

LSI MegaRAID SAS Controller Tasks

This chapter describes the LSI WebBIOS Configuration Utility (CU), which operates independently of the operating system and enables you to create and manage RAID configurations on LSI MegaRAID SAS controllers.

If your server has an integrated RAID SAS controller, see [Chapter 4, “Configuring the LSI SAS2 Integrated RAID Controller.”](#)

This chapter contains the following sections:

- [LSI MegaRAID Controller Management Utilities, page 5-1](#)
- [LSI WebBIOS CU, page 5-2](#)
- [Managing RAID, page 5-16](#)
- [Determining Which Controller is in Your Server, page 5-22](#)
- [Limitation on Importing Foreign Configuration To a Virtual Disk That is Under Construction, page 5-25](#)

LSI MegaRAID Controller Management Utilities

LSI offers these three main MegaRAID controller management utilities:

- LSI WebBIOS Configuration Utility
- MegaRAID Command Tool
- MegaRAID Storage Manager Software

The LSI utilities have help documentation for more information about using the utilities.

For basic information on RAID and how to use the LSI utilities, see the following documentation:

[Broadcom 12Gb/s MegaRAID SAS Software User Guide](#)

LSI WebBIOS Configuration Utility

The LSI WebBIOS Configuration Utility (CU) operates independently of the operating system and enables you to create and manage RAID configurations on LSI MegaRAID SAS controllers.

MegaRAID Command Tool

The LSI MegaRAID Command Tool (CT) is a command-line interface (CLI) utility that enables you to configure, monitor, and maintain LSI MegaRAID SAS controllers and the devices connected to them.

MegaRAID Storage Manager

The LSI MegaRAID Storage Manager (MSM) is a software application with a graphical user interface that enables you to configure, monitor, and maintain storage configurations on LSI MegaRAID SAS controllers.

LSI WebBIOS CU

The LSI WebBIOS CU, unlike the LSI MegaRAID Storage Manager software, resides in the SAS controller BIOS and operates independently of the operating system.

You can use the WebBIOS CU to do the following tasks:

- Create drive groups and virtual drives for storage configurations.
- Display controller, virtual drive, physical drive, and battery backup unit (BBU) properties, and change parameters.
- Delete virtual drives.
- Migrate a storage configuration to a different RAID level.
- Detect configuration mismatches.
- Import a foreign configuration.
- Scan devices connected to the controller.
- Initialize virtual drives.

The WebBIOS CU provides a configuration wizard to guide you through the configuration of virtual drives and drive groups.

Starting the WebBIOS CU

Step 1 When the host computer is booting, the following text appears:

```
Copyright© LSI Corporation  
Press <Ctrl><H> for WebBIOS
```

Press **Ctrl+H**.

The **Controller Selection** window appears.

Step 2 If the system has multiple SAS controllers, choose a controller.

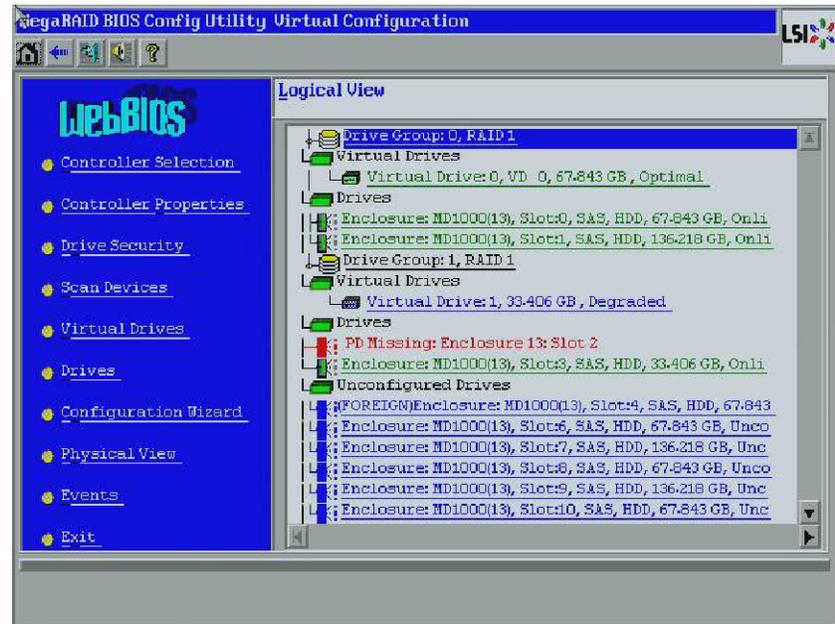
Step 3 Click **Start** to continue.

The main **WebBIOS CU** window appears. See [Figure 5-1](#).

WebBIOS CU Main Menu Window Options

Figure 5-1 shows the window that appears when you start the WebBIOS CU and choose a controller.

Figure 5-1 Main WebBIOS CU Window



Listed in the right pane are the following:

- Virtual drives configured on the controller
- Drives that are connected to the controller
- Drives that are foreign or missing



Note

In the list of virtual drives, the drive nodes are sorted based on the order in which you added the drives to the drive group, rather than the physical slot order that displays in the physical trees.



Note

The minimum screen resolution for WebBIOS is 640x480.

Click **Physical View** or **Logical View** in the left pane to toggle between the physical view and logical view of the storage devices connected to the controller. When Physical View appears in the right pane, it shows the drive groups that are configured on this controller.

For drives in an enclosure, the window shows the following drive information:

- Enclosure
- Slot
- Interface type (such as SAS or SATA)
- Drive type (HDD or SSD)
- Drive size

- Drive status (such as Online or Unconfigured Good)

Toolbar

The toolbar at the top of the WebBIOS CU contains the following tools, from left to right (see [Figure 3-1](#)):

- Main window—Click to return to the main window from any other WebBIOS CU window.
- Previous—Click to return to the previous window that you were viewing.
- Exit—Click to exit the WebBIOS CU program.
- Turn off alarm—Click to turn off the sound from the onboard controller alarm.
- About—Click to display information about the WebBIOS CU version, browser version, and HTML interface engine.

