Xserve G5

Using the Hardware RAID PCI Card

Instructions for using the software provided with the Hardware RAID PCI Card

Using the Hardware RAID PCI Card

You use command-line tools to configure and manage RAID operations with the Hardware RAID PCI Card.

With the optional Hardware RAID PCI Card installed in your Xserve G5 system, you can create a RAID array that uses the system's drive modules and takes advantage of the enhanced reliability and performance for data transfer and storage in the server.

Using the software supplied with the Hardware RAID PCI Card, you can:

- Create a RAID array using RAID 0, 1, or 5
- Designate an available drive module as a hot spare
- Rebuild a drive module
- Change the cache policy settings for a RAID array
- Check the status of the RAID card, drive modules, and RAID arrays
- Display and save a log of operations, showing time and date
- Turn the audible alarm on or off

About RAID Levels and the Hardware RAID PCI Card

The Hardware RAID PCI Card supports three different RAID levels—RAID 0, 1, and 5. These levels have different architecture and provide varying degrees of performance and fault tolerance. Understanding the differences in RAID levels will help you set up your RAID array to best meet your data performance and security needs.

RAID 0: Data Striping Without Fault Tolerance

Also known as striping, RAID level 0 is a performance-oriented mapping technique for disk sets. Uniform subsets of the array's logical volume, called stripes, are mapped in regular sequence to a set's drives. RAID 0 provides high I/O performance at low cost, but it offers no redundancy and so is not recommended for use with the Xserve G5 system.

RAID 0 is not recommended for mission-critical applications. It may be acceptable as a working environment or "scratch disk."

RAID 0 requires a minimum of two drives.

RAID 1: Data Mirroring

RAID level 1, mirroring, is a long-established, popular scheme because of its simplicity and high levels of reliability and availability. In the Xserve G5 system, RAID 1 uses two drives. Each drive stores identical data; thus, RAID 1 provides very high data reliability and improved performance for read-intensive applications. However, this RAID level has a high capacity cost because it retains a full copy of data on each drive.

Common applications of RAID 1 are accounting and payroll, or any application requiring high data availability. In a RAID 1 configuration, the capacity of the smallest drive is the maximum storage area.

RAID 5: Independent Data Disks With Distributed Parity

RAID 5 uses parity, a specialized protection method, to rebuild data should a problem occur with a drive module. Parity protects stored information without requiring data duplication. In a RAID 5 array, parity information is distributed across all three drives.

With a RAID 5 array, a failed drive module can be replaced and the data from that drive will be rebuilt using parity information. Although performance is degraded during rebuilding, data remains intact and available.

Logical Drives and Cache

Whichever RAID level you choose, you can create one or more logical drives. Each logical drive can have its own RAID level and cache policies.

For example, you might create one logical drive that is 50 gigabytes (GB) in size on which to install the OS, and use the remainder of the array as a second large logical drive for data storage.

Note: When you create logical drives, it is not a good idea to mix RAID levels on the same physical drive module. Using multiple logical drives to set up two or more RAID arrays of the same level is a good use of this feature.

You can also use a cache to improve performance.

See "Sample Configuration Commands" on page 6 for examples of commands to create RAID arrays.

Choosing a RAID Level

You can set up a variety of configurations using RAID level 0, 1, or 5. The following scenarios are examples of good use of the Xserve G5 system's three drive modules and the RAID card.

RAID 5 Using Three Drive Modules

This large RAID array provides efficient storage and excellent protection for data. If a drive fails, the process of rebuilding data on the replacement drive will take 2 to 3 hours and the array's performance will be degraded during that period.

RAID 5 Using Three Drive Modules and Two Logical Drives

This configuration offers the benefits of a large RAID 5 array with a second logical drive that can be used for the operating system or another separate purpose.

RAID 1 Using Two Drive Modules (and an Optional Hot Spare)

A RAID 1 configuration provides complete redundancy of data through mirroring. If the system has a third drive module, it can serve as a hot spare, which will automatically be used if either of the array's drives fails. The third drive cannot be used for other data-storage purposes if it serves as a hot spare, however.

Installing Mac OS X Server on a Hardware RAID Volume

Because configuring the Hardware RAID PCI Card erases the operating system on the server's startup drive, you must reinstall Mac OS X Server after you configure the card.

Here's an overview of the procedure for reinstalling Mac OS X Server on a volume created with the hardware RAID card.

Step 1: Start up the server from the Install Disc 1 that came with the system

Step 2: Configure the Hardware RAID PCI Card

Use one of the methods discussed in "Configuring the Hardware RAID PCI Card" on page 5.

Step 3: Format the volume created with the card using Disk Utility or the diskutil command-line tool

Step 4: Install Mac OS X Server on a new volume

See Chapter 5, "Updating or Installing Software on Xserve G5 Systems," in the *Xserve G5 User's Guide* for instructions.

Configuring the Hardware RAID PCI Card

The Hardware RAID PCI Card 's default configuration is JBOD mode (for "just a bunch of disks") in which each drive module is presented as an independent drive with no RAID configured. You must configure the card and drive modules to use its RAID features.

