



Xserve RAID

Technology Overview
January 2004



Contents

| | |
|----------------|---|
| Page 3 | Introduction |
| Page 4 | Product Overview Key Features High-Availability Design |
| Page 7 | RAID Basics RAID Techniques RAID Levels |
| Page 12 | Deploying Xserve RAID Xserve RAID Applications |
| Page 14 | Management and Monitoring Setting Up and Managing Arrays Monitoring Xserve RAID Systems Managing Volumes |
| Page 17 | System Architecture Drive Technology and Capacity Passive Midplane Data Path RAID Controllers Fibre Channel Interconnect High-Availability Features |
| Page 23 | Service, Support, and Training Options |
| Page 25 | Purchasing Information |
| Page 27 | Technical Specifications |

Introduction

What's New?

- Larger, 250GB drive modules for even more capacity—up to 3.5TB of storage per 14-drive system¹
- Industry-standard SFP (Small Form-factor Pluggable) connectors for additional deployment flexibility in optical or copper infrastructures
- Advanced capabilities such as LUN slicing and online expansion
- Certification of compatibility by leading storage infrastructure vendors
- Qualification for use in Linux, Windows, and mixed-platform environments



Last year, Apple introduced Xserve RAID, a high-availability, high-performance storage solution at a groundbreaking price. With massive capacity and industry-leading remote management capabilities in a compact 3U enclosure, Xserve RAID put powerful RAID features within easy reach—an ideal solution for near-line and disk-to-disk storage, as well as for protecting business- and mission-critical data.

The new Xserve RAID builds on the success of its predecessor with even more capacity, industry-standard SFP connectors, and advanced management functions. Apple has worked with leading storage infrastructure vendors to certify Xserve RAID for integration with existing Fibre Channel hardware and data management solutions. What's more, Xserve RAID is now qualified for use in Linux, Windows, and mixed-platform environments.²

At the heart of Xserve RAID is an innovative Apple-engineered architecture, with 14 high-performance drive channels, dual independent RAID controllers, and a dual 2Gb Fibre Channel host interface. Together they provide up to 3.5TB of storage¹ that can “grow as you go” and throughput of up to 400 megabytes per second.³ Robust monitoring and notification features and hot-swappable components keep your data online and available. And with intuitive tools for quick configuration of protected storage volumes, this revolutionary RAID solution delivers ease of use that could come only from Apple.

High-performance, high-availability storage has never been so affordable. With a low price, simplified administration, easy serviceability, and flexible deployment options, Xserve RAID is designed to protect your organization's bottom line, as well as your digital assets.

Product Overview



Xserve RAID

Xserve RAID combines leading-edge storage technologies for massive capacity, fast performance, and superior data protection.



Xserve G5

Xserve RAID works seamlessly with the new Xserve G5, Apple's high-density 1U rack-optimized server. Equipped with single or dual PowerPC G5 processors, Xserve packs phenomenal power and a rich feature set into an affordable, easy-to-deploy system.

Key Features

Xserve RAID is a cost-effective answer to the growing storage requirements of businesses and institutions everywhere. Its high-performance, high-availability features include:

Massive storage. Xserve RAID holds up to 14 hot-swappable Apple Drive Modules—now with 3.5TB of total storage—in a rack-optimized 3U enclosure. Each 250GB hard drive connects to a dedicated ATA drive channel, eliminating a traditional source of bottlenecks and maximizing the 400MB/s Fibre Channel host connection. By adding more Xserve RAID systems, you'll have virtually limitless expansion capabilities: A standard 42U rack can hold over 49TB of Xserve RAID storage.⁴

Advanced data protection. The high-availability architecture and dual independent RAID controllers support RAID levels 0, 1, 3, 5, and 0+1. Xserve RAID also supports hybrid RAID levels 10, 30, and 50 when used in conjunction with host-based software RAID.

High-availability design. To ensure availability of your critical data, Xserve RAID is designed for nonstop operation. Redundant hot-swappable power and cooling modules allow the system to keep functioning even if one module fails. All the active components are modular, making it easy to replace them in seconds—usually without any interruption of service—and with no tools required. In the event of a failed drive, a global hot spare provides automatic rebuilding of data, without administrator intervention.⁵

Fast data access. With sustained throughput of up to 336MB/s, the advanced Xserve RAID architecture delivers fast access to storage without compromising data integrity.⁶ Dual independent RAID controllers provide protected storage with unprecedented performance. In fact, Xserve RAID boasts a throughput at RAID level 5 that's fast enough to support real-time, 10-bit high-definition (HD) 1080i video editing.

Intuitive management and monitoring tools. Sophisticated remote management capabilities dramatically simplify setup and monitoring of RAID storage. The Java-based RAID Admin application can build RAID sets on the fly, allowing administrators to bring protected storage online instantly, without waiting for initialization. RAID Admin also provides continuous feedback on system activity and health. If a problem is detected, the remote monitoring software automatically sends notification via email or pager, so administrators can quickly identify problems and repair them without downtime or data loss.

Server-class support products. To minimize downtime, Apple offers a suite of enterprise-class support products, including onsite hardware repairs, advanced software support, and convenient spares kits.



Apple Fibre Channel PCI Card

Xserve RAID connects to a host Xserve or Power Mac system using the dual-port 2Gb Apple Fibre Channel PCI Card (sold separately) with throughput of up to 400MB/s.³ The Fibre Channel interconnect technology supports multiple application environments using point-to-point, loop, and fabric technologies.

High-Availability Design

The 3U Xserve RAID enclosure is built for reliability, availability, and serviceability with high-quality, Apple-engineered construction. LEDs on the front panel provide continuous visual feedback on system status and activity levels; and 14 drive bays support hot-swapping of drive modules. The back panel features easy access to redundant RAID controllers and field-replaceable power, cooling, and battery modules, as well as connectivity to one or more host computers and up to two uninterruptible power supply (UPS) devices.

Front View

The diagram shows the front panel of the Xserve RAID enclosure. It features two rows of seven drive bays each. Above the drive bays, there are several indicator lights and buttons. Red numbered callouts (1-6) point to specific features: 1 points to the drive modules, 2 to a lock, 3 to unit identifier buttons, 4 to an alarm silencer button, 5 to system status indicators, and 6 to system activity indicators.

- ① **Apple Drive Modules.** Xserve RAID holds up to 14 hot-swappable Apple Drive Modules. Each 250GB drive is on an independent channel to maximize data throughput and increase system reliability. LEDs on the front of each module indicate drive health and activity.
- ② **System lock.** A built-in lock secures the drive modules in the system. The remote monitoring application shows the status of enclosure security and can notify the administrator if someone unlocks it.
- ③ **Unit identifier.** Duplicate system identifier buttons on the front and back of the system can be turned on manually or by using the remote monitoring software, making it easy to locate a particular Xserve RAID in a rack with multiple devices. The buttons also illuminate if a system event occurs.
- ④ **Alarm silencer.** In the event of a component failure, the alarm system is triggered; with a touch of this button, the alarm is silenced. The remote monitoring software provides detailed event information about the affected system and notifies the administrator via email or pager.
- ⑤ **System status indicators.** Indicator lights display status for power supplies, cooling modules, RAID controllers, and system temperature.
- ⑥ **System activity indicators.** Forty-six blue LEDs provide at-a-glance activity levels for each host channel, and Fibre Channel indicators provide link status information.



Convenient replacement modules

Problem resolution is fast with AppleCare Service Parts Kits for Xserve RAID (sold separately). Each kit has an Apple Drive Module, power supply module, cooling module, and RAID controller module.