Menu Options

These options are located in the left pane of the main WebBIOS CU window (the hotkey shortcut for each option is shown in parentheses next to the option name):

- Controller Selection (Alt-c)—Views the Controller Selection window, where you can choose a different SAS controller. You can then view information about the controller and the devices connected to it, or create a new configuration on the controller.
- Controller Properties (Alt-p)—Views the properties of the currently chosen SAS controller. For more information, see [Table 5-1](#).
- Scan Devices (Alt-s)—Allows the WebBIOS CU to rescan the physical and virtual drives for any changes in the drive status or the physical configuration. The WebBIOS CU displays the results of the scan in the physical and virtual drive descriptions.
- Virtual Drives (Alt-v)—Views the Virtual Drives window, where you can change and view virtual drive properties, delete virtual drives, initialize drives, and perform other tasks. For more information, see [Viewing Virtual Drive Properties, Policies, and Operations, page 5-13](#).
- Drives (Alt-d)—Views the Drives window, where you can view drive properties, create hot spares, and perform other tasks. For more information, see [Viewing Physical Drive Properties and Operations, page 5-14](#).
- Configuration Wizard (Alt-o)—Starts the Configuration Wizard and creates a new storage configuration, clear a configuration, or add a configuration. For more information, see [Configuring RAID Drive Groups and Virtual Drives, page 5-5](#).
- Logical View (Alt-l)—Toggles between the Physical View window and the Logical View window.
- Physical View (Alt-h)—Toggles between the Physical View window and the Logical View window.
- Events (Alt-e)—Views system events in the Event Information window. For more information, see [Managing RAID, page 5-16](#).
- Exit (Alt+x)—Exits the WebBIOS CU and continue with system boot.

Configuring RAID Drive Groups and Virtual Drives

This section describes how to use the WebBIOS CU Configuration Wizard to configure RAID drive groups and virtual drives to create storage configurations.

Choosing the Configuration with the Configuration Wizard

Step 1 Click **Configuration Wizard** in the left pane of the WebBIOS main window.

The first **Configuration Wizard** window opens.

Step 2 Choose a configuration type:



Caution

If you choose **Clear Configuration** or **New Configuration**, all existing data in the configuration is deleted. Back up data that you want to keep before you choose a configuration type.

- **Clear Configuration**—Clears the existing configuration.
- **New Configuration**—Clears the existing configuration and allows you to create a new configuration.
- **Add Configuration**—Retains the existing storage configuration and adds new drives to it (this option does not cause any data loss).

Step 3 Click **Next**.

The **WebBIOS Configuration Method** window opens.

Step 4 Choose a configuration method:

- **Manual Configuration**—Allows you to control all attributes of the new storage configuration as you create drive groups and virtual drives, and set their parameters.
- **Automatic Configuration**—Automatically creates an optimal RAID configuration.

If you choose Automatic Configuration, you can choose whether to create a redundant RAID drive group or a nonredundant RAID 0 drive group. Choose one of the following from the **Redundancy** drop-down list:

Redundancy when possible

No redundancy

Step 5 Click **Next** to continue.

- For Automatic Configuration, continue with [Using Automatic Configuration, page 5-5](#).
 - For Manual Configuration, continue with [Using Manual Configuration, page 5-6](#).
-

Using Automatic Configuration

Step 1 When WebBIOS displays the proposed new configuration, review the information in the window, and click **Accept** or click **Back** to change the configuration.

One of the following configuration appears:

- **RAID 0**—If you choose **Automatic Configuration** and **No Redundancy**, WebBIOS creates a RAID 0 configuration.
- **RAID 1**—If you choose **Automatic Configuration** and **Redundancy when possible**, and only two drives are available, WebBIOS creates a RAID 1 configuration.
- **RAID 5**—If you choose **Automatic Configuration** and **Redundancy when possible**, and three or more drives are available, WebBIOS creates a RAID 5 configuration.
- **RAID 6**—If you choose **Automatic Configuration** and **Redundancy when possible**, and the RAID 6 option is enabled, and three or more drives are available, WebBIOS creates a RAID 6 configuration.

Step 2 When you are prompted to save the configuration, click **Yes**.

Step 3 When you are prompted to initialize the new virtual drives, click **Yes**.

WebBIOS CU begins a background initialization of the virtual drives.

New RAID 5 virtual drives and new RAID 6 virtual drives require a minimum number of drives for a background initialization to start. If there are fewer drives, the background initialization does not start. The following number of drives is required:

- New RAID 5 virtual drives must have at least five drives for a background initialization to start.
- New RAID 6 virtual drives must have at least seven drives for a background initialization to start.

Using Manual Configuration

This section contains the procedures for creating RAID drive groups for RAID levels 0, 1, 5, 6, 00, 10, 50, and 60.

Using Manual Configuration (RAID 0, 1, 5, 6)

For more information about RAID levels, see [RAID Levels, page 1-8](#).



Note

RAID level 0 is not fault tolerant. If a drive in a RAID 0 drive group fails, the whole virtual drive (all drives associated with the virtual drive) fails.

Step 1 Choose **Manual Configuration** and click **Next**.

The **Drive Group Definition** window appears. Use this window to choose drives to create drive groups.

Step 2 Press and hold **Ctrl** while choosing two or more ready drives in the **Drives** pane.

Choose all of the drives for the drive group.

Step 3 Click **Add To Array** to move the drives to a proposed drive group configuration in the **Drive Groups** pane.

If you need to undo the changes, click **Reclaim**.

Step 4 From the **Encryption** drop-down list, choose an encryption option.

Step 5 If needed, finish adding drives to the **Drive Groups** pane and click **Accept DG**.

Step 6 Click **Next**.

The **Virtual Drive Definition** window appears. This window lists the possible RAID levels for the drive group. Use this window to choose the RAID level, strip size, read policy, and other attributes for the new virtual drives.

Step 7 Choose the virtual drive options using the drop-down list in the left pane (the default values are shown).

- From the **RAID Level** drop-down list, choose the desired RAID option (RAID 0, 1, 5, or 6).

All of the possible RAID levels for the virtual drive are listed.

- From the **Strip Size** drop-down list, choose one of the following:
8, 16, 32, 64, 128, 256, 512, 1024 KB. The default is 64 KB.

The strip size is the portion of a stripe that resides on a single drive in the drive group. The stripe consists of the data segments that the RAID controller writes across multiple drives, not including parity drives.

For example, consider a stripe that contains 64 KB of drive space and has 16 KB of data residing on each drive in the stripe. In this case, the stripe size is 64 KB and the strip size is 16 KB. A larger strip size produces higher read performance. If your server regularly performs random read requests, choose a smaller strip size.

