Xserve RAID

Technology Overview
October 2005

Contents

Introduction

Features at a Glance

- Larger, 500GB drive modules for even more capacity—up to 7TB of storage per 14-drive system[1]
- Industry-standard SFP (Small Form-factor Pluggable) connectors for additional deployment flexibility in optical or copper infrastructures
- Advanced capabilities such as LUN slicing, online expansion, and background scan
- Certification of compatibility by leading storage infrastructure vendors
- Qualification for use in Linux, Windows, NetWare, and mixed-platform environments

The Xserve RAID is 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 puts 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.

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 qualified for use in Linux, Windows, NetWare, 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 7TB of storage[1] 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 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 7TB of total storage—in a rack-optimized 3U enclosure.[1] Each 500GB hard drive connects to a dedicated Ultra 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 98TB 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. Dual independent RAID controllers provide protected storage with unprecedented performance. In fact, Xserve RAID boasts an average read throughput of over 380MB/s[6] and write throughput of up to 301MB/s.[7] While throughput rates are measured differently in video applications, Xserve RAID is fast enough to support real-time, uncompressed, high definition (HD) and multiple-stream (SD) video editing using Apple's Final Cut Pro without dropping a frame.

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 to complete. 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.

Comprehensive service and support. To ensure rapid issue resolution for your server and storage deployments, choose from a full range of AppleCare products designed to provide integrated expert support.

Apple Fibre Channel PCI-X Card
Xserve RAID connects to a host Xserve or Power Mac system using the dual-port 2Gb, 133MHz Apple Fibre Channel PCI-X Card (sold separately) with throughput of up to 400MB/s.[3] Xserve RAID boasts an average read throughput of over 380MB/s[6] and write throughput of up to 301MB/s.[7] 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

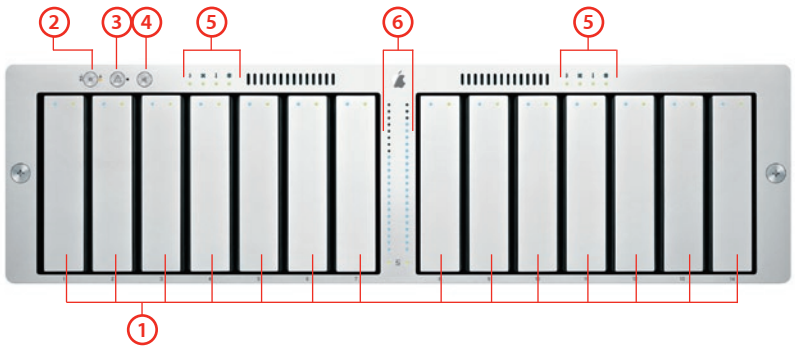

1. **Apple Drive Modules.** Xserve RAID holds up to 14 hot-swappable Apple Drive Modules. Each drive is on an independent channel to maximize data throughput and increase system reliability. LEDs on the front of each module indicate drive health and activity.
2. **System lock.** A built-in lock secures 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

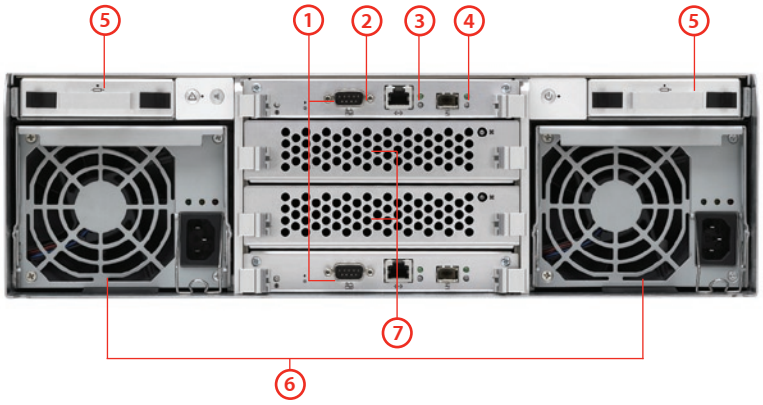

1. **RAID controller modules.** Two independent storage processor units manage RAID functions, data transfers, and failure protection for each set of seven drives. The environment management coprocessor in each controller simplifies configuration and management of RAID sets.
2. **Serial ports.** Standard 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 assured, 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.