Back View

- 1 RAID controller modules.** Two independent storage processor units manage RAID functions, data transfers, and failure protection for each set of seven drives. The environment management coprocessor in each controller simplifies configuration and management of RAID sets.
- 2 Serial ports.** Standard DB-9 serial ports allow connection to UPS units for protection from brownout or over-voltage conditions.
- 3 Ethernet ports.** The 10/100BASE-T Ethernet interfaces allow you to manage, monitor, and diagnose Xserve RAID systems over TCP/IP.
- 4 Fibre Channel ports.** Each RAID processor connects to the host system via a 2Gb Fibre Channel interface with throughput of up to 200MB/s per port.³ Throughput is guaranteed, which means bandwidth remains constant, even as more devices are added in a fabric configuration. Industry-standard SFP connectors support both optical and copper infrastructures.
- 5 Optional batteries.** Cache Backup Battery Modules can provide more than 72 hours of backup power to protect the integrity of data in the RAID controller cache during a power outage.
- 6 Power supply modules.** Either of the redundant, load-sharing power supplies can power Xserve RAID should the other one fail. A failed power supply can be replaced in seconds without tools and without shutting down the system.
- 7 Cooling modules.** Redundant, hot-swappable cooling modules provide automatic front-to-rear cooling for rack environments.

RAID Basics

Redundant array of independent disks, or RAID, is a grouping of multiple physical hard drives into an array, or RAID set, that appears to the host computer as a single logical storage unit. Use of RAID technology enables organizations to increase storage capacity while achieving levels of performance, reliability, and data protection not possible from a single hard drive. All drives in the array can operate simultaneously for dramatically faster overall throughput. RAID systems can also improve data availability and fault tolerance, because redundant data can be stored on multiple physical drives. Even in the event of a drive failure, the system can continue to operate without loss of data and with no interruption in service.

RAID Techniques

Software versus hardware RAID

RAID functionality can be implemented in hardware or software. Software RAID—available in Mac OS X and Mac OS X Server—enables drive striping and drive mirroring, whether for basic RAID 0 or RAID 1 configurations or for robust configurations using hybrid RAID levels 10, 30, or 50. Software RAID is not recommended for parity RAID levels, such as RAID 3 and 5, which involve compute-intensive calculations that can overtax the system processor.

RAID technology is based on three practices: *Striping* to improve storage performance, and *mirroring* and *parity* to provide redundancy for increased data protection. Most RAID configurations, or RAID levels, combine these techniques to provide a balance of data protection and performance.

Striping

Increased performance in a set of multiple drives is accomplished by data striping. As the name implies, striping divides a logical drive into data blocks, or stripes, that are distributed across the array of physical drives. Data is then laid down according to the stripe paths, so that each file is spread across multiple drives. Striping a set of disks greatly improves overall storage performance because the drives operate in parallel. While one drive is writing or reading a data block, another is seeking the next block. Striping alone, known as RAID level 0, offers no data protection.

Mirroring

The simplest method of achieving data redundancy, mirroring involves writing identical copies of all data to a pair of physical drives. This results in very high data reliability: If one drive fails, the data is still available on the remaining disk drive. However, it also results in storage efficiency of only 50 percent, because two physical drives are required to create a protected storage volume with the equivalent capacity of a single drive. Mirroring alone is known as RAID level 1.

Parity

A more sophisticated method of creating redundancy, parity provides data protection for an array of drives without requiring complete duplication of the drive contents. Parity information can be used—along with the data on the surviving drives—to reconstruct the contents of a failed drive. The parity data can be stored on a dedicated drive, as in RAID 3, or distributed across an array of drives, as in RAID 5. In either case, parity provides much greater storage efficiency than mirroring—up to 85 percent for a set of seven drives.

RAID Levels

Each RAID level offers a unique balance of I/O performance, data protection, and storage efficiency. Xserve RAID supports all popular RAID levels, so you can select the best configuration for your application and your budget.

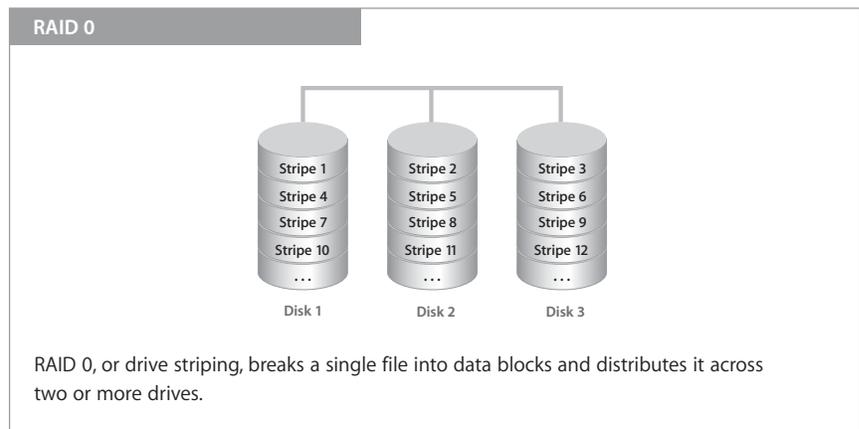
Comparison of RAID levels

| RAID level | Minimum number of drives required | Storage efficiency | Read performance | Write performance | Data redundancy |
|------------|-----------------------------------|--------------------|------------------|-------------------|-----------------|
| RAID 0 | 2 | Highest | Very high | Highest | No |
| RAID 1 | 2 | Low | High | Medium | Yes |
| RAID 3 | 3 | High to very high | Medium | Medium | Yes |
| RAID 5 | 3 | High to very high | High | High | Yes |
| RAID 0+1 | 4 | Low | High | High | Yes |
| RAID 10 | 4 | Low | High | High | Yes |
| RAID 30 | 6 | High to very high | High | High | Yes |
| RAID 50 | 6 | High to very high | Highest | Very high | Yes |

Each RAID level offers a unique balance of performance, data protection, and drive efficiency. The efficiency of parity RAID levels increases with the number of drives in an array.

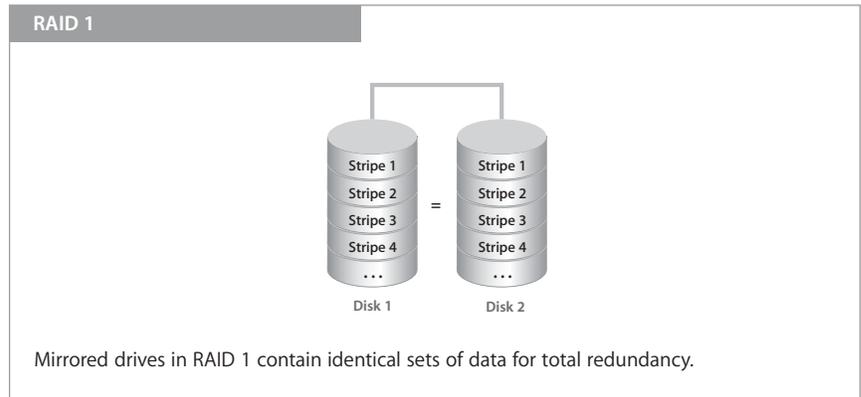
RAID 0: Striping

RAID 0 distributes data evenly in horizontal stripes across an array of drives. While RAID 0 offers substantial speed enhancements, it provides no data protection: If one drive fails, all of its data is lost and all drives must be reformatted. RAID 0 provides the most efficient use of drive capacity, because no storage space is dedicated to redundant data or parity information.



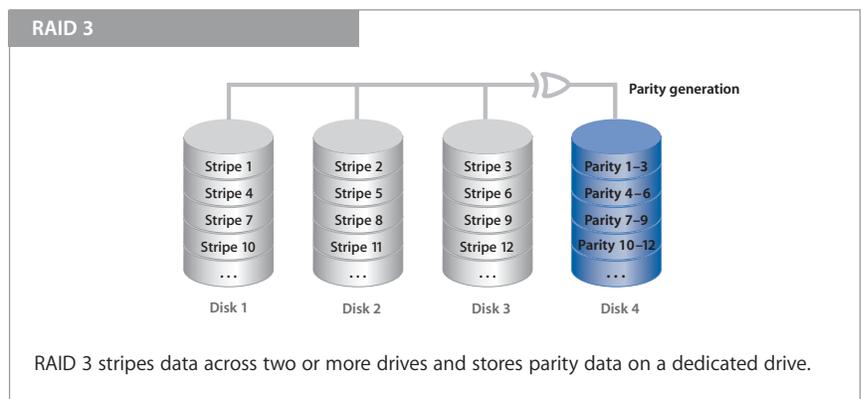
RAID 1: Mirroring

RAID 1 creates a pair of mirrored drives with exactly the same data. It provides a high level of data availability and, by using both drives simultaneously, offers fast read performance. Write speeds are lower, however, because all data is written twice—once to each drive. RAID 1 is also relatively expensive to deploy because it offers the poorest level of drive efficiency.