Important: Configuring the hardware RAID card erases Mac OS X Server on the system's startup drive. You must reinstall Mac OS X Server from the system discs after configuring the card. See "Installing Mac OS X Server on a Hardware RAID Volume" above for an overview of software installation and Chapter 5, "Updating or Installing Software on Xserve G5 Systems," in the *Xserve G5 User's Guide* for detailed instructions.

You can access and configure the RAID card in one of three ways:

- Use a remote computer on the network to connect to the server
- Connect a terminal to the serial port of the server
- Use the Terminal application on a server that has a monitor and keyboard connected to it

All of the methods use command-line tools, as described below.

To configure the RAID card using a remote computer:

1. Start up the Xserve G5 system from Mac OS X Server Install Disc 1.
2. Open the Terminal application on a remote computer.
3. Type the following command to log in to the remote server:

   ```
   ssh -l root <server IP address>
   ```

   Note: You will be prompted to enter the server password, which is the first eight digits of the server's serial number (from a label on the back of the server or on the server's interior). The password is case-sensitive. Also, the "ssh" command uses a lowercase "L," not the number one.
4. Type the desired configuration command.

   See "Sample Configuration Commands" on page 6.

To configure the RAID card using a terminal connected to the server's serial port:

1. Connect a terminal to the serial port on the back panel of the server.
2. Use the terminal to connect to the server with settings of 57.6 kilobytes per second, 8 bits, no parity.
3. Type the desired configuration command.

   See "Sample Configuration Commands" on page 6.

To configure the RAID card using the server with a monitor and keyboard connected to it:

Note: The server must have a video card to connect a monitor or an available keyboard-video-monitor switch with monitor and keyboard connected.

1. Start up the server using Mac OS X Install Disc 1.
2. Start the Terminal application from the Installer menu.
3. Type the desired configuration command.

   See "Sample Configuration Commands" on page 6.

Important: When entering most commands for RAID configuration, you must be the root user. To log in as the root user, you can type "ssh" at the beginning of a command (as in "To configure the RAID card using a remote computer:" on page 5) or type "sudo" as the first part of the command. You will also need to provide the password for the server to which you are logging in.

Sample Configuration Commands

Examples of a full configuration command for RAID levels 0, 1, and 5 are shown below.

Note: The command-line tool you use to configure and manage RAID arrays is named "megaraid"; its files are stored in these locations:

/usr/bin/megaraid (megaraid command tool)

/usr/share/man/man8/megaraid.8 (man page)

/System/Library/Extensions/megaraid.kext (kernel extension)

To create a RAID 5 array with a cache for writing data and reading ahead:

```
megaraid -create R5 -drive 0 1 2 -writecache enable -readahead on
     -iopolicy cached
```

To create a RAID 5 array with a logical drive of 50 GB:

```
megaraid -create R5 -drive 0 1 2 -size 51200
```

To create a RAID 1 array with a hot spare (in two steps):

```
megaraid -create R1 -drives 0 1
megaraid -spare 2 -create
```

To create 6 logical drives automatically:

```
megaraid -create auto -numld 6
```

To rebuild a RAID 5 or RAID 1 array with a replacement drive:

1. To verify the drive that failed, check the list of devices by typing the following command.

   ```
   megaraid -showdevices
   ```

2. Remove the failed drive module and insert a replacement.

   Rebuilding begins automatically.

3. If rebuilding does not begin, you can specify rebuilding by typing this command.

   ```
   megaraid -rebuild <drive number>
   ```

   where <drive number> is the number of the drive module you replaced

RAID Management Commands

You can use several commands to manage or check the status of the RAID card or an array and to set up a log of operations. Examples of these commands follow.

For a complete list of commands to configure and manage the Hardware RAID PCI Card, display the man page for the card by typing the following command:

```
man megaraid
```

To change logical drive characteristics:

```
megaraid -changepolicy ld [-writecache enable or disable] [-readahead
    on or off or adaptive] [-iopolicy direct or cached]
```

To rebuild a drive:

```
megaraid -rebuild <drive number> -start (or -stop)
```

To display details of drives:

```
megaraid -showdevices
```

To display details of the hardware RAID card configuration:

```
megaraid -showconfig
```

To display details of RAID array configuration:

```
megaraid -showconfig <n>
```

where <n> is the logical drive number (often 0)

To turn the audible alarm on or off:

```
megaraid -alarm -on (or -off)
```

Adding a File System to the Logical Drives in RAID Arrays

Once you've configured the hardware RAID card and created logical drives, you must put a file system on those drives. You can use the Disk Utility application or the command-line tool diskutil for this purpose. The examples below show several diskutil options.

To show all commands provided by diskutil:

```
man diskutil
```

To display a list of both logical and physical drives available to Mac OS X and their share points:

```
diskutil list
```

To create an HFS+ volume:

```
diskutil eraseDisk "Journaled HFS+" <volume name> <device mount
    point>
```

where, for example, the volume name might be "Macintosh HD" and the device mount point might be /dev/disk0

Default Settings for the Hardware RAID PCI Card

The following settings are the default values for the hardware RAID card:

- Stripe size is 64 kilobytes (k) if no other size is specified
- Write cache is off
- Readahead is off
- I/O policy is direct
- RAID level for easy configuration (using the "-auto" command) is RAID 1 if the system has two drive modules or RAID 5 if the system has three drive modules; a single logical drive of the array's entire capacity is available