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 over 380MB/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 AX100	IBM DS4100 FAST 100	Sun StorEdge 3511
Capacity	7TB	3.5TB	7TB	4.8TB
Size	3U	2U	3U	2U
Price*	$12,999	$14,866	$23,995	$20,995
Price per GB	$1.86	$4.25	$3.43	$4.37

* Based on prices published on manufacturers' websites as of August 31, 2005. Configurations vary. Features are similar in class.

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 7TB 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. 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:

- QLogic
- Brocade
- McDATA
- Emulex
- Cisco
- VERITAS
- ATTO Technology
- LSI Logic

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
- Windows 2000 Server
- Windows 2000 Professional
- Red Hat Enterprise Linux v2.1 and v3
- Novell NetWare v5.x and v6.x
- SUSE Enterprise Server 9
- Yellow Dog Linux v3

Primary storage for Mac, Windows, Linux, and NetWare 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 7TB,[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 (NAS) devices. Offering the compatibility of open standards and Apple's legendary manageability, this high-performance, high-capacity solution allows Mac, Windows, Linux, and NetWare 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 (SAN).

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—an average read throughput of over 380MB/s[6] and write throughput up to 301MB/s[7]—required for video post-production. In fact, Xserve RAID is fast enough to support real-time, uncompressed, high definition (HD) and multiple-stream (SD) video editing without dropping a frame. In addition, with up to 7TB 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 over 13 hours of protected RAID level 5 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.

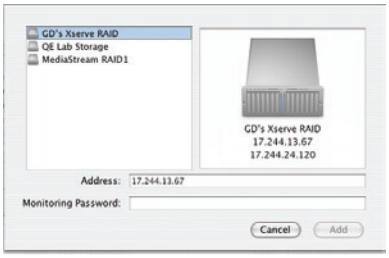

Automatic discovery with Bonjour
Using Apple's Bonjour technology, RAID Admin automatically discovers Xserve RAID systems on your network. This makes it easy to add them to your management window, without having to enter each IP address manually.

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.

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.

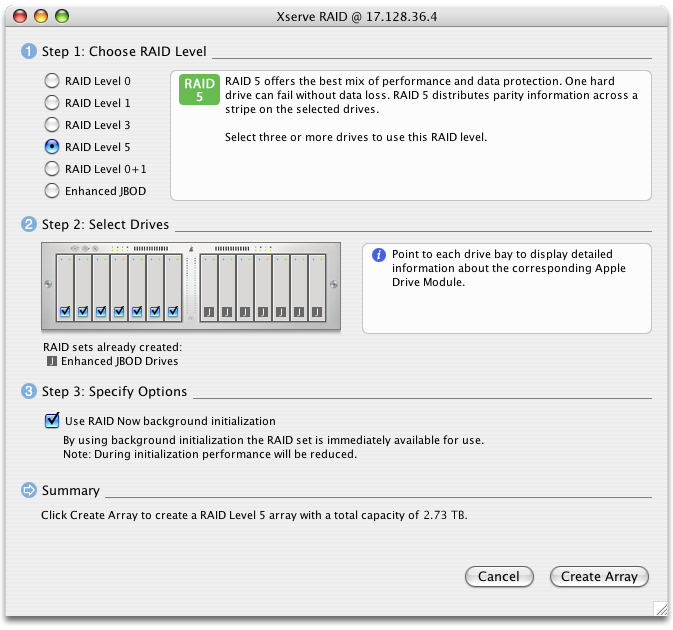

RAID Admin provides an easy-to-use interface for creating RAID sets.