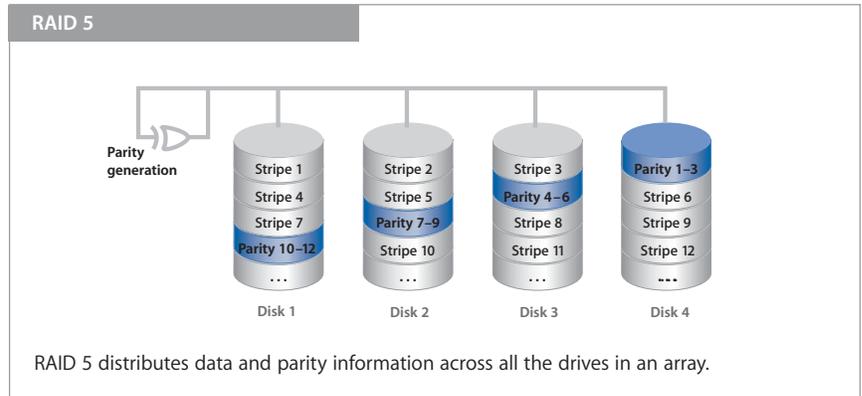
RAID 3: Striping with parity

RAID 3 stripes data across two or more drives and stores parity data on a dedicated drive. In the event of a disk failure, the redundant parity bits can be used to reconstruct data on any drive. RAID 3 offers fast read rates and high data availability, with greater storage efficiency than RAID 1. Write performance in RAID 3 is faster than in RAID 1, but not as fast as in RAID 5, because it depends on the availability of a separate parity drive. RAID 3 requires a minimum of three drives; using more drives results in faster performance and greater drive efficiency.



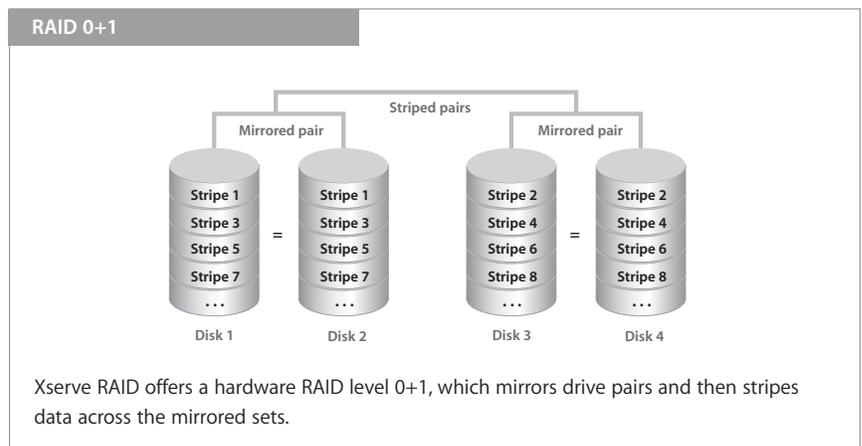
RAID 5: Striping with distributed parity

RAID 5 distributes both data and parity information across an array of drives one block at a time, with each drive operating independently. This enables maximum read performance when accessing large files and improves performance in a transaction-processing environment. Write performance is also improved because parity information is striped across the drives, removing the bottleneck of a single parity drive. RAID 5 is the most popular configuration for high-throughput protected storage. Like RAID 3, RAID 5 requires a minimum of three drives; using more drives provides faster performance and higher storage efficiency.



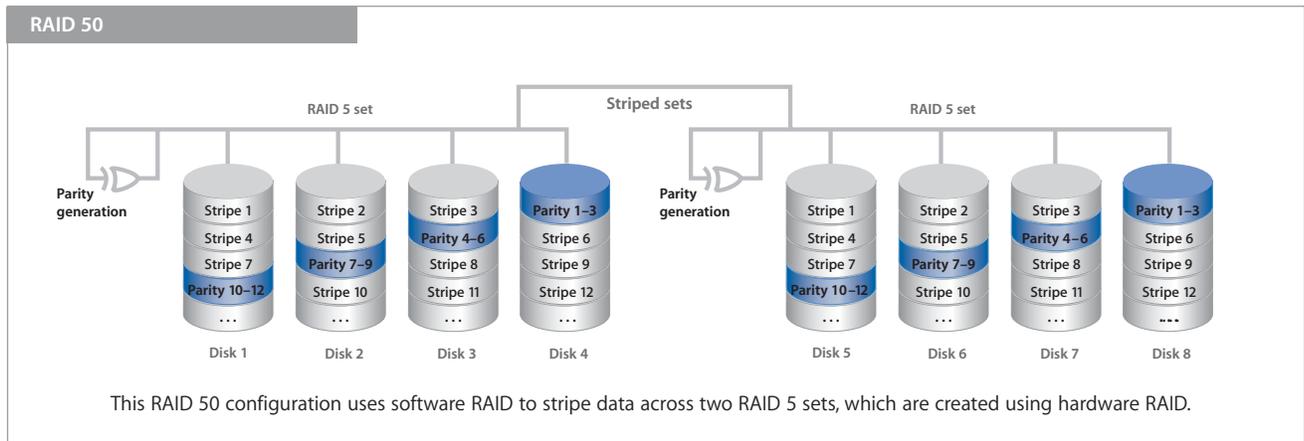
RAID 0+1: Striping over mirroring

This hybrid RAID level is created by striping data across multiple pairs of mirrored drives. With Xserve RAID, both striping and mirroring are set up in hardware, which means the array can run independently of the server processor. RAID 0+1 provides higher throughput with simpler setup than protected configurations that use software RAID for striping, such as RAID 10.



RAID 10, 30, and 50

Hybrid RAID levels 10, 30, and 50 use hardware RAID to create two or more sets in RAID 1, 3, or 5, and software RAID to stripe across the sets. This creates a single data volume with the best balance of performance and data protection. In the case of Xserve RAID, you can combine all 14 drives for high performance at maximum capacity.



Deploying Xserve RAID



Xserve RAID: A unique blend of features

A robust feature set and highly flexible deployment options make it possible to use Xserve RAID in a wide range of environments, including:

- Mission-critical data storage deployments, where zero downtime is paramount at virtually any cost
- Business-critical data storage deployments, which require high capacity and high availability at a reasonable cost
- Near-line archive deployments, which seek a balance of high capacity, scalability, and reasonable throughput at the lowest possible cost
- Rich media storage deployments, which require superfast throughput, data protection, large capacity, and a reasonable cost

Exponential growth in the creation and distribution of digital content is driving demand for high-capacity storage solutions. Large databases, digital video footage, immense scientific data sets, and expanding archives of financial information and employee records require terabytes of storage. At the same time, data protection and near-instant data access are crucial to many organizations. In response to these pressing requirements, customers have been forced to make tradeoffs among data protection, performance, and capacity—or spend hundreds of thousands of dollars on high-end storage solutions that combine these features.

Xserve RAID is a revolutionary storage product that eliminates the need for these tradeoffs. With massive capacity and high-availability features previously available only in much more expensive storage systems, Xserve RAID offers unmatched versatility at an unprecedented price. What's more, it meets the most demanding performance requirements, delivering a sustained throughput of up to 336MB/s—the highest of any RAID system in its class.⁶ These robust capabilities make Xserve RAID ideally suited for protecting mission- and business-critical data, yet it's affordable enough for near-line storage deployments and fast enough for media production environments.

