

HP Embedded SATA RAID Controller

User Guide



Legal notices

© Copyright 2006 Hewlett-Packard Development Company, L.P.

The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Microsoft and Windows are U.S. registered trademarks of Microsoft Corporation.

Part number: 433600-001

First Edition: June 2006

Abstract

This guide provides reference information for using the HP Embedded SATA RAID Controller to configure SATA drive arrays.

Audience assumptions

This guide is for the person who installs, administers, and troubleshoots servers. HP assumes you are qualified in the servicing of computer equipment and trained in recognizing hazards in products with hazardous energy levels.

Contents

1	About the controller	
	Installing the driver.....	4
	Methods for configuring an array.....	4
2	Using the RAID Configuration Utility	
	Creating an array	5
	Managing arrays	6
	Adding or deleting a hotspare	7
	Configuring a drive	7
	Rebuilding an array	7
	Verifying an array.....	8
3	Using the Command Line Utility	
	Starting the command line utility	9
	Using the command line utility	9
	HRCONF commands	9
4	Understanding RAID	
	RAID technology overview	14
	Understanding drive segments	14
	Stripe-unit size	14
	Selecting a RAID level and tuning performance	14
	RAID 0	15
	RAID 1	15
	RAID 10	16

1 About the controller

The HP Embedded SATA RAID Controller enables you to create arrays with RAID levels 0, 1, and 10 on SATA hard drives in specified HP ProLiant server models.



NOTE: Some servers that use the HP Embedded SATA RAID Controller do not support RAID 10. For information about supported RAID levels on a particular server, see the user guide for that server.

For a complete list of the controller attributes, see the server page on the HP website (<http://www.hp.com/support>). For information about the compatibility of particular operating systems with the controller, see the user guide for the server in which the controller is used.

Installing the driver

A link to the downloadable driver for the controller is provided on the server-specific page of the HP website, <http://www.hp.com/support>. Installation instructions are provided with the driver.

Methods for configuring an array

Three utilities are provided to enable you to configure arrays on the controller:

- The menu-based RAID Configuration Utility, located on the controller BIOS, enables you to configure an array even before the driver or operating system is installed. However, this utility can be opened and used only while the system is rebooting, before any installed operating system is loaded. Also, you cannot change the RAID level, stripe size, or ID label of an existing array using this tool.

For detailed information about using this tool, see “Using the RAID Configuration Utility” on page 5.

- The GUI-based HP Storage Manager is provided on the HP website as an executable file that you install after you have installed the driver and operating system. This utility enables you to create new arrays and change the properties of an existing array. You can use this utility at any convenient time after the system has booted and the operating system has loaded.

For detailed information about using this tool, see the *HP Storage Manager User Guide* provided with the utility.

- The command line utility, HRCNF, is also provided on the HP website. This utility enables you to create new arrays and change the properties of an existing array. You can use this utility at any convenient time after the system has booted and the operating system has loaded.

For detailed information about using this tool, see “Using the Command Line Utility” on page 9.

2 Using the RAID Configuration Utility

The RAID Configuration Utility is loaded on the controller ROM. To run the utility, press the **Ctrl+A** keys or the **F8** key (depending on the version of the ROM) when prompted by one of the following messages during system startup:

Press <Ctrl><A> for Adaptec RAID Configuration Utility

Press <F8> for HP Embedded SATA Setup

To select an option from any of the menus within the RAID Configuration Utility, browse with the arrow keys, and then press the **Enter** key. In some cases, selecting an option displays another menu. To return to the previous menu at any time, press the **Esc** key.

The main option on the RAID Configuration Utility menu is the array configuration utility, which enables you to perform the following tasks:

- Create an array
- Manage an array
- Add or delete a hotspare
- Configure a drive
- Rebuild an array
- Verify drives

The procedures for performing each of these tasks are described in the following sections.

Creating an array

Before creating an array, be sure that the drives to be used for the array are installed in your system and have been configured (see “Configuring a drive” on page 7). If the utility displays a drive in gray, the drive does not have any usable space, and it cannot be used to create an array.



CAUTION: Before adding a drive to an array, back up any data that is already on the drive. Otherwise, the data will be lost.