- From the **Access Policy** drop-down list, choose one of the following:
 - **RW**—Allows read/write access. This is the default.
 - **Read Only**—Allows read-only access.
 - **Blocked**—Does not allow access.
- From the **Read Policy** drop-down list, choose one of the following:
 - **Normal**—Disables the read-ahead capability. This is the default.
 - **Ahead**—Enables read-ahead capability, which allows the controller to read sequentially ahead of requested data and to store the additional data in cache memory, anticipating that the data will be needed soon. This option speeds up reads for sequential data but there is little improvement when accessing random data.
- From the **Write Policy** drop-down list, choose one of the following:
 - **WBack**—In Writeback mode, the controller sends a data transfer completion signal to the host when the controller cache has received all of the data in a transaction. This setting is recommended in Standard mode.
 - **WThru**—In Writethrough mode, the controller sends a data transfer completion signal to the host when the drive subsystem has received all of the data in a transaction. This is the default.
 - **Bad BBU**—Choose this mode if you want the controller to use Writeback mode, but the controller has no BBU or the BBU is bad. If you do not choose this option, the controller firmware automatically switches to Writethrough mode if it detects a bad or missing BBU.



Caution

The LSI WebBIOS CU allows Writeback mode to be used with or without a battery. We recommend that you use either a battery to protect the controller cache or an uninterruptable power supply (UPS) to protect the entire system. If you do not use a battery or a UPS, and there is a power failure, you risk losing the data in the controller cache.

- From the **IO Policy** drop-down list, choose one of the following:
 - **Direct**—In direct I/O mode, reads are not buffered in the cache memory. Data is transferred to the cache and the host concurrently. If the same data block is read again, it comes from the cache memory. This is the default.

- **Cached**—In cached I/O mode, all reads are buffered in the cache memory.

The IO Policy applies to reads on a specific virtual drive. It does not affect the read-ahead cache.

- From the **Drive Cache** drop-down list, choose one of the following:
 - **Enable**—Enables the drive cache.
 - **Disable**—Disables the drive cache.
 - **NoChange**—Leaves the current drive cache policy as is. This is the default.
- From the **Disable BGI** drop-down list, choose the background initialization status from the following:
 - **No**—Leaves background initialization enabled, which means that a new configuration can be initialized in the background while you use WebBIOS to do other configuration tasks. This is the default.
 - **Yes**—Does not allow background initializations for configurations on this controller.
- From the **Select Size** drop-down list, choose the size of the virtual drive in **MB**.
Normally, this would be the full size for RAID 0 shown in the Configuration pane. You can specify a smaller size if you want to create other virtual drives in the same drive group.

Step 8 Click **Accept** to accept the changes to the virtual drive definitions.

If you need to undo the changes, click **Reclaim**.

Step 9 After you finish the virtual drive definitions, click **Next**.

The **Configuration Preview** window appears.

Step 10 Check the virtual drive configuration in the **Configuration Preview** window and choose one of the following:

- If the virtual drive configuration is acceptable, click **Accept** to save the configuration.
- Click **Back** to return to the previous windows and change the configuration.

Step 11 Click **Yes** at the prompt to save the configuration.

The main **WebBIOS CU** window appears.

Using Manual Configuration (RAID 00, 10, 50, 60)

For more information about RAID levels, see [RAID Levels, page 1-8](#).

Step 1 Choose **Manual Configuration** and click **Next**.

The **Drive Group Definition** window appears. Use this window to choose drives to create drive groups.

Step 2 Press and hold **Ctrl** while choosing ready drives in the **Drives** pane.

Choose all of the drives for the drive group.

Step 3 Click **Add To Array** to move the drives to a proposed drive group configuration in the **Drive Groups** pane.

If you need to undo the changes, click **Reclaim**.

Step 4 If needed, finish adding drives to the **Drive Groups** pane and click **Accept DG** to create a RAID drive group.

An icon for the next drive group appears in the **Drive Groups** pane.

Step 5 Choose the drive group created in [Step 4](#), and press and hold **Ctrl** while choosing more ready drives in the **Drives** pane to create a second RAID drive group.

Step 6 Click **Add To Array** to move the drives to a second drive group configuration in the **Drive Groups** pane. If you need to undo the changes, click **Reclaim**.



Note RAID 00 supports a maximum of eight spans with a maximum of 32 drives per span.



Note RAID 10 supports a maximum of eight spans with a maximum of 32 drives per span. You must use an even number of drives in each RAID 10 drive group in the span.

Step 7 From the **Encryption** drop-down list, choose an encryption option.

Step 8 If needed, finish adding drives to the **Drive Groups** pane and click **Accept DG** to create a RAID 0 drive group.

Step 9 Repeat [Step 4](#) through [Step 6](#) until you have chosen all the drives you want for the drive groups.

Step 10 After you finish adding drives to the **Drive Groups** pane, choose each drive group, and click **Accept DG** after each drive group choice.

Step 11 Click **Next**.

The **Span Definition** window appears. This window shows the drive group holes that you can choose to add to the **Span** pane.

Step 12 Press and hold **Ctrl** while you choose a drive group in the **Array With Free Space** pane, and click **Add to SPAN**.

The drive group you chose appears in the **Span** pane.

Step 13 Press and hold **Ctrl** while you choose a second drive group, and click **Add to SPAN**.

Step 14 Repeat [Step 12](#) and [Step 13](#) until you have chosen all of the drive groups that you need.

Step 15 Click **Next**.

The **Virtual Drive Definition** window appears. This window lists the possible RAID levels for the drive group. Use this window to choose the RAID level, strip size, read policy, and other attributes for the new virtual drives.

Step 16 Press and hold **Ctrl** to choose the drive group in the **Configuration** pane, and then choose the virtual drive options using the drop-down lists in the left pane (the default values are shown):

- From the **RAID Level** drop-down list, choose The desired RAID level (RAID 00, 10, 50, or 60).

All of the possible RAID levels for the virtual drive are listed.

- From the **Strip Size** drop-down list, choose one of the following:
8, 16, 32, 64, 128, 256, 512, 1024 KB. The default is 64 KB.

The strip size is the portion of a stripe that resides on a single drive in the drive group. The stripe consists of the data segments that the RAID controller writes across multiple drives, not including parity drives.

For example, consider a stripe that contains 64 KB of drive space and has 16 KB of data residing on each drive in the stripe. In this case, the stripe size is 64 KB and the strip size is 16 KB. A larger strip size produces higher read performance. If your server regularly performs random read requests, choose a smaller strip size.