Expanding arrays

Adding more disks to an existing array is accomplished using RAID Admin. Expanding an array preserves the existing array while making available the capacity of the added drives. Expansion in RAID Admin restripes the data across the entire set of new disks to provide optimal performance and capacity utilization.[8]

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 8 LUNs per RAID controller, or up to 16 LUNs per Xserve RAID system.

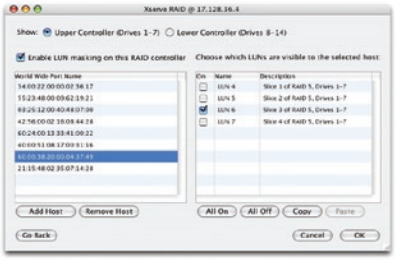

An intuitive interface makes it easy to set LUN mapping as well as other storage options.

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.

Enhanced JBOD mode

Enhanced JBOD mode allows each drive to become a separate LUN. This mode is especially useful for high-end database applications that have specific needs on how data is addressed. These applications provide protection using specialized host mirroring and striping. Xserve RAID provides JBOD performance that is enhanced by using the RAID controller's sophisticated adaptive Write Cache and Read Pre-fetch. When using enhanced JBOD mode, LUN masking and hardware slicing are disabled; therefore, this mode is recommended only for applications that provide these features.

Steady Streaming mode

Steady Streaming mode allows Xserve RAID users to provide consistent read latency for applications requiring it. Steady Streaming mode can compensate for the mechanical instabilities and external vibrations that can cause a hard drive to deliver data to the system with undesirable latency.

Background scan and repair

Xserve RAID enables remote background scans to repair or map out bad blocks on a RAID set. You can set a load-sharing priority, allowing the background scan to operate without causing a significant degradation in overall RAID set performance.

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.

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.

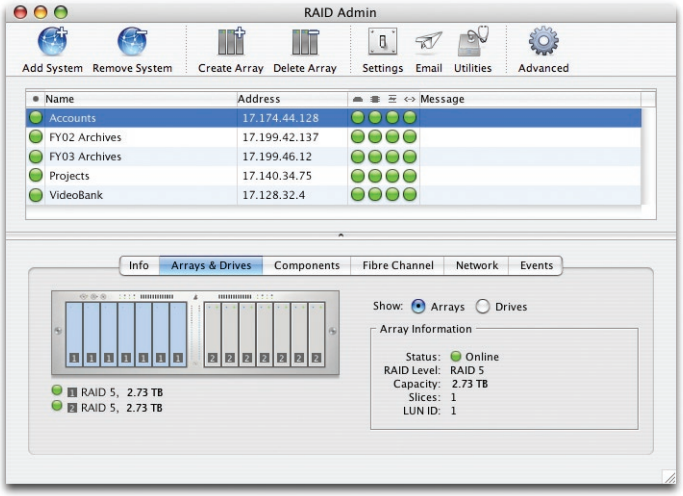

RAID Admin software displays summary health and activity information for multiple Xserve RAID systems.

Managing Volumes

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.