To create an array:

1. Power down the computer, and then restart it.
2. When the appropriate prompt appears during POST, press the **Ctrl+A** keys or the **F8** key to access the RAID Configuration Utility.
3. From the menu, select **Array Configuration Utility**.
4. From the array configuration utility menu, select **Create Array**.
5. Highlight a drive to be used in the new array, and then press the **Insert** key.
6. Repeat the previous step to select another drive. (To deselect a drive, highlight the drive and press the **Delete** key.)
7. When all drives for the new array are selected, press the **Enter** key.

The Array Properties menu appears, showing the RAID levels (types of array) that can be configured for the selected number of drives. (For this controller, RAID 0 and RAID 1 are shown, and RAID 10 is shown if the server supports it. RAID 0 can have from two to four drives, RAID 1 requires two drives, and RAID 10 requires four drives.)
8. Select the RAID level that you want to create, and press the **Enter** key.
9. (Optional) Enter a label of no more than 15 characters to identify the array.
10. If you are creating a RAID 0 or RAID 10 array, you can change the stripe size that the array uses, although in general it is best to stay with the default size of 64 KB. Available stripe sizes are 16, 32, and 64 KB.

11. Select the method that the controller is to use to build the array.

CAUTION: Do not use the Migrate or Build methods to create an array on Microsoft® Windows® dynamic disks (volumes), or you will lose data.

Table 1 Choosing a method for creating the array

RAID level	Method	Method is appropriate for...
0	Quick Init	Creating a RAID 0 array on new drives.
0	Migrate	Creating a RAID 0 array when you want to preserve data on an existing drive. If you select this option, you are asked to designate the source drive. The contents of the source drive are preserved, and any data on the new drive is lost. When migrating from single volume to a RAID 0 array, migrating from a larger drive to a smaller drive is allowed. However, the destination drive must have at least half the capacity of the source drive.
1	Build	Creating a RAID 1 array when you want to preserve data on an existing drive. If you select this option, you are asked to designate the source drive. The contents of the source drive are preserved, and any data on the new drive is lost. Also, if you use drives of different sizes, only the smaller drive can be selected as the source drive.
1 or 10	Clear	Creating a RAID 1 or RAID 10 array on new drives or when you want to be sure that the new array contains no existing data.
1 or 10	Quick Init	Creating a RAID 1 or RAID 10 array on new drives. This method is the fastest way to create a RAID 1 or RAID 10 array.

12. Click **Done**.

Other factors to consider when building an array:

- If you pause a Build or Clear process on a RAID 1 array, or if you pause a Migrate process on a RAID 0 array, you can restart the process by pressing the **Ctrl+R** keys.
- A RAID 1 or RAID 10 array created using the Quick Init option might return some data mismatches if you later run a consistency check. This behavior is normal and is not a cause for concern.
- When you create an array, write cache is disabled by default. If you operate an array with write cache off, disk operation might be very slow. To enable the write cache, press the **Ctrl+W** key combination.

Managing arrays

With the Manage Arrays option, you can perform the following tasks:

- View array properties
- Make an array bootable
- Change the disk write cache setting
- Delete an array

Viewing array properties

1. From the main menu of the array configuration utility, select **Manage Arrays**, and press the **Enter** key.
2. From the List of Arrays dialog box, select the array whose properties you want to view, and press the **Enter** key.
The Array Properties dialog box appears, showing detailed information about the array. The physical disks associated with the array appear here.
3. Press the **Esc** key to return to the previous menu.

Making an array bootable

1. From the main menu of the array configuration utility, select **Manage Arrays**, and press the **Enter** key.
2. From the List of Arrays, select the array that you want to make bootable, and press the **Ctrl+B** keys. An asterisk appears next to the array to indicate that it is bootable.

To make an array nonbootable, select the boot array (the one with the asterisk) and press the **Ctrl+B** keys.

Changing the disk write cache setting

To enable or disable the drive write cache for an array:

1. From the main menu of the array configuration utility, select **Manage Arrays**, and press the **Enter** key.
2. From the List of Arrays, select the array for which you want to modify the write cache setting, and press the **Ctrl+W** keys. A dialog box appears asking you to confirm that you want to modify the setting.
3. Press the **Y** key to change the current setting.