- From the **Access Policy** drop-down list, choose one of the following:

- **RW**—Allows read/write access. This is the default.
- **Read Only**—Allows read-only access.
- **Blocked**—Does not allow access.
- From the **Read Policy** drop-down list, choose one of the following:
 - **Normal**—Disables the read-ahead capability. This is the default.
 - **Ahead**—Enables read-ahead capability, which allows the controller to read sequentially ahead of requested data and to store the additional data in cache memory, anticipating that the data will be needed soon. This speeds up reads for sequential data but there is little improvement when accessing random data.
- From the **Write Policy** drop-down list, choose one of the following:
 - **WBack**—In Writeback mode, the controller sends a data transfer completion signal to the host when the controller cache has received all of the data in a transaction. This setting is recommended in Standard mode.
 - **WThru**—In Writethrough mode, the controller sends a data transfer completion signal to the host when the drive subsystem has received all of the data in a transaction. This is the default.
 - **Bad BBU**—Choose this mode if you want the controller to use Writeback mode but the controller has no BBU or the BBU is bad. If you do not choose this option, the controller firmware automatically switches to Writethrough mode if it detects a bad or missing BBU.

**Caution**

LSI allows Writeback mode to be used with or without a battery. LSI recommends that you use *either* a battery to protect the controller cache or an uninterruptable power supply (UPS) to protect the entire system. If you do not use a battery or a UPS, and there is a power failure, you risk losing the data in the controller cache.

- From the **IO Policy** drop-down list, choose one of the following:
 - **Direct**—In direct I/O mode, reads are not buffered in the cache memory. Data is transferred to the cache and the host concurrently. If the same data block is read again, it comes from the cache memory. This is the default.
 - **Cached**—In cached I/O mode, all reads are buffered in the cache memory.

The IO Policy applies to reads on a specific virtual drive. It does not affect the read-ahead cache.
- From the **Drive Cache** drop-down list, choose one of the following:
 - **Enable**—Enables the drive cache.
 - **Disable**—Disables the drive cache.
 - **NoChange**—Leaves the current drive cache policy as is. This is the default.
- From the **Disable BGI** drop-down list, choose the background initialization status from the following:
 - **No**—Leaves background initialization enabled, which means that a new configuration can be initialized in the background while you use WebBIOS to do other configuration tasks. This is the default.
 - **Yes**—Does not allow background initializations for configurations on this controller.
- From the **Select Size** drop-down list, choose the size of the virtual drive in **MB**.
Normally, this would be the full size for RAID 00 shown in the Configuration pane. You can specify a smaller size if you want to create other virtual drives in the same drive group.

- Step 17** Click **Accept** to accept the changes to the virtual drive definitions.
If you need to undo the changes, click **Reclaim**.
- Step 18** After you finish the virtual drive definitions, click **Next**.
The **Configuration Preview** window appears.
- Step 19** Check the virtual drive configuration in the **Configuration Preview** window and choose one of the following:
- If the virtual drive configuration is acceptable, click **Accept** to save the configuration.
 - Click **Back** to return to the previous windows and change the configuration.
- Step 20** Click **Yes** at the prompt to save the configuration.
The main **WebBIOS CU** window appears.
-

Viewing and Changing Device Properties

This section describes how you can use the WebBIOS CU to view and change the properties for controllers, virtual drives, physical drives, and BBUs.

Viewing Controller Properties

WebBIOS displays information for one LSI SAS controller at a time. If your computer system has multiple LSI SAS controllers, you can view information for a different controller by clicking **Controller Selection** in the main window. When the Controller Selection window appears, choose the controller you want from the list.

To view the properties of the current controller, follow these steps:

-
- Step 1** Click **Controller Properties** in the main **WebBIOS CU** window.
There are three Controller Properties windows. The information in the first window is read-only and cannot be modified directly. The window lists the number of virtual drives that are already defined on this controller and the number of drives connected to the controller.
- Step 2** Click **Next** to view the second **Controller Properties** window.
- Step 3** Click **Next** to view the third **Controller Properties** window.
[Table 5-1](#) describes the entries and options listed in the second and third Controller Properties windows. LSI recommends that you leave these options at their default settings to achieve the best performance, unless you have a specific reason for changing them.
- Step 4** If you make changes to the options on these windows, click **Submit** to register them. If you change your mind, click **Reset** to return the options to their default values.
-

Table 5-1 Controller Properties Menu Options

Option	Description
Battery Backup	Indicates whether the chosen controller has a BBU. If present, you can click Present to view information about the BBU. For more information, see Viewing and Changing Device Properties, page 5-11 .
Set Factory Defaults	Loads the default MegaRAID WebBIOS CU settings. The default is No.
Rebuild Rate	Chooses the rebuild rate for drives connected to the chosen controller. The default is 30 percent. The rebuild rate is the percentage of system resources dedicated to rebuilding a failed drive. The higher the number, the more system resources are devoted to a rebuild.
BGI Rate	Chooses the amount of system resources dedicated to background initialization of virtual drives connected to the chosen controller. The default is 30 percent.
CC Rate	Chooses the amount of system resources dedicated to consistency checks of virtual drives connected to the chosen controller. The default is 30 percent.
Reconstruction Rate	Chooses the amount of system resources dedicated to reconstruction of drives connected to the chosen controller. The default is 30 percent.
Controller BIOS	Enables or disables the BIOS for the chosen controller. The default is Enabled. If the boot device is on the chosen controller, the BIOS must be enabled. Otherwise, the BIOS should be disabled or it might not be possible to use a boot device elsewhere.
NCQ	Native Command Queuing (NCQ) gives an individual drive the ability to optimize the order in which it executes the read and write commands. The default is Enabled.
Coercion Mode	Drive coercion is a tool for forcing drives of varying capacities to the same size so they can be used in a drive group. The coercion mode options are None, 128MB-way, and 1GB-way. The default is None. The number you choose depends on how much the drives from various vendors vary in their actual size. LSI recommends that you use the 1-GB coercion mode option.
SMART Polling	Determines how frequently the controller polls for drives reporting a Predictive Drive Failure (SMART: Self-Monitoring Analysis and Reporting Technology error). The default is 300 seconds (5 minutes).
Alarm Control	Enables, disables, or silences the onboard alarm tone generator on the controller. The default is Disabled.
Patrol Read Rate	Chooses the rate for patrol reads for drives connected to the chosen controller. The default is 30 percent. The patrol read rate is the percentage of system resources dedicated to running a patrol read. See “Patrol Read-Related Controller Properties” in the Chapter 5 of the <i>LSI MegaRAID Software User Guide</i> for additional information about patrol read.
Cache Flush Interval	Controls the interval (in seconds) at which the contents of the onboard data cache are flushed. The default is 4 seconds.
Spinup Drive Count	Controls the number of drives that spin up simultaneously. The default is 2 drives.