Comparison of storage solutions

| | Apple Xserve RAID | Dell EMC CX200 | IBM FASt200 3542-1R | HP StorageWorks 1000 | Sun StorEdge 6120 |
|--------------|----------------------|-------------------|---------------------------|----------------------------|-------------------------|
| Capacity | 3.5TB | 2.1TB | 2.1TB | 2.1TB | 2.04TB |
| Size | 3U | 3U | 3U | 3U | Two 3U |
| Price* | \$10,999 | \$18,999 | \$51,895 | \$23,925 | \$74,600 |
| Price per GB | \$3.14 | \$9.05 | \$24.71 | \$11.39 | \$36.57 |

* Based on prices published on resellers' websites on December 10, 2003.

Compared with leading storage products, Xserve RAID offers the lowest cost per gigabyte. At the same time, it packs more gigabytes in less space—up to 3.5TB in a 3U enclosure.¹

Xserve RAID Applications

Whether in all-Apple or heterogeneous environments, Xserve RAID can fit into existing storage networks or provide a foundation for new deployments. The industry-standard SFP-based Fibre Channel interface supports point-to-point, loop, and fabric topologies for integrating Xserve RAID into Fibre Channel storage infrastructures. In addition, a platform-independent design and Java-based administrative tools make setup and monitoring easy from virtually any Internet-connected computer. Now certified by leading storage network manufacturers and qualified to work with popular operating systems, Xserve RAID is a flexible solution for a full range of storage applications.

Third-party certifications

Leading storage infrastructure vendors have certified Xserve RAID for integration with existing Fibre Channel hardware and data management solutions, including:

- Veritas
- QLogic
- Brocade
- Chaparral Network Storage
- ATTO
- Emulex
- Candera

In addition to Mac OS X and Mac OS X Server, Xserve RAID has been qualified for use on these operating systems:

- Windows Server 2003
- Red Hat Enterprise Linux v2.1 and v3
- Yellow Dog Linux v3

Primary storage for Mac, Windows, and Linux servers

With high-availability features such as redundant power and cooling, protected RAID storage, global drive hot sparing, and hot-swap components, Xserve RAID delivers fast, reliable data access—meeting the requirements of the most demanding business-critical and mission-critical storage deployments. It also provides the capacity and deployment flexibility required for High Performance Computing (HPC) environments. At the same time, Xserve RAID offers industry-leading capacity, up to 3.5TB,¹ and can support dozens of servers using advanced LUN slicing and mapping capabilities. Best of all, this outstanding feature set is available at a cost far below that of competing primary storage products.

Network Attached Storage (NAS)

Xserve RAID teams with Xserve G5, Apple's easy-to-deploy 1U server, to provide an affordable alternative to traditional Network Attached Storage devices. Offering the compatibility of open standards and Apple's legendary manageability, this high-performance, high-capacity solution allows Mac, Windows, and Linux clients to share terabytes of information—with no per-user licensing fees. At the same time, Xserve RAID has the flexibility to fit easily into existing mixed-platform Storage Area Networks (SANs).

Near-line and disk-to-disk archive

Xserve RAID offers outstanding value for near-line applications, while providing the ready data access required by many governmental regulations for email and HR records. Fibre Channel connectivity provides simple and flexible attachment to existing Fibre Channel networks and SAN infrastructures. Near-line storage users get the benefit of low-cost, high-capacity storage, without sacrificing high availability and data protection. Xserve RAID also offers superior throughput and scalability compared with other systems typically used for near-line storage. With Xserve RAID, organizations can reduce their overall storage costs, as well as their dependence on inflexible, proprietary systems.

Rich media storage

With a video-optimized architecture and dual independent 2Gb Fibre Channel interface, Xserve RAID is the perfect complement to a Power Mac G5 video editing workstation. It provides the fast, consistent performance—up to 336MB/s sustained throughput—required for video post-production. In fact, Xserve RAID delivers enough bandwidth for working in high-definition (HD) video or supporting multiple streams of uncompressed standard-definition (SD) video with real-time effects. In addition, with up to 3.5TB of online storage, Xserve RAID makes it easy to manage multiple video tasks. Even when used for editing in uncompressed HD 1080i, this robust storage system can hold nearly 16 hours of footage. What's more, Xserve RAID scales in both throughput and capacity: Users need only purchase the performance level and capacity they require today, with the assurance that their investment will be protected as application and storage needs grow.

Management and Monitoring

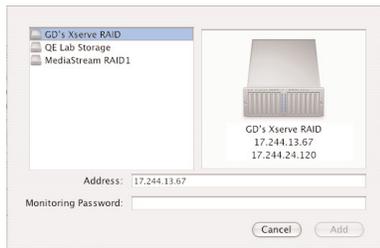


Remote management using RAID Admin

The Java-based RAID Admin application works over TCP/IP, making it easy to set up, manage, and monitor Xserve RAID systems from virtually any Internet-connected computer.

Xserve RAID comes with RAID Admin, a powerful remote management application that dramatically simplifies setup and monitoring of multiple Xserve RAID systems. This Java-based tool provides an intuitive interface for creating protected storage volumes, managing preferences, and monitoring storage hardware. To protect storage systems from unauthorized access, RAID Admin uses robust, multilevel password authentication.

RAID Admin integrates with hardware to continuously monitor system activity and status. The environment management coprocessors in the RAID controllers connect to the host system via Ethernet, providing dedicated, out-of-band system management capabilities. This ensures optimum system performance and reliability, because there is no contention for bandwidth with the RAID processors or the Fibre Channel interface.

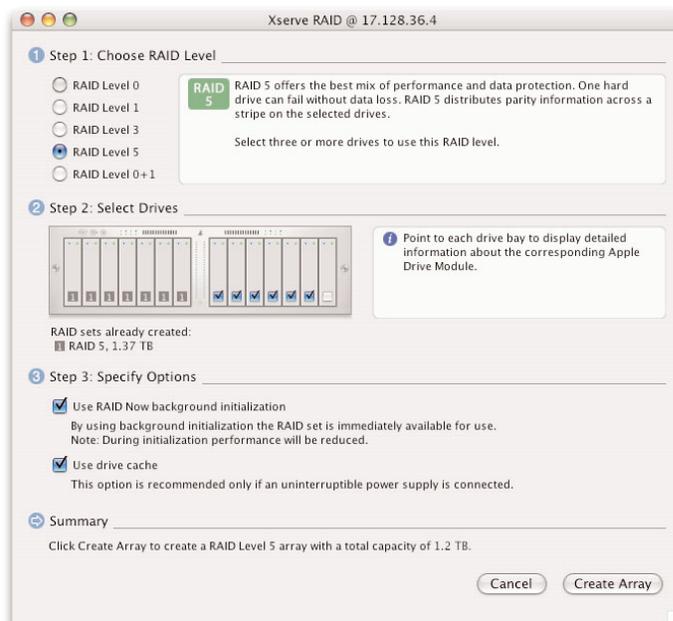


Automatic discovery with Rendezvous

Using Apple's Rendezvous technology, RAID Admin automatically discovers Xserve RAID systems on your network. This makes it easy to add them to your management window, without needing to enter each IP address manually.

Setting Up and Managing Arrays

RAID Admin lets you create arrays in three easy steps. First, choose a RAID level for information about the benefits and limitations of that level. Next, click the drives you want to assign to the array. RAID Admin even gives you information about each of the drives in your Xserve RAID system, so there's no guesswork. Any remaining drives are assigned as global hot spares for each controller. Last, choose background initialization to make the newly created array available immediately.



RAID Admin provides an easy-to-use interface for creating RAID sets on the fly.

Expanding arrays

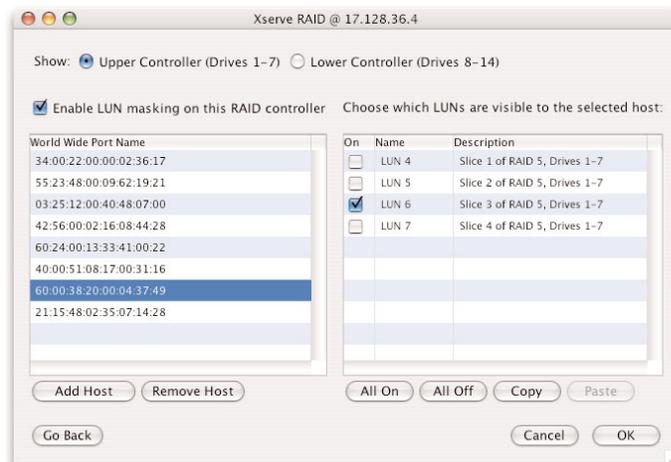
Once an array is established, you can easily add more disks using RAID Admin. Expanding an array preserves the existing array, while making available the capacity of the added drives. Expansion in RAID Admin works “on the fly,” so the expanded capacity is available within minutes.