Deleting an array

1. Back up any important data on the array.
2. From the main menu of the array configuration utility, select **Manage Arrays**, and press the **Enter** key.
3. Select the array that you want to delete, and press the **Delete** key.
4. In the Array Properties dialog box, select **Delete**, and press the **Enter** key. The following prompt appears:
Warning!! Deleting the array will render the array unusable.
Do you want to delete the array? (Yes/No):
5. **Enter** Yes. The following prompt appears:
To delete the partition table, choose which member:
member #0, member #1, both, none
6. Select the member to be deleted.
7. Press the **Esc** key to return to the previous menu.

Adding or deleting a hotspare

1. From the main menu of the array configuration utility, select **Add/Delete Hotspares**.
2. Use the arrow keys to highlight the drive that you want to assign as a hotspare, and press the **Insert** key.
3. Press the **Enter** key. The following prompt appears:
Do you want to create spare? (Yes/No)
4. Enter Yes. The spare that you have selected appears in the Selected Drive menu.

Configuring a drive



CAUTION: If you configure a drive that is already part of an array, the array might become unusable. Do not configure a drive that is part of a boot array. To determine which drives are associated with a particular array, see "Viewing array properties" on page 6.

1. From the main menu of the array configuration utility, select **Configure Drives**.
2. Use the arrow keys to highlight the drive that you want to configure, and press the **Insert** key.
3. If you want to add another drive to be configured, repeat the previous step.
4. Press the **Enter** key.
5. Read the warning message, and confirm that you have selected the correct drives to configure.
6. Press the **Y** key to continue.

Rebuilding an array

If a drive in a fault-tolerant (RAID 1 or RAID 10) array fails, you can return the array to optimal status by performing a rebuild.

There are two ways to perform a rebuild:

- Shut down the system, replace the failed drive, and begin the rebuild at the next system restart.
- Add a hotspare to the array if necessary, and select the array for a manual rebuild.

System shutdown rebuild

1. Power down the system.
2. Replace the failed drive with a new one of equal or greater capacity.
3. Restart the system.

4. At the appropriate prompt during POST, press the **Ctrl+A** keys or the **F8** key to access the RAID Configuration Utility.
5. From the menu, select **Array Configuration Utility**.
6. From the main menu of the array configuration utility, select **Add/Delete Hotspares**.
7. Assign the new drive as a hotspare.

The Rebuild task begins. All the data from the good drive is copied to the new one, and the original array is re-created.

Manual rebuild



IMPORTANT: If a hard drive in the array fails and the array does not have a hotspare, you must add a hotspare before you can rebuild the array. See “Adding or deleting a hotspare” on page 7.

1. From the main menu of the array configuration utility, select **Manage Arrays**.
2. From the List of Arrays, select the array that you want to rebuild.
3. Press the **Ctrl+R** keys to begin the rebuild process.

Verifying an array

If you are notified of a data mismatch during RAID 1 or RAID 10 array creation by a Build operation, you can verify the array to determine a possible cause of the mismatch. However, you cannot verify an array if it has failed.

To verify the array, highlight the array and then press the **Ctrl+S** keys.

When the Verify operation is complete, you are notified of any errors found.

3 Using the Command Line Utility

The command line utility, HRCONF, enables you to perform the following tasks:

- Create and delete logical drives
- Display and modify a limited set of configuration settings
- Recover from a failed physical drive and rebuild an affected logical drive
- Detect the removal and connection of any disk drives
- Restore the controller configuration

HRCONF is provided on the HP website at <http://www.hp.com/support>. The file is also called the HP Embedded SATA RAID Controller Command Line Interface. Installation instructions are provided with the utility.

Starting the command line utility

To start HRCONF, enter one of the following commands:

Windows: `c:\<install_dir*>\hrconf.exe` (where `install_dir*` is the directory where HRCONF is installed)

Linux: `/usr/StorMan/hrconf.exe`

To see a list of available commands, enter `HRCONF` at the prompt. The utility command functions are detailed in “Using the command line utility” on page 9.