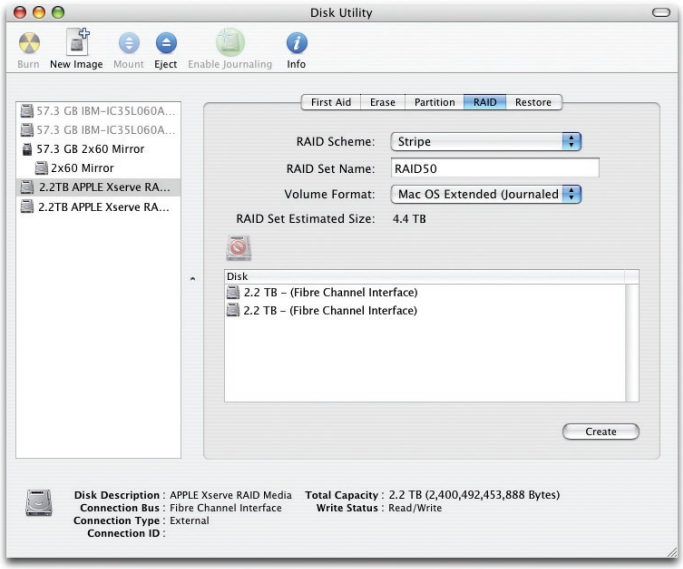

Disk Utility makes it easy to use software RAID to create hybrid RAID levels 10, 30, and 50.

Managing Volumes with Xsan

Together Xserve RAID and Xsan provide enterprise-class storage management. Xsan, Apple's high-performance 64-bit SAN file system for Mac OS X, enables concurrent sharing of terabytes of data over an ultrafast Fibre Channel network. Xsan is a powerful and scalable solution for enterprise storage consolidation and demanding post-production workflows.

With Xsan, it is easy to configure a storage network to meet the needs of your organization. It features sophisticated volume management and monitoring tools.

Xsan features Xsan Admin for setup, management of your storage systems, monitoring, and notifications. Xsan Admin provides an easy-to-use graphical user interface that enables administrators to perform complex administration tasks that would otherwise be accessible only from a command-line interface.

Xsan is easy to install and configure on Mac OS X and Mac OS X Server systems. Installation and configuration can be performed locally or remotely over your local IP network.

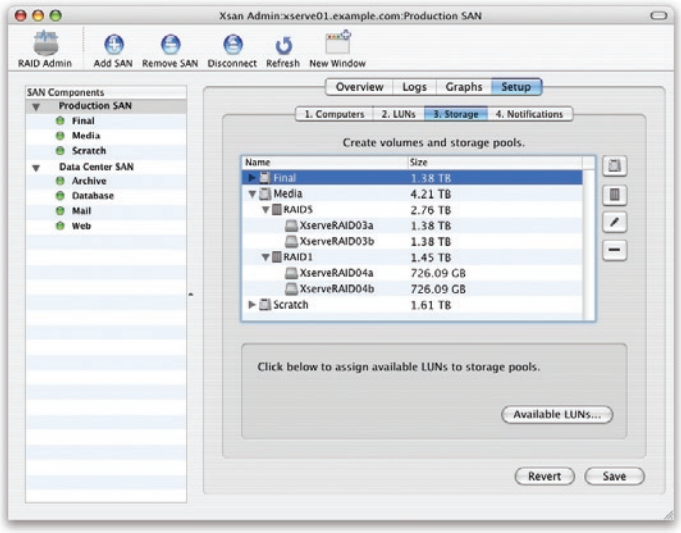

Xsan Admin includes real-time monitoring of Xsan systems and notifies the Xsan administrator of various events using email or pager. Xsan Admin provides the following status information:

- Free space in a volume or storage pool
- User quota usage
- Graphs of CPU and network utilization
- Setup status notifications
- Status of file system processes
- Log file
- Connected clients
- Fibre Channel failures

Xsan Admin tools can be used to map volumes to specific systems on the SAN, allowing you to limit which systems can see which volumes. Keep in mind that all Mac OS X file permissions still apply to Xsan volumes. Volume mapping simply provides an added layer of control and security.

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 Ultra 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.