Slicing arrays

To share a single RAID set across multiple servers, RAID Admin allows you to divide it into smaller segments, or “slices.” Each slice becomes a separate LUN, or logical unit number, so the host system can manage it as a discrete volume. Instead of creating many small arrays for individual servers, this “pooling and distributing” technique enables administrators to leverage the storage efficiencies of a large consolidated array. RAID Admin can slice up to eight LUNs per RAID controller, or up to 16 LUNs per Xserve RAID system.

LUN mapping

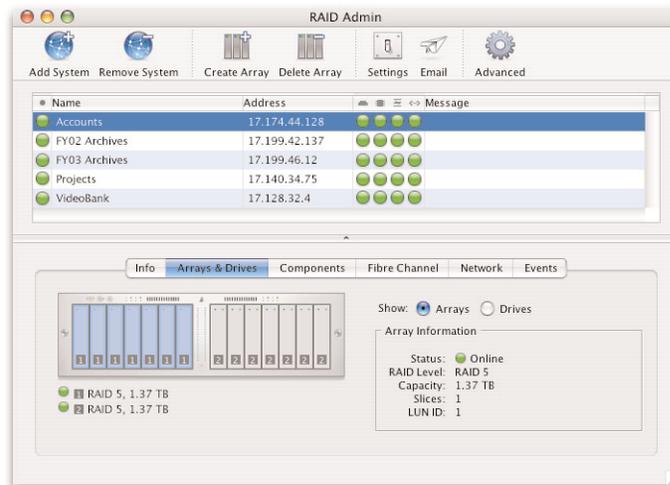
RAID Admin offers advanced Fibre Channel networking features such as LUN mapping and masking. When you connect an Xserve RAID to more than one host system on a Fibre Channel network, you can “map” each LUN on the RAID system to a single host and “mask” it from any other host. By masking arrays and slices, you avoid the possibility that more than one host will write data to an array, eliminating the risk of array corruption or data loss.



An intuitive interface makes it easy to set LUN mapping, as well as other storage options.

Monitoring Xserve RAID Systems

Xserve RAID features robust self-monitoring technology that maximizes storage availability while minimizing the pressure on system administrators. RAID Admin software reports at-a-glance information about all Xserve RAID systems on the network, with green, yellow, or red icons indicating the health of each subsystem. Administrators can easily “drill down” for detailed information about status and performance of each component. If operating conditions for any subsystem exceed predefined thresholds, RAID Admin can automatically notify the administrators via email, pager, or email-capable cell phone—allowing them to stay in touch with storage deployments and resolve issues without downtime or data loss.



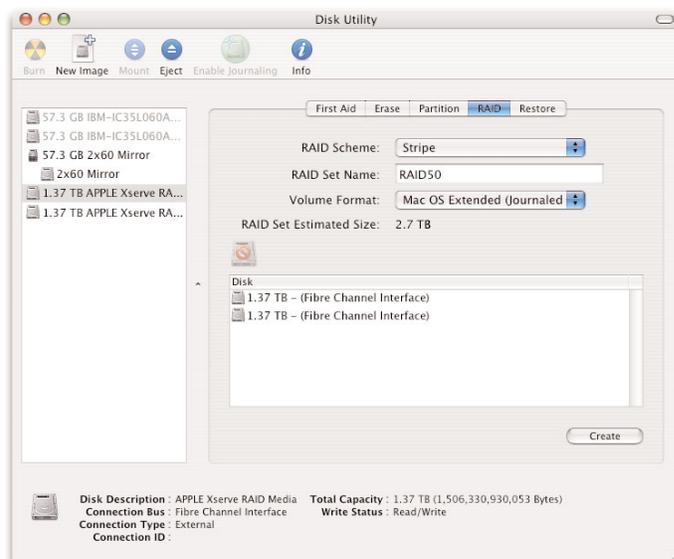
RAID Admin software displays summary health and activity information for multiple Xserve RAID systems.

Managing Volumes

File system journaling in Mac OS X Server

A robust file system journaling feature in Mac OS X and Mac OS X Server enhances the availability and fault resilience of servers and server-attached storage devices. Journaling protects the integrity of the Mac OS Extended (HFS+) file system in the event of an unplanned shutdown or power failure and maximizes uptime by expediting repairs to the affected volumes when the system restarts.

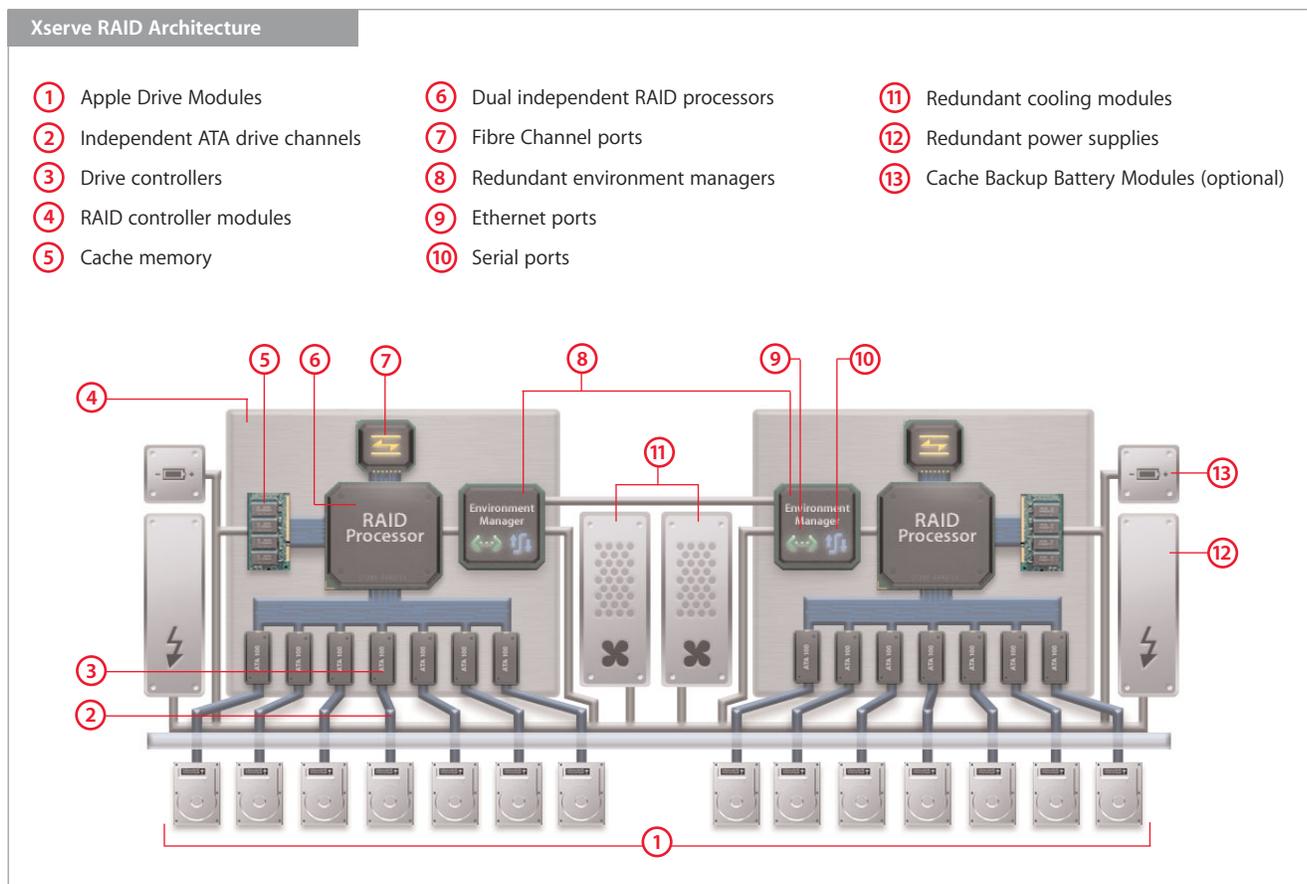
Administrators can use host-based volume management tools to mount storage volumes on host systems. In Mac OS X or Mac OS X Server, they can use Disk Utility to initialize arrays in Mac OS Extended or UNIX File System formats; turn on file system journaling; and set up software RAID arrays.