Using the command line utility

This section explains how to use the command line utility interactively or in batch mode. With interactive mode, enter commands at the prompt. In batch mode, create scripts and run the script in the appropriate shell. For example:

Table 2 Shells for running scripts

Environment	Batch file	Run script
Windows	.bat	CMD.EXE
Linux	.sh	sh / bash

In either mode, if your command fails, you immediately see an error message of `Command failed`. Other script messages that you can get are `Command completed successfully` or `Command aborted`.

Available commands are described on the following pages, in alphabetical order.

HRCONF commands

Perform the following functions from the command line:

Table 3 HRCONF commands

backup	getstatus	setconfig
create	getversion	setstate
delete	rescan	task
getconfig	restore	

hrconf backup

For large-scale deployments, this command stores the current controller and disk drive configuration setting to a specific file. Stored files can be used with the `RESTORE` command to restore to another controller or disk drive. To restore, the controller or disk drive must have the same configuration as it did before the backup. In other words, you must use the same type of controller with the same number and type of disk drives, which have the same IDs and channels.

Syntax

```
hrconf backup <Controller#> <Filename>
```

Parameters

Controller# is the controller number.

Filename is the relative or absolute path with file name.

Example

```
hrconf backup 1 c:\windows\hr2200
```

Return values

SUCCESS: 0x00 indicates the command completed successfully.

FAILURE: 0x01 indicates the command failed.

hrconf create

This command creates logical drives. You must provide the channel and device ID of the physical drives. On redundant logical drives, HRCONF performs autosynchronization.

Syntax

```
hrconf create <Controller#> LOGICALDRIVE [Options] <Size> <RAID#> <CHANNEL#  
DRIVE#> [CHANNEL# DRIVE#] ... [noprompt]
```

Parameters

Controller# is the controller number.

Options indicates the logical drive with the following options:

- **Stripesize**—Optional parameters for specifying a stripe size in kilobytes. Values of 16, 32, and 64 are supported.
- **Name**—Optional parameter for specifying the name of the logical drive to be created.
- **Init_Priority**—Initialization priority for logical drive to be created. Valid parameters are HIGH, MED, or LOW.
- **Init_Method**—Initialization method for the logical drive. Valid options are CLEAR and QUICK.

Size indicates the size of the logical drive. MAX is the only size option available.

RAID# indicates the RAID level for the logical drive (0, 1, 10 volume).

Channel# is the channel number for the device.

Drive# is the device number for the device.

noprompt: No prompt for confirmation.

Example

```
hrconf create 1 logicaldrive stripesize 64 max 0 1 0 1 1 1 2
```

Return values

SUCCESS: 0x00 indicates successful termination.

FAILURE: 0x00 indicates a bad arguments or internal error.

hrconf delete

This command deletes a logical drive. All data stored on the logical drive will be lost.

Syntax

```
hrconf delete <Controller#> LOGICALDRIVE <LogicalDrive#> [noprompt]
```

Parameters

Controller# is the controller number.

LogicalDrive# is the number of the logical drive to be deleted.

noprompt: No prompt for confirmation.

Example

```
hrconf delete 1 logicaldrive 1 noprompt
```

Return values

SUCCESS: 0x00 indicates successful termination.

FAILURE: 0x01 indicates bad arguments or internal error.

hrconf getconfig

This command lists information about the controllers, logical drives, and physical drives. This information can include (but is not limited to) the following items:

- Controller type
- Logical drive status, RAID level, and size
- Physical drive type, device ID, and the presence of predictive failure analysis (PFA)
- Physical drive state

Syntax

```
hrconf getconfig <Controller#> [AD/LD/PD/AL]
```

Parameters

Controller# is the controller number.

The options are defined as follows:

- AD—Adapter information only
- LD—Logical drive information only
- PD—Physical device information only
- AL—All information (optional)

Example

```
hrconf getconfig 1 ad
```

hrconf getstatus

The GETSTATUS function displays the status of any background command that is currently running.

Syntax

```
hrconf getstatus <Controller#>
```

Parameters

Controller# is the controller number.