Table 5-1 Controller Properties Menu Options (continued)

Option	Description
Spinup Delay	Controls the interval (in seconds) between the spinup of drives connected to this controller. The delay prevents a drain on the system power supply that would occur if all drives spun up at the same time. The default is 12 seconds.
StopOnError	When enabled, this option stops the boot process when the controller BIOS encounters an error during boot-up. The default is Disabled.
Spin Down Delay Time	Controls the interval (in seconds) between the spindown of drives connected to this controller. The delay prevents a drain on the system's power supply that would occur if all drives spun down at the same time. The default is 30 minutes.
Stop CC on Error	When enabled, this option stops a consistency check when the controller BIOS encounters an error. The default is No.
Maintain PD Fail History	Maintains the history of all drive failures. The default is Disabled.
Schedule CC	Indicates whether the option to schedule the date and time for a consistency check is supported.

Viewing Virtual Drive Properties, Policies, and Operations

Step 1 Click the **Virtual Drives** icon in the right pane in the main **WebBIOS CU** window.

The **Virtual Drive** window displays the following information:

- **Properties** area—Displays the virtual drive's RAID level, state, capacity, and strip size information.
- **Policies** area—Lists the virtual drive policies that were defined when the storage configuration was created. For information about these policies, see [Using Manual Configuration, page 5-6](#).
 - Choose an option from any of the drop-down lists and click **Change** to change any of these policies.
- **Operations** area—Lists operations that can be performed on the virtual drive.

Choose from the following options and click **Go** to perform an operation:

- **Delete**—Deletes this virtual drive. For more information, see [Deleting a Virtual Drive, page 5-18](#).
- **Fast Init**—Initializes this virtual drive. A fast initialization quickly writes zeroes to the first and last 10-MB regions of the new virtual drive and completes the initialization in the background.
- **Slow Init**—Initializes this virtual drive. A slow initialization is not complete until the entire virtual drive has been initialized with zeroes. It is seldom necessary to use this option, because the virtual drive was already initialized when you created it.



Caution

Before you run an initialization, back up any data on the virtual drive that you want to save. All data on the virtual drive is lost when you initialize the drive.

- **CC**—Runs a consistency check on this virtual drive. For more information, see [Using Manual Configuration, page 5-6](#). (This option is not available for RAID 0 virtual drives.)

- **AdvOpers**—Accesses windows to remove drives, migrates RAID levels (that is, changes the virtual drive configuration by adding a drive and changing the RAID level).

See [Migrating the RAID Level of a Virtual Drive, page 5-20](#) for information about adding a drive to a virtual drive or migrating its RAID level.

- **Expand**—Increases the size of a virtual drive to occupy the remaining capacity in the drive group. In addition, you can add drives to the virtual drive to increase the capacity.



Note Drives can be added to a drive group only when there is a *single* virtual drive present on the drive group. If there are multiple virtual drives present on the drive group, the option to add an extra disk is not available.

See [Expanding a Virtual Drive, page 5-16](#) for the procedure you can use to expand a virtual drive.



Caution

Before you change a virtual drive configuration, back up any data on the virtual drive that you want to save.

Viewing Physical Drive Properties and Operations

The Physical Drive window displays the properties of a chosen drive and then enables you to perform operations on the drive.

To view drive properties and perform subsequent operations, follow these steps:

- Step 1** To access the **Physical Drive** window, choose one of the following options:
- In the main **WebBIOS CU** window, in the **Physical Drives** area (in the right side), click a drive. The **Physical Drive** window appears.
 - In the main **WebBIOS CU** window, click **Physical Drives** in the left pane. Then click a drive in the right pane. The **Physical Drive** window appears.

- Step 2** Click **Properties**, and click **Go**. The drive properties window for the chosen drive appears.

The drive properties are view-only and include the state of the drive.

- Step 3** If the drive state is **Online**, choose one of the following operations and click **Go**:

- **MakeDriveOffline**—Forces the drive offline.



Note If you force offline a good drive that is part of a redundant drive group with a hot spare, the drive rebuilds to the hot spare drive. The drive you forced offline goes into the Unconfigured Bad state. Access the BIOS utility to set the drive to the Unconfigured Good state.

- **Locate**—Makes the LED flash on the drive, which works only if the drive is installed in a drive enclosure.

- Step 4** If the drive state is **Unconfigured Good**, choose one of the four following operations and click **Go**:

- **Make Global HSP**—Makes a global hot spare available to all of the virtual drives.
- **Make Dedicated HSP**—Makes a hot spare dedicated to a specific virtual drive.

WebBIOS displays the global hot spare as **Global** and the dedicated hot spare as **Ded**. The icon for the dedicated hot spare displays under its associated virtual drive. The drive number, drive state, drive capacity, and drive manufacturer display.

- **Enclosure Affinity**—Specifies where the hot spare is used first.

If there are drive failures present on a split backplane configuration, then the hot spare will be used first on the backplane side that it resides in.

- **Prepare for Removal**—Prepares the drive for removal from the enclosure.

The **Prepare for Removal** feature is different from spinning a drive down into powersave mode, because it also involves flagging the drive as ready to remove. Therefore, if you choose to prepare a drive for removal, **Ready to Remove** displays in the device tree for that drive, instead of **Powersave**.

Viewing and Changing Battery Backup Unit Information

If your SAS controller has a battery backup unit (BBU), you can view information and change settings.

Step 1 Click **Controller Properties** in the main **WebBIOS CU** window. The first **Controller Information** window appears.

Step 2 Click **Next** to view the second **Controller Properties** window.

The second **Controller Properties** window appears. The **Battery Backup** field indicates whether the BBU is present.

Step 3 Click **Present** in the **Battery Backup** field.

The **Battery Module** window appears. This window contains the following information:

- Battery information
- Design information
- Capacity information
- Auto Learn properties and settings

Most of the Battery Module properties are view-only and are self-explanatory.

In the lower right corner of the window are the auto learn options. A *learning cycle* is a battery calibration operation performed by the controller periodically to determine the condition of the battery. You can change the learn delay interval (the length of time between automatic learning cycles) and the auto learn mode.



Note

LSI recommends leaving the learn delay interval and the auto learn mode at their default settings.

Managing RAID

This section includes information about maintaining and managing storage configurations.

Expanding a Virtual Drive

You can increase the size of a virtual drive to occupy the remaining capacity in a drive group.


Note

You can expand a virtual drive only if it is the *only* virtual drive present in the drive group. When multiple virtual drives are present in a drive group, you cannot expand the virtual drives.

In addition, you can add drives to the virtual drive to increase capacity.


Note

Drives can be added to a drive group only when there is a *single* virtual drive present on the drive group. When multiple virtual drives present on the drive group, the option to add an extra disk is not available.

