service Builder, Shadow Tape, Simplicity, Simulate ONTAP, SnapCopy, Snap Creator, SnapDirector, SnapDrive, SnapFilter, SnapIntegrator, SnapLock, SnapManager, SnapMigrator, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapSuite, SnapValidator, SnapVault, StorageGRID, StoreVault, the StoreVault logo, SyncMirror, Tech OnTap, The evolution of storage, Topio, VelocityStak, vFiler, VFM, Virtual File Manager, VPolicy, WAFL, Web Filer, and XBB are trademarks or registered trademarks of NetApp, Inc. in the United States, other countries, or both.

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. A complete and current list of other IBM trademarks is available on the web at www.ibm.com/legal/copytrade.shtml.

Apple is a registered trademark and QuickTime is a trademark of Apple, Inc. in the United States and/or other countries. Microsoft is a registered trademark and Windows Media is a trademark of Microsoft Corporation in the United States and/or other countries. RealAudio, RealNetworks, RealPlayer, RealSystem, RealText, and RealVideo are registered trademarks and RealMedia, RealProxy, and SureStream are trademarks of RealNetworks, Inc. in the United States and/or other countries.

All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

NetApp, Inc. is a licensee of the CompactFlash and CF Logo trademarks.

NetApp, Inc. NetCache is certified RealSystem compatible.

How to send your comments

You can help us to improve the quality of our documentation by sending us your feedback. Your feedback is important in helping us to provide the most accurate and high-quality information.

If you have suggestions for improving this document, send us your comments by email to doccomments@netapp.com. To help us direct your comments to the correct division, include the product name, version, and operating system in the subject line.

You can also contact us in the following ways:

- NetApp, Inc., 495 East Java Drive, Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277