Example

```
hrconf getstatus 1
```

hrconf getversion

This command returns the version information for all controllers or the optionally specified controller.

Syntax

```
hrconf getversion
```

Parameters

Optional controller# prints the version information for the specified controller.

Example

```
hrconf getversion
```

Return values

SUCCESS: 0x00

hrconf rescan

This command enables the controller to detect the removal of any disk drives in the ready state and to detect the connection of any new disk drives to the controller. The command returns when the rescan is complete.

Syntax

```
hrconf rescan <Controller#>
```

Parameters

Controller# is the controller number.

Example

```
hrconf rescan 1
```

Return values

SUCCESS: 0x00 indicates the command completed successfully.

FAILURE: 0x01 indicates the command failed.

hrconf restore

This command restores the controller configuration by importing its configuration settings from a specified file and deletes the current configuration. The file must have been saved through the BACKUP command from a controller of the same type with the same number and type of physical drives and having the same channels and device IDs. A reboot is required for the configuration change to take effect.

Syntax

```
hrconf restore <Controller#> <Filename> [noprompt]
```

Parameters

Controller# is the controller number.

Filename is the name of the file from which to read the configuration.

noprompt: No prompt for confirmation.

Example

```
hrconf restore 1 c:\windows\hr2200 noprompt
```

Return values

SUCCESS: 0x00 indicates the command completed successfully.

FAILURE: 0x01 indicates the command failed.

BAD_PATH: 0xFF indicates that the path provided is incorrect.

FAILURE: 0xFE indicates an error occurred writing the configuration file.

hrconf setconfig

This command resets the configuration of the controller.

Syntax

```
hrconf setconfig <Controller#> DEFAULT [noprompt]
```

Parameters

Controller# is the controller number.

Default: Logical drives are deleted, disk drives are reset to the ready state, and any controller settings are reset to default values.

noprompt: No prompt for confirmation.

Example

```
hrconf setconfig 1 default noprompt
```

Return values

SUCCESS: 0x00 indicates the command completed successfully.

FAILURE: 0x01 indicates the command failed.

hrconf setstate

This command redefines the state of a physical device from its current state to the designated state.

Syntax

```
hrconf setstate <Controller#> <Channel#> <Device#> <State>
```

Parameters

Controller# is the controller number.

channel# is the channel number for the drive.

device# is the device number for the drive.

The following State options are available:

- HSP—Create a hot spare from a ready drive
- RDY—Remove a hot spare designation
- RBL—Rebuild drive

Example

```
hrconf setstate 1 0 1 HSP
```

```
hrconf setstate 1 0 2 RDY
```

```
hrconf setstate 1 0 2 RBL
```

Return values

SUCCESS: 0x00 indicates the command completed successfully.

FAILURE: 0x01 indicates the command failed.

hrconf task

This command performs a task on a logical drive.

Syntax

```
hrconf TASK START <controller#> LOGICALDRIVE <LogicalDrive#> <options>[noprompt]
```

```
hrconf TASK STOP <controller#> LOGICALDRIVE <LogicalDrive#>
```

Parameters

Logical drive# is the number of the logical drive in which the task is to be performed.

Options specifies the tasks to be started or performed. The following options are available:

- verify_fix (Verify with fix)
- verify
- clear

noprompt: No prompt for confirmation.

4 Understanding RAID

When you create arrays (or logical drives), you can assign a RAID level to protect your data. Each RAID level offers a unique combination of performance and redundancy. RAID levels also vary by the number of disk drives they support.

This section describes the RAID levels supported by the HP Embedded SATA RAID Controller and provides a basic overview of each level to help you select the best level of protection for your data storage needs.

RAID technology overview

RAID is the technology of grouping several physical drives in a computer into an array that you can define as one logical drive. Each logical drive appears to the operating system as a single drive. This grouping technique greatly enhances logical drive capacity and performance beyond the physical limitations of a single physical drive.

When you group multiple physical drives into a logical drive, the controller can transfer data in parallel from the multiple drives in the array. This parallel transfer yields data transfer rates that are many times higher than with non-arrayed drives, enabling the system to better meet the throughput (amount of data processed in a given amount of time) or productivity needs of the multiple-user network environment.