-
- Step 1** Click a virtual drive icon in the right pane of the main **WebBIOS CU** window to access the **Virtual Drive** window. The **Virtual Drive** window appears.
 - Step 2** Choose the **Expand** radio button and click **Go**. The **Expand Virtual Drive** window appears.
 - Step 3** Enter the percentage of the available capacity that you want the virtual drive to use.
For example, if 100 GB capacity is available and you want to increase the size of the virtual drive by 30 GB, choose 30 percent.
 - Step 4** Click **Calculate** to determine the capacity of the virtual drive after expansion.
 - Step 5** Click **Ok**. The virtual drive expands by the chosen percentage of the available capacity.
-

Monitoring Array Health

You should periodically run a consistency check on fault-tolerant virtual drives. A consistency check verifies that the redundancy data is correct and available for RAID 1, RAID 5, RAID 6, RAID 10, RAID 50, and RAID 60 drive groups.

-
- Step 1** In the main WebBIOS CU window, choose a virtual drive.
 - Step 2** Click **Virtual Drives**.
 - Step 3** When the **Virtual Drive** window appears, click the **CC** radio button, and click **Go**.
The consistency check begins.

**Note**

If the WebBIOS CU finds a difference between the data and the parity value on the redundant drive group, it assumes that the data is accurate and automatically corrects the parity value. Be sure to back up the data before running a consistency check if you think the data might be corrupted.

Recovery

MegaRAID Recovery, also known as Snapshot, offers a simplified way to recover data and provides automatic protection for the boot volume. You can use the Recovery feature to take a snapshot of a volume and to restore a volume or file. Snapshot functionality allows you to capture data changes to the volume, and, if data is deleted accidentally or maliciously, you can restore the data from the view or roll back to a snapshot at a previous point-in-time (PiT). MegaRAID Recovery supports up to eight snapshots of PiTs for each volume.

Each Recovery PiT volume snapshot is typically a fraction of the original volume size, because it tracks only the changes that are made to a volume after the PiT is created. The disk space for PiTs is reserved in the Snapshot Repository virtual drive, and the PiT is expanded in small increments as new data is written to the volume. Multiple PiTs of each volume can be retained online, enabling frequent snapshots to be stored in a space-efficient manner.

**Caution**

Do not select the virtual drive containing the operating system (OS) as the Snapshot Repository. Updates to the operating system or operating system crashes could destroy data on that virtual drive.

Three primary scenarios in which to use the Recovery feature are as follows:

1. Restore the missing or deleted files (restore from view).
 - a. Discover that files are missing or deleted.
 - b. Review the Snapshot views of the file content (also known as “mounting” a snapshot) from each PiT until you find the missing file.

A Snapshot view contains the content from the Point-in-Time at which the snapshot was made.
 - c. Drag and drop the missing file from the Snapshot view back into the online storage volume that was the source of the Snapshot.
2. If there are corrupt operating system files in a volume, roll back the volume to a previous state:
 - a. Reboot the system and run WebBIOS.
 - b. Select the most recent snapshot that does not contain the corrupted or malicious file to roll back to. Select the most recent PiT snapshot to roll back to.
 - c. Reboot the system.

The system automatically rolls back to its previous state based on the selected PiT snapshot.
3. Reduce the risk of extended downtime during application updates/upgrades in the IT center:
 - a. When the application is offline, take a snapshot of the application volume.
 - b. Install each patch individually and test for any new defects that might have been introduced.
 - c. Take a snapshot after you test each patch and determine that it is clean.

- d. If a defect is introduced, roll back to the previous installation and bypass the installation of the defective patch.

**Note**

If the volume is still damaged, continue to select from the next most current PiT snapshot to the oldest.

You can enable the Recovery advanced software in WebBIOS. After you enable Recovery, you create two virtual drives: one as a Snapshot Base or source and the other as a Snapshot Repository. The Snapshot Base virtual drive contains the data that is stored in the repository virtual drive.

Deleting a Virtual Drive

You can delete any virtual drive on the controller if you want to reuse that space for a new virtual drive. The WebBIOS CU provides a list of configurable drive groups where there is a space to configure. If multiple virtual drives are defined on a single drive group, you can delete a virtual drive without deleting the whole drive group.

**Caution**

Back up any data that you want to keep before you delete the virtual drive.

- Step 1** Click a virtual drive icon in the right pane in the main **WebBIOS CU** window to access the **Virtual Drive** window. The **Virtual Drive** window appears.
 - Step 2** Click the **Delete** radio button in the **Operations** area, and click **Go**.
 - Step 3** When the message appears, confirm that you want to delete the virtual drive.
-

Migrating an Array to a New Server

A *foreign configuration* is a storage configuration that already exists on a set of drives that you install in a server. To migrate an existing array from one server to another, as in the case of a server replacement, the new server must have the same controller type as the server on which the array was initially created.

In addition, if one or more drives are removed from a configuration, by a cable pull or drive removal, for example, the configuration on those drives is considered a foreign configuration by the RAID controller.

The BIOS CU allows you to import the foreign configuration to the RAID controller, or to clear the configuration so you can create a new configuration using these drives.

**Note**

When you create a new configuration, the WebBIOS CU shows only the unconfigured drives. Drives that have existing configurations, including foreign configurations, do not appear. To use drives with existing configurations, you must first clear the configuration on those drives.

If WebBIOS CU detects a foreign configuration, the Foreign Configuration window appears.

- Step 1** In the **Foreign Configuration** window, click the **Select Configuration** drop-down list to show the configurations.

The GUID (Global Unique Identifier) entries in the drop-down list are OEM names and will vary from one installation to another.

Step 2 Choose a configuration or **All Configurations**.

Step 3 To preview or clear a foreign configuration, do one of the following:

- Click **Preview** to preview the foreign configuration. The **Foreign Configuration Preview** window appears.
- Click **Clear** to clear the foreign configuration and reuse the drives for another virtual drive.

If you click **Cancel**, it cancels the importation or preview of the foreign configuration.

The right pane shows the virtual drive properties of the foreign configuration. The left pane shows the drives in the foreign configuration.

Step 4 Click **Import** to import this foreign configuration and use it on this controller.

If you click **Cancel**, you return to the **Foreign Configuration** window.

Foreign Configurations in Cable Pull and Drive Removal Scenarios

If one or more drives are removed from a configuration, for example, by a cable pull or drive removal, the configuration on those drives is considered a foreign configuration by the RAID controller.

Use the **Foreign Configuration Preview** window to import or clear the foreign configuration in each case. The import procedure and clear procedure are described in [Migrating an Array to a New Server, page 5-18](#).