Disk Utility makes it easy to use software RAID to create hybrid RAID levels 10, 30, and 50.

System Architecture

Fast throughput and data availability are paramount for dedicated storage systems. That's why Xserve RAID combines leading-edge Fibre Channel, RAID, and ATA technologies in an innovative architecture designed for superior performance and reliability. Redundant, hot-swappable components and a passive midplane data path protect your digital assets from single points of failure, while independent controllers and 14 drive channels provide high-speed throughput for delivering data quickly to the host system.



Drive Technology and Capacity

Instead of relying on expensive SCSI or Fibre Channel hard drives, Apple developed a high-throughput ATA to Fibre Channel storage architecture that delivers superior performance and reliability at a much lower cost. Xserve RAID achieves its massive 3.5TB capacity with 14 hot-swappable 250GB Apple Drive Modules. The high-capacity 7200-rpm ATA/100 drive modules provide affordable expansion and configuration flexibility as storage needs grow.



Apple Drive Modules

Xserve RAID uses reliable, high-performance Apple Drive Modules with true hot-swapping capabilities.

Fourteen independent drive channels

The multithreaded Xserve RAID architecture features a dedicated ATA/100 channel for each of its 14 drive bays. Drive independence prevents data bottlenecks and maximizes throughput while enhancing availability. Because each hard drive is isolated on its own bus, a drive failure doesn't degrade the accessibility or performance of the surviving drives. In addition, independent drive channels reduce the complexity and cost of high-availability storage, since loop redundancy circuits and signal amplifiers aren't required, as in multidrive Fibre Channel and SCSI implementations.

Apple Drive Modules

Hot-plug Apple Drive Modules feature a unique handle design for fast, easy installation and a positive locking mechanism to hold them tightly in place after insertion. Sophisticated connectors called SCA II, or Single Connector Attachment II, protect the connecting pins from bending when a drive is inserted or removed. Carefully tested and qualified to ensure maximum performance and reliability, Apple Drive Modules work seamlessly with RAID Admin, Apple's innovative remote management and monitoring software for Xserve RAID.

On-drive cache

Each Apple Drive Module is equipped with 8MB of disk cache to accelerate read and write operations in performance-sensitive applications such as video editing. On-drive read cache, commonly called "read-ahead" cache, allows the drive controller to retrieve a block of data before it is required by the host system. The RAID controller provides instructions to each drive in an array, indicating what data block is required next and allowing the individual drive controllers to queue up the data for instant availability.

Similarly, on-drive cache can serve as a holding tank for write data. When this option is enabled, the RAID controller distributes data to each drive and stores it in cache—allowing the drive controller to notify the RAID processor immediately that transmission was successful, also known as a "committed write." This frees up the storage pipeline for other needs while the data is being written to the drive platters at maximum speed.

Using on-drive write cache is inherently risky: In the event of a power failure, the cached data could be lost. For this reason, it is turned off by default in protected RAID configurations.

Hot-swapping capabilities

The hardware RAID controllers and Apple-designed drive carriers enable true hot-swapping of Apple Drive Modules. This means a failed drive can be removed and replaced without interruption in service or loss of data. The system continues to operate while the contents of the failed drive are rebuilt on a replacement drive using redundant or parity information. (This feature is not applicable to a RAID 0 configuration, since it doesn't offer data protection.)

Monitoring of drive health

Xserve RAID hardware and remote management software work together to provide industry-leading monitoring and alerting capabilities. The RAID controller automatically reads Self-Monitoring, Analysis, and Reporting Technology (SMART) data, which allows each hard drive to report its health. This enables the operating system to warn the administrator of a prefailure condition, providing the opportunity to replace the hard drive before the failure occurs. Each drive module has two LEDs for local monitoring of drive activity and health.

Passive Midplane Data Path

The Xserve RAID architecture is designed to avoid vulnerability to a single point of failure. This means that the failure of any single component can't result in a system-wide failure and loss of data availability. For highest reliability, Apple built Xserve RAID around a midplane that serves as a passive data path. The midplane is the central connector between the drives, RAID controllers, power supplies, and cooling modules. Most RAID systems depend on the midplane to relay data and instruction sets between drives, and a failure in the midplane can impair data availability. In Xserve RAID, all data passes through the independent drive channels, which are simply held in place by the midplane. This design improves system reliability and protects the availability of stored data.

RAID Controllers

Xserve RAID features two controllers that independently manage storage operations for a set of up to seven drives. Dual independent controllers allow for simultaneous storage processing. Because each controller needs to manage only seven drives, Xserve RAID delivers outstanding performance that scales as capacity increases. These controller modules offer advanced availability and manageability features, including the environment manager, which manages RAID functions and monitors status and activity of system components.

RAID processor

Each Xserve RAID controller has a powerful processor that manages all the data transfers between the hard drives and host computer. It performs the complex parity calculations required for RAID levels 3 and 5, generating parity data on the fly while managing write operations to each drive in the set. In the event of a drive failure, the RAID processor uses redundant data to rebuild its contents on a spare drive.⁵ Xserve RAID supports RAID levels 0, 1, 3, 5, and 0+1 using the RAID processor, as well as hybrid RAID levels 10, 30, and 50 by combining hardware RAID capabilities with host-based software RAID.

An innovative process ensures that RAID sets are error-free and tuned for maximum throughput. When creating a RAID set, Xserve RAID verifies the integrity of the array by checking every block on every disk for potential errors. If a bad block is discovered, it simply chooses another block that's in close proximity to the original—eliminating unnecessary head movement that can result in lower performance.

To maintain optimal throughput, the RAID processor uses another special technique that retrieves data on the first revolution of the RAID set. This avoids multiple retries (an activity that can hamper performance on RAID systems), while allowing Xserve RAID to maintain a nearly constant sustained throughput over the entire platter of each hard disk. Outer and inner tracks perform comparably, providing consistent, fast delivery of data to the host system. This is particularly critical in high-bandwidth applications such as HD video editing, which requires fast storage throughput for clean video capture and smooth playback.

Controller cache memory

Each RAID controller is equipped with 128MB or 512MB of cache memory for enhanced storage throughput. For optimal read performance, streaming data can be queued into cache, and sophisticated algorithms measure data traffic patterns and re-allocate read cache dynamically. Depending on write performance and data protection requirements, cache memory can be configured as “write-back” or “write-through.” In either case, the write cache configuration does not affect read cache settings.

Write-back cache. For performance-critical applications, cache memory can accelerate write speeds using a configuration called write-back cache. In this mode, data is considered committed, or successfully received, as soon as the RAID controller writes back to the host that the information is stored in cache memory. However, as with drive cache, write-back cache is vulnerable to data loss. Since cached data is stored in SDRAM, which requires a continuous supply of power, a power outage or major system failure could result in data loss. This risk can be minimized by connecting a UPS system or using the optional Cache Backup Battery Modules. With a minimum of 72 hours of power, the battery modules preserve the contents of cache memory until system power is restored.

Write-through cache. When write-through cache is enabled, the RAID controller writes data through the cache, directly to the disks, before informing the host that the write is committed. Although this may reduce performance, write-through cache is the recommended setting for protected RAID without a battery or UPS device. For high-bandwidth applications or when multiple servers are attached to Xserve RAID, write-back cache is recommended. If a power outage occurs while using write-back cache with a UPS system, Xserve RAID automatically reverts to write-through cache to protect the data in case of a UPS failure.