The ability to respond to multiple data requests provides not only an increase in throughput, but also a decrease in response time. The combination of parallel transfers and simultaneous responses to multiple requests enables disk arrays to provide a high level of performance in network environments.

Understanding drive segments

A drive segment is a disk drive or portion of a disk drive that is used to create an array. A disk drive can contain RAID segments (segments that are part of an array) and unused space.

Each segment can be part of only one logical device at a time. If a disk drive is not part of any logical device, the entire disk is an available segment.

Stripe-unit size

With RAID technology, data is striped across an array of physical drives. This data-distribution scheme complements the way the operating system requests data. The granularity at which data is stored on one drive of the array before subsequent data is stored on the next drive of the array is called the stripe-unit size.

You can set the stripe-unit size to 16, 32, or 64 KB. You can maximize the performance of the controller by setting the stripe-unit size to a value that is close to the size of the system I/O requests. For example, performance in transaction-based environments, which typically involve large blocks of data, might be optimal when the stripe-unit size is set to 32 or 64 KB. However, performance in file and print environments, which typically involve multiple small blocks of data, might be optimal when the stripe-unit size is set to 16 KB.

The collection of stripe units, from the first drive of the array to the last drive of the array, is called a stripe.

Selecting a RAID level and tuning performance

Disk arrays are used to improve performance and reliability. The amount of improvement depends on the application programs that you run on the server and the RAID levels that you assign to the logical drives.

Each RAID level provides different levels of fault tolerance (data redundancy), utilization of physical drive capacity, and read and write performance. In addition, the RAID levels differ in regard to the minimum and maximum number of physical drives that are supported.

When selecting a RAID level for your system, consider the following factors.

Table 4 Comparison of RAID level features

	RAID 0	RAID 1	RAID 10
Data redundancy	No	Yes	Yes
Amount of physical drive capacity used for data storage	100%	50%	50%
Read performance	Superior	Very high	Very high
Write performance	Superior	Very high	Very high
Built-in spare drive	No	No	No
Minimum number of drives	2	2	4
Maximum number of drives	4	2	4

Physical drive utilization, read performance, and write performance depend on the number of drives in the array. Generally, the more drives in the array, the better the performance.


RAID 0

A RAID 0 array uses from two to four disk drives and provides data striping, where data is distributed evenly across the disk drives in equal-sized sections.

RAID 0 arrays do not maintain redundant data, so they offer no data protection. However, compared to an equal-sized group of independent disks, a RAID 0 array provides improved I/O performance.

The size of a drive segment is the size of the smallest disk drive in the array. For instance, an array that consists of two 250-GB disk drives and two 400-GB disk drives has a drive segment size of 250 GB. This array thus has 1,000 GB (= 4 x 250 GB) of usable drive space, all of which can be used to store data when the array is used to build a RAID 0 logical drive. The array also has 300 GB of drive space that is unusable in the context of this particular collection of drives.


The following figure illustrates this example, with HD1-HD4 representing the disk drives, LD representing the logical drive, D1-D9 representing data fragments, and (ND) representing sections of the drive that are unused.



RAID 1

A RAID 1 array is built from two disk drives, where one disk drive is a mirror of the other (the same data is stored on each disk drive). Compared to independent disk drives, RAID 1 arrays provide improved performance, with twice the read rate and an equal write rate of single disks. However, capacity is only 50% of independent disk drives. If a RAID 1 array is built from different-sized disk drives, the size of the drive segment is the size of the smaller disk drive, as with RAID 0.

The following figure shows a RAID 1 array being built from two disk drives of dissimilar size.



RAID 10

A RAID 10 array is built from two equal-sized RAID 1 arrays. Data in a RAID 10 array is both striped and mirrored. Mirroring provides data protection, and striping improves performance.

Using as an example the same set of drives as in the RAID 0 discussion, the total usable drive space is still 1,000 GB if these drives are configured into a RAID 10 array. However, only 500 GB can be used to store data in this case. The remaining 500 GB is used to store the mirrored data that provides fault tolerance.

This example is illustrated in the following figure.