Note

If you want to import the foreign configuration in any of the following scenarios, you should have all of the drives in the enclosure before you perform the import operation.

The following scenarios can occur with cable pulls or drive removals:

- Scenario 1: If all of the drives in a configuration are removed and reinserted, the controller considers the drives to have foreign configurations.

Import or clear the foreign configuration. If you choose **Import**, automatic rebuilds occur in redundant virtual drives.



Note

Start a consistency check immediately after the rebuild is complete to ensure data integrity for the virtual drives.

- Scenario 2: If some of the drives in a configuration are removed and reinserted, the controller considers the drives to have foreign configurations.

Import or clear the foreign configuration. If you choose **Import**, automatic rebuilds occur in redundant virtual drives.



Note

Start a consistency check immediately after the rebuild is complete to ensure data integrity for the virtual drives.

- Scenario 3: If all of the drives in a virtual drive are removed, but at different times, and reinserted, the controller considers the drives to have foreign configurations.

Import or clear the foreign configuration. If you choose **Import**, all drives that were pulled *before* the virtual drive became offline are imported and automatically rebuilt. Automatic rebuilds occur in redundant virtual drives.

- Scenario 4: If the drives in a nonredundant virtual drive are removed, the controller considers the drives to have foreign configurations.

Import or clear the foreign configuration. No rebuilds occur after the import operation because there is no redundant data with which to rebuild the drives.

Importing Foreign Configurations from Integrated RAID to MegaRAID

The LSI Integrated RAID solution simplifies the configuration options and provides firmware support in its host controllers. LSI offers two types of Integrated RAID (IR): Integrated Mirroring (IM) and Integrated Striping (IS).

You can import an IM or IS RAID configuration from an IR system into a MegaRAID system. The MegaRAID system treats the IR configuration as a foreign configuration. You can import or clear the IR configuration.



Note

For more information about Integrated RAID, see [Chapter 4, “Configuring the LSI SAS2 Integrated RAID Controller.”](#)

Troubleshooting Information

An IR virtual drive can have either 64 MB or 512 MB available for metadata at the end of the drive. This data is in LSI Data Format (LDF). MegaRAID virtual drives have 512 MB for metadata at the end of the drive in the Disk Data Format (DDF).

To import an IR virtual drive into MegaRAID, the IR virtual drive must have 512 MB in the metadata, which is the same amount of metadata as in a MegaRAID virtual drive. If the IR virtual drive has only 64 MB when you attempt to import it into MegaRAID, the import fails because the last 448 MB of your data is overwritten and the data is lost.

If your IR virtual drive has only 64 MB for metadata at the end of the drive, you cannot import the virtual drive into MegaRAID. You need to use another upgrade method, such as backup/restore to the upgraded virtual drive type.

To import an IR virtual drive into a MegaRAID system, use the **Foreign Configuration Preview** window to import or clear the foreign configuration. The import procedure and the clear procedure are described in [Migrating an Array to a New Server, page 5-18](#).

Migrating the RAID Level of a Virtual Drive

As the amount of data and the number of drives in your system increase, you can use RAID-level migration to change a virtual drive from one RAID level to another. You do not have to power down or reboot the system. When you migrate a virtual drive, you can keep the same number of drives, or you can add drives. You can use the WebBIOS CU to migrate the RAID level of an existing virtual drive.



Note

While you can apply RAID-level migration at any time, LSI recommends that you do so when there are no reboots. Many operating systems issue I/O operations serially (one at a time) during boot. With a RAID-level migration running, a boot can often take more than 15 minutes.

Migrations are allowed for the following RAID levels:

- RAID 0 to RAID 1
- RAID 0 to RAID 5
- RAID 0 to RAID 6
- RAID 1 to RAID 0
- RAID 1 to RAID 5
- RAID 1 to RAID 6
- RAID 5 to RAID 0
- RAID 5 to RAID 6
- RAID 6 to RAID 0
- RAID 6 to RAID 5

Table 5-2 lists the number of additional drives required when you change the RAID level of a virtual drive.

Table 5-2 Additional Drives Required for RAID-Level Migration

From RAID Level to RAID Level	Original Number of Drives in the Drive Group	Additional Drives Required
RAID 0 to RAID 1	RAID 0: 1 drive	1
RAID 0 to RAID 5	RAID 0: 1 drive	2
RAID 0 to RAID 6	RAID 0: 1 drive	3
RAID 1 to RAID 5	RAID 1: 2 drives	1
RAID 1 to RAID 6	RAID 1: 2 drives	1



Caution

Back up any data that you want to keep before you change the RAID level of the virtual drive.

- Step 1** In the main **WebBIOS CU** window, choose a virtual drive.
- Step 2** Click **Virtual Drives**. The **Virtual Drive** window appears.
- Step 3** Click the **AdvOpers** radio button. The **Advanced Operations** window appears.
- Step 4** Choose either **Change RAID Level** or **Change RAID Level and Add Drive**.
- If you choose **Change RAID Level**, change the RAID level from the drop-down list.
 - If you choose **Change RAID Level and Add Drive**, change the RAID level from the drop-down list and choose one or more drives to add from the list of drives.
- The available RAID levels are limited, based on the current RAID level of the virtual drive plus the number of drives available.
- Step 5** Click **Go**.
- Step 6** When the message appears, confirm that you want to migrate the RAID level of the virtual drive.

A reconstruction operation begins on the virtual drive. You must wait until the reconstruction is completed before you perform any other tasks in the WebBIOS CU.

Determining Which Controller is in Your Server

There is a dedicated SAS riser slot for the RAID controller card in a C-series chassis. There is also a mounting point inside the chassis for the optional RAID battery backup unit that is available when using the appropriate LSI controller.

Supported RAID controllers for all models are listed in [RAID Controllers in UCS Servers, page 3-10](#).

If you do not have a record of which option is used in your server, you can read the on-screen messages that are displayed during system bootup

Information about the card models installed are displayed as part of the verbose boot. You are also prompted to press **Ctrl-H** to launch configuration utilities for those cards. For servers running CIMC firmware earlier than release 1.2(1), see also [Disabling Quiet Boot for CIMC Firmware Earlier than Release 1.2\(1\), page 5-22](#).

Disabling Quiet Boot for CIMC Firmware Earlier than Release 1.2(1)

For CIMC firmware and BIOS release 1.2(1) and later releases, Quiet Boot has been removed and is no longer needed. If you are running CIMC firmware and BIOS earlier than release 1.2(1), you can use the following procedure to disable Quiet Boot.