Environment managers

Each Xserve RAID controller module contains a special coprocessor that monitors the enclosure environment and manages configuration of the main processor’s RAID functions. This innovative environment manager simplifies management and boosts reliability by ensuring continuous feedback on the status and activity levels of system components.

The environment managers integrate seamlessly with RAID Admin software and the RAID processor, providing advanced capabilities such as RAID Now for background initialization during RAID creation. Each environment manager operates on an independent bus and communicates with the host system via Ethernet. All the management and monitoring processes are handled “out of band,” which means Fibre Channel bandwidth is not specified for management chores.

Each environment manager communicates with both RAID controllers and all hard drives. If one of the RAID environment managers fails, the other takes over and continues the monitoring and management functions to avoid any interruption in service. At the same time, it immediately notifies the administrator via email of the component failure.

Because the RAID environment managers connect to the host computer over Ethernet using standard Internet protocols, it’s easy to connect to and monitor Xserve RAID systems from any computer on the Internet. In addition to an Ethernet port, each environment manager has a serial port for connecting an optional UPS device, providing an extra measure of protection against power surges or electrical outages.



Apple Fibre Channel PCI Card

Xserve RAID connects to an Xserve or Power Mac host system using the dual-port 2Gb Apple Fibre Channel PCI Card (sold separately). Two 2.9-meter copper cables are included with the card for connecting to the two RAID controllers in the Xserve RAID system.

Fibre Channel Interconnect

Xserve RAID features dual independent 2Gb Fibre Channel ports for fast data transport between the storage system and the host computer. This industry-standard technology provides superior bandwidth, availability, and deployment flexibility over SCSI technology. To ensure interoperability, Xserve RAID uses the latest industry-standard SCSI 3 protocol for data transfer. In addition, SFP (Small Form-factor Pluggable) connectors support low-cost copper cables, long-distance optical cables, or copper cables with embedded SFP transceivers and work with industry-standard Fibre Channel switches.

400MB/s throughput

Each 2Gb Fibre Channel port offers 200MB/s bandwidth for a total throughput of up to 400MB/s.³ Fibre Channel is the only storage connectivity technology that provides guaranteed bandwidth, so the host computer receives data at the same high speed as the RAID system sends it.

Superior scalability

Unlike with Gigabit Ethernet or SCSI technologies, adding storage units to a Fibre Channel network doesn't lead to congestive data loss or reduced overall bandwidth. The highly scalable Fibre Channel technology can address up to 126 devices per loop, with each device experiencing a full 200MB/s throughput. In contrast, SCSI allows for a maximum of only 15 devices per channel, resulting in added expense and complexity to deploy large RAID arrays.

Long-distance connectivity

Xserve RAID connects to a 2Gb Fibre Channel PCI Card in the host computer via a thin, flexible copper or fiber-optic cable. The Fibre Channel interface also provides support for much longer distances between the storage units and the host system than is possible using SCSI cables. The Apple Fibre Channel PCI Card accepts optical transceivers allowing it to use fiber-optic cables that can extend 500 meters—particularly valuable for organizations using Xserve RAID to back up or archive critical data. What's more, long-distance connectivity provides the added protection of physical separation in the event of a major disaster.

High-Availability Features

A modular design that features redundant components and easy serviceability makes Xserve RAID ideal for operations where data availability is crucial. All active components are field-replaceable and most are hot-swappable. LEDs on the front and rear indicate the health of each component: green means OK, yellow indicates warning, and red signals failure. In the event of a component failure, the remote monitoring system notifies the administrator via email, allowing repairs to be accomplished in seconds—usually without shutting down the system.

Redundant components

The environment managers monitor the Xserve RAID enclosure and the status of system components. If one of the power supplies fails, the other takes over and powers the system alone until the failed module is replaced. The environment managers themselves are redundant: Both monitor the entire enclosure and record status information on all components, enabling automatic adjustments and remote notifications as needed. If an environment manager fails, the other one takes charge of monitoring the system.

UPS support

A single Xserve RAID power supply can sustain the system in the event of a power outage or electrical surge. Each RAID controller has its own DB-9 serial port, allowing Xserve RAID to connect to and monitor up to two UPS sources. When the system detects that input power has been switched to a UPS source, it changes the cache mode from high-performance write-back to safer write-through cache—protecting data transactions should the UPS fail. If the UPS system is capable of reporting a low power level, Xserve RAID can alert the administrator of the condition, permitting a manual shutdown before the power runs out.

Cache Backup Battery Modules

With optional Cache Backup Battery Modules in Xserve RAID, data transactions in the RAID controller cache are protected for more than 72 hours after an input power failure. Once power is restored, Xserve RAID writes the transactions and begins recharging the batteries. If battery power is low, the system remains in write-through cache mode until the battery charge exceeds 50 percent, when it returns to write-back mode.

Hot sparing

For each RAID controller, any drives not assigned to an array are automatically used as global hot spares. If a drive fails, the RAID controller rebuilds its data on the spare drive without requiring intervention by the administrator. The rebuild operation occurs in the background while the controller processes normal host reads and writes—so service continues uninterrupted.⁵ To promote the best mix of performance and rebuild speed, the rebuild rate varies with the amount of traffic, or fabric load, on the RAID controller. As soon as the administrator replaces the failed drive, Xserve RAID configures it as a new global hot spare for the array.

Easy serviceability

Power supplies, cooling modules, RAID controller modules, and hard drives can be removed and replaced in seconds without tools. These components are all hot-swappable, which means they can be replaced without interrupting data availability. In addition, Xserve RAID has a shutdown feature that allows for planned maintenance without running down the cache backup batteries. This feature can also be beneficial if Xserve RAID is stored for long periods of time.

Service, Support, and Training Options

Every Xserve RAID comes with a one-year limited warranty and 90 days of up-and-running telephone support. In addition, the AppleCare website publishes in-depth product information, training on hardware and software installation and configuration, and technical resources, including the AppleCare Knowledge Base, discussions, and downloadable software on Apple's Featured Software site.

For critical data storage deployments, Apple also offers a comprehensive range of service and support options for Xserve, Xserve RAID, and Mac OS X Server. For more information about these AppleCare products, see www.apple.com/server/support.

AppleCare Premium Service and Support Plan

This server-class support product provides up to three years of up-and-running telephone and email support and onsite hardware service. Apple technical support experts are available 24 hours a day to help you determine whether you're experiencing a hardware failure or a RAID Admin configuration issue. In either case, Apple will work to get systems up and running quickly.

The AppleCare Premium Service and Support Plan delivers up-and-running telephone and email support within 30 minutes—24 hours a day, seven days a week. The hardware repair coverage provides onsite response within four hours during business hours and next-day onsite response when you contact Apple after business hours (terms apply).⁷ For added peace of mind, you'll have the assurance that Apple-authorized technicians will perform repairs using genuine Apple parts.

The AppleCare Premium Service and Support Plan can be purchased at any time while Xserve RAID is still under its original one-year warranty. However, since coverage ends three years after the hardware purchase date, you'll get maximum advantage when you make both purchases at the same time.

AppleCare Service Parts Kit

Xserve RAID is designed for quick and easy swapping of crucial parts; no special tools or training certifications are needed. AppleCare Service Parts Kits let system administrators keep key components handy to address the most common hardware failures. Each kit has an Apple Drive Module, a power supply, a cooling module, a RAID controller module, and a Fibre Channel cable (SFP to SFP). When the AppleCare Premium Service and Support Plan is combined with an AppleCare Service Parts Kit, technical support experts can often help troubleshoot and fix systems right over the phone—day or night—eliminating the need for an onsite technician.

Mac OS X Server Software Support

Apple also offers expert consultative phone and email support for Mac OS X Server integration and migration issues, as well as help with command-line configuration. Three levels of Mac OS X Server Software Support are available—Select, Preferred, and Alliance—depending on the number of incidents supported and desired response time. Each plan provides one year of coverage.

Training and Certification Programs

Apple offers comprehensive instruction on Mac OS X and Mac OS X Server, covering such topics as client management, system troubleshooting, and cross-platform network configuration. A combination of lectures, demonstrations, and hands-on exercises, classes are taught by Apple Certified Trainers with real-world experience and dynamic presentation skills. Classes are offered at Apple Authorized Training Centers, as well as at customer locations.