-
- Step 1** Boot the server and watch for the F2 prompt.
 - Step 2** Press **F2** when prompted to enter the BIOS Setup utility.
 - Step 3** On the Main page of the BIOS Setup utility, set Quiet Boot to **Disabled**.
This action allows non default messages, prompts, and POST messages to display during bootup instead of the Cisco logo window.
 - Step 4** Press **F10** to save your changes, and to exit the utility.
-

Launching an Option ROM-Based Controller Utility

To alter the RAID configurations on your hard drives, you can use your host-based utilities that you install on top of your host OS, or you can use the LSI option ROM-based utilities that are installed on the server.

When you boot the server and you have quiet boot disabled, information about your controller is displayed. The prompts for the key combination to launch the option ROM-based utilities for your controller are also displayed.

-
- Step 1** Watch for the prompt for your controller during verbose boot:
- Step 2** Press **Ctrl-H** for the LSI controller utility.
-

**Note**

Cisco has also developed the Cisco Server Configuration Utility for C-Series servers, which can assist you in setting up some RAID configurations for your drives. This utility is shipped with new servers on CD. You can also download the ISO from Cisco.com. See the user documentation for this utility at the following URL:

http://www.cisco.com/en/US/docs/unified_computing/ucs/sw/ucsscu/user/guide/20/SCUUG20.html

LSI MegaRAID Card Beep Codes

Table 5-3 contains a summary of the LSI MegaRAID card beep codes. These beep codes indicate activity and changes from the optimal state of your RAID array. For full documentation on the LSI MegaRAID cards and the LSI utilities, refer to the LSI documentation for your card.

Table 5-3 Summary of LSI MegaRAID Card Beep Codes

Beep Code	LSI Firmware State	Cause (Depending on RAID Level)
3 seconds on, 1 second off	SPEAKER_OFFLINE_ENTRY	<ul style="list-style-type: none"> RAID 0: One or more drives offline. RAID 1: Two drives offline. RAID 5: Two or more drives offline. RAID 6: More than two drives offline.
1 second on, 1 second off	SPEAKER_DEGRADED_ENTRY	<ul style="list-style-type: none"> RAID 1: A mirrored drive failed. RAID 5: One drive failed. RAID 6: One or two drives failed.
1 second on, 3 seconds off	SPEAKER_HOTSPARE_ENTRY	A hot spare drive has completed the rebuild process and has been brought into the array.

Restoring the RAID Configuration After Replacing a RAID Controller

When you replace a RAID controller, the RAID configuration that is stored in the controller is lost. Use the following procedure to restore your RAID configuration to your new RAID controller.

-
- Step 1** Replace your RAID controller. See your server documentation for specific steps.
- Step 2** If you are performing a full chassis swap, replace all drives into the drive bays, in the same order that they were installed in the old chassis.
- Step 3** If Quiet Boot is enabled, disable it in the system BIOS. See [Disabling Quiet Boot for CIMC Firmware Earlier than Release 1.2\(1\)](#), page 5-22.
- Step 4** Reboot the server and watch for the message to press F.
- Step 5** Press **F** when you see the following on-screen message:

Foreign configuration(s) found on adapter.
Press any key to continue or 'C' load the configuration utility,
or 'F' to import foreign configuration(s) and continue.

Step 6 Press any key (other than C) to continue when you see the following on-screen message:

All of the disks from your previous configuration are gone. If this is
an unexpected message, then please power of your system and check your cables
to ensure all disks are present.
Press any key to continue, or 'C' to load the configuration utility.

Step 7 Watch the subsequent windows for confirmation that your RAID configuration was imported correctly.

- If you see the following message, your configuration was successfully imported. The LSI virtual drive is also listed among the storage devices.

```
N Virtual Drive(s) found on host adapter.
```

- If you see the following message, your configuration was not imported which can happen if you do not press F quickly enough when prompted. In this case, reboot the server and try the import operation again when you are prompted to press F.

```
0 Virtual Drive(s) found on host adapter.
```

Limitation on Importing Foreign Configuration To a Virtual Disk That is Under Construction

Limitation

When the reconstruction operation of a virtual disk (VD) is in progress and the controller fails or the VD changes to an offline state, you cannot import the foreign configuration and the VD cannot be recovered.

Design

The Avago MegaRAID firmware provides the ability to migrate a VD from any basic RAID level to any basic RAID level without affecting the system availability or disrupting any other functionality. This is called RAID level migration (RLM). The firmware verifies that the disk group that contains the virtual drives that are being migrated has sufficient space to complete the migration before the migration operation starts. The user data size of the target virtual drive must be greater than or equal to the source virtual drive.

The limitations of this design are as follows.

- The RLM operation cannot be aborted once started.
- If all of the VDs that are undergoing the RLM operation change their state to offline (drives are pulled out or the enclosure powers off) and then if drives return, the VD is foreign and in offline state.
- The MegaRAID firmware stores some of the reconstruction-specific information in the NVRAM but does not store the reconstruction check point data or metadata in the disk data format (DDF) area. Then the VD enters into the morphing state.

During the next reboot, if the VD is still in the offline mode, the user is notified with a boot message that some of the virtual drives might become offline. A VD cannot be imported when it is in the morphing state because the firmware cannot determine the state or type of the morphing activity. These VDs are unusable and the user cannot import them.

Prerequisites For Reconstruction to Start

Firmware performs checks before starting the reconstruction operation on a VD. It fails the reconstruction command if any of the following conditions exist:

- Reconstruction is already active (multiple reconstructions are not allowed).
- Rebuild operation is active.
- Copyback operation is active.
- Check Consistency is active.
- Background Initialization is active.
- Virtual drive initialization is active.
- All the virtual drives (maximum supported) are configured.
- All the arrays (maximum supported) are configured.
- Logical volume to be reconstructed is OFFLINE.
- If the new RAID level is not equal to 0, 1, 5 and 6.
- Operation other than INSERT, DELETE and NONE.
- New Row Size or New PD Count is greater than maximum disks per array.
- Premium feature RAID5 is disabled and the new RAID level will be RAID5 or premium feature RAID6 is disabled and the new RAID level will be RAID6.
- Pinned cache is present.
- More than one VD in the disk group and total number of spans for the VD is greater than 1.
- R1E virtual drive which is shown as R10.

The firmware blocks the following operations if reconstruction is in progress:

- Any descriptor command (DCMD) related to configuration changes
- Logical drive initialization
- Rebuild Drive
- Second Reconstruction
- Flashing RAID Firmware
- Check Consistency
- Background Initialization
- Hot spare commissioning
- Patrol read
- Turning off Protection Information (PI), which will remove protection from any existing logical disks.