Once IT professionals have acquired the requisite skills, Apple certification programs provide tangible evidence of their technical expertise. For more information about Apple training and certification programs, visit www.apple.com/training.

Purchasing Information

Standard Configurations

The following standard Xserve RAID configurations are available. You can also order custom-configured systems with additional Apple Drive Modules and more cache memory.

| Order number | M9271LL/A | M9272LL/A | M9273LL/A |
|--------------------------------|---|---------------------------------|------------------------------------|
| Total available storage—RAID 0 | 1TB ¹ | 1.75TB ¹ | 3.5TB ¹ |
| Usable storage—RAID 1 | 500GB ¹ | 750GB ¹ | 1.5TB ¹ |
| Usable storage—RAID 3 and 5 | 750GB ¹ | 1.5TB ¹ | 3TB ¹ |
| ATA/100 Apple Drive Modules | Four 250GB drives ¹ | Seven 250GB drives ¹ | Fourteen 250GB drives ¹ |
| On-drive cache | 8MB per drive | 8MB per drive | 8MB per drive |
| Cache memory | 128MB per controller | 128MB per controller | 128MB per controller |
| Expansion | Fourteen drive bays with independent 100MB/s ATA channels for up to 3.5TB of storage ¹ | | |
| Also included | Mounting screws with M5, M6, and 10/32-inch threads; caged nuts; two agency-approved 12-foot power cables | | |
| Software | Xserve RAID Admin Tools CD | | |
| Service and support | 90 days of free telephone support and one-year limited warranty; optional extended service and support products | | |

System Requirements

Xserve RAID connects to the following Apple systems using the Apple Fibre Channel PCI Card (cables included):

- Xserve G5 or Xserve with Mac OS X v10.2.4, Mac OS X Server v10.2.4, or later
- Power Mac G5 with Mac OS X v10.2.7, Mac OS X Server v10.3, or later
- Power Mac G4 (800MHz or faster) with Mac OS X v10.2.4, Mac OS X Server v10.2.4, or later

Xserve RAID connects to Windows- and Linux-based systems using a compatible Fibre Channel PCI card and cables; check with manufacturer for details.

Related Products

- Apple Fibre Channel PCI Card with two copper SFP to SFP cables
- 250GB Ultra ATA Apple Drive Modules
- Cache Backup Battery Modules (two)
- Xserve G5, available in single- or dual processor server configurations and dual processor cluster node configuration
- AppleCare Premium Service and Support Plan
- AppleCare Service Parts Kit
- Third-party products, including racks, fiber-optic cables, SFP transceivers, and Fibre Channel switches

For up-to-date information on these and other products that enhance Xserve RAID deployments, visit www.apple.com/store or call 800-MY-APPLE.

Technical Specifications

Storage

- Fourteen drive bays on independent 100MB/s channels for up to 3.5TB of total storage¹
- Four, seven, or fourteen 250GB 7200-rpm ATA/100 Apple Drive Modules with 8MB disk cache (standard configurations); empty drive bays contain blank modules
- Support for reading SMART data from Apple Drive Modules for prefailure notification

RAID controllers and cache memory

- Dual independent controllers, each with an environment management coprocessor for out-of-band remote management and monitoring
- 128MB of cache per controller (256MB total); support for up to 512MB per controller (1GB total)
- Cache Backup Battery Modules (sold separately) for over 72 hours of memory protection

RAID operation

- Support for RAID levels 0, 1, 3, 5, 0+1, 10, 30, 50 (10, 30, and 50 using host-based software RAID)
- Support for multiple RAID sets, multiple hosts, and LUN masking and mapping
- Background RAID set creation; automatic variable background rebuilding⁵; online expansion; LUN slicing; global drive hot sparing (per RAID controller)

Fibre Channel storage-to-host connection

- Dual 2Gb Fibre Channel ports (SFP); 200MB/s throughput per channel with guaranteed bandwidth (400MB/s full duplex)³
- Host connectivity using 2Gb Apple Fibre Channel PCI Card (sold separately) or compatible third-party PCI card
- Support for point-to-point, loop, and switched fabric topologies
- Dual 10/100BASE-T Ethernet interfaces for remote management

Apple Fibre Channel PCI Card (sold separately)

- 64-bit, 66MHz card with two SFP 2Gb Fibre Channel ports; compatible with 32-bit, 66MHz PCI slots and 64-bit, 100MHz or 133MHz PCI-X slots
- Two 2.9-meter Fibre Channel copper cables with embedded SFP transceivers; compatible with SFP fiber-optic cables and transceivers

Cooling

- Redundant, hot-swappable cooling modules with self-regulating speeds and front-to-back cooling
- Environmental monitoring system for automatically maintaining optimal ambient temperature

Electrical

- Redundant, load-sharing hot-swappable power supplies (450W); universal input (100V to 240V AC), power factor corrected
- Maximum input current: 7.6A (100V to 127V) or 3.6A (200V to 240V)
- Power usage: 300W typical continuous power, 400W maximum continuous power
- Dual DB-9 serial ports for UPS systems
- Frequency: 50Hz to 60Hz, single phase

Environmental requirements and approvals

- Operating temperature: 50° to 95° F (10° to 35° C)
- Storage temperature: -40° to 116° F (-40° to 47° C)
- Relative humidity: 5% to 95% noncondensing
- Maximum thermal output: 1365 BTUs per hour
- Maximum altitude: 10,000 feet
- FCC Class A approved

Size and weight

- Height: 5.25 inches (13.3 cm)
- Width: 17 inches (43.2 cm)
- Depth: 18.4 inches (46.7 cm)
- Fits EIA-310-D-compliant, industry-standard 19-inch-wide four-post racks from 24 to 36 inches deep; deeper racks require third-party extender
- 60 to 110 pounds (27 to 45 kg), depending on configuration

For More Information

For more information about Xserve RAID and other Apple server solutions, visit www.apple.com/server.

¹For hard drive capacity measurements, 1GB = 1 billion bytes and 1TB = 1 trillion bytes; actual formatted capacity less. Maximum capacity of 3.5TB achieved through use of fourteen 250GB Apple Drive Modules. Usable capacity depends on drive configuration and RAID level. ²See www.apple.com/xserve/raid for information on third-party certifications and qualifications. ³Actual rates will vary depending on drive configuration and RAID level. ⁴Host operating system limitations apply. ⁵Automatic rebuild requires a designated hot spare; not available for RAID level 0. ⁶Throughput based on testing performed by Apple in January 2004 using Iometer (version 2003.05.10) sequential read tests with 4MB block size and 16 workers. Xserve RAID was configured with 14 drives and RAID 50; the host was a Power Mac G5 with 1GB SDRAM, Mac OS X Server v10.3.1, HFS+ (journalled) file system, and Apple Fibre Channel PCI Card. ⁷A separate AppleCare Premium Service and Support Plan must be purchased for each Xserve RAID system to be covered. To qualify, systems must be within the one-year hardware warranty. Coverage ends three years after date of Xserve RAID purchase. Actual onsite response time and availability of onsite service depend on location; see www.apple.com/support/products/premium for details. Support also covers Apple-qualified Xserve RAID solutions; see www.apple.com for a list of supported third-party solutions. Local telephone fees may apply; telephone numbers may vary and are subject to change.

© 2004 Apple Computer, Inc. All rights reserved. Apple, the Apple logo, Apple Store, Mac, Macintosh, Mac OS, Power Mac, and Xserve are trademarks of Apple Computer, Inc., registered in the U.S. and other countries. Rendezvous is a trademark of Apple Computer, Inc. AppleCare is a service mark of Apple Computer, Inc., registered in the U.S. and other countries. Java and all Java-based trademarks are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries. PowerPC is a trademark of International Business Machines Corporation, used under license therefrom. Other product and company names mentioned herein may be trademarks of their respective companies. Product specifications are subject to change without notice. This material is provided for information purposes only; Apple assumes no liability related to its use. January 2004 L301297A