Xserve RAID Architecture

1. Apple Drive Modules
2. Independent ATA drive channels
3. Drive controllers
4. RAID controller modules
5. Cache memory
6. Dual independent RAID processors
7. Fibre Channel ports
8. Redundant environment managers
9. Ethernet ports
10. Serial ports
11. Redundant cooling modules
12. Redundant power supplies
13. Cache Backup Battery Modules (optional)

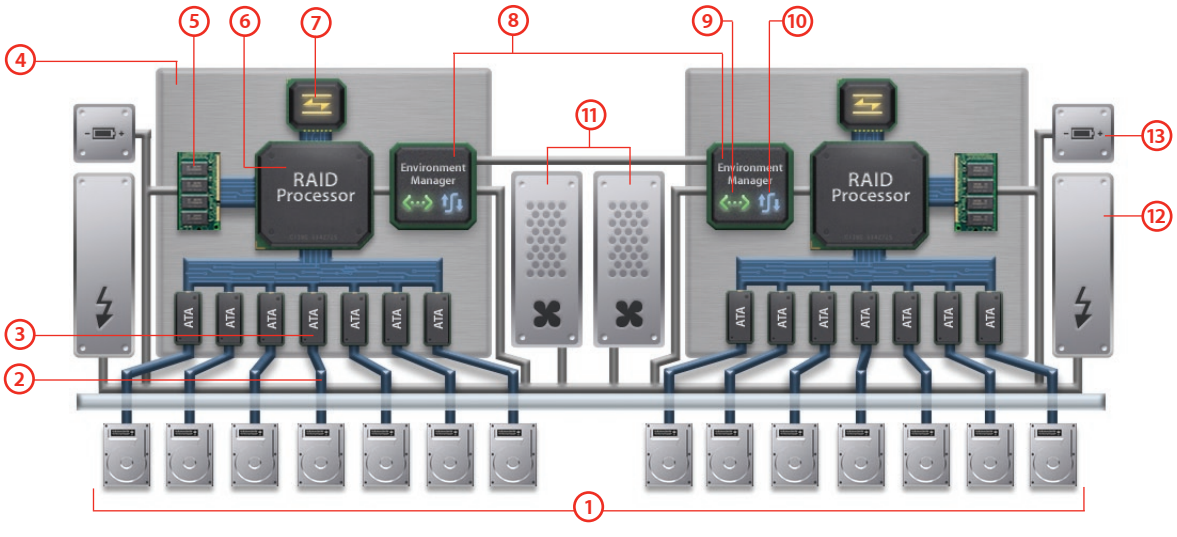

Apple Drive Modules
Xserve RAID uses reliable, high-performance Apple Drive Modules with true hot-swapping capabilities.

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 7TB capacity with 14 hot-swappable 500GB Apple Drive Modules. The high-capacity 7200-rpm Ultra ATA drive modules provide affordable expansion and configuration flexibility as storage needs grow.

Fourteen independent drive channels
The multithreaded Xserve RAID architecture features a dedicated Ultra ATA 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 they are 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 or enhanced JBOD configuration, since they do not 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 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-X Card
Xserve RAID connects to an Xserve or Power Mac host system using the dual-port 2Gb, 133MHz Apple Fibre Channel PCI-X 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, short- and long-haul SFP transceivers, and fiber-optic cables, 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 or PCI-X 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-X Card accepts short- and long-haul SFP 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, with the exception of the RAID controllers, which are warm swappable, meaning 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

To keep your Xserve RAID system up and running, Apple offers the AppleCare Premium Service and Support Plan. If you experience a problem, whether a hardware failure or a software configuration issue, Apple's dedicated Server Technical Support team will help diagnose the problem and work to get your systems back online quickly. And because Apple hardware and software are uniquely integrated, there's only one AppleCare number to call.

This server-class support product delivers up to three years of 24/7 telephone and email support—with 30-minute response.[9] The plan covers server administration and network management issues using the graphical user interface of Mac OS X Server or RAID Admin software. It also covers connectivity issues between your Mac or Xserve and your Xserve RAID storage system.

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.[9] With the AppleCare Premium Service and Support Plan, you'll have the assurance that Apple-authorized technicians will perform Xserve RAID repairs using genuine Apple parts.

The AppleCare Premium Service and Support Plan can be purchased at any time while your 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, and a RAID controller module. 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 IT department-level support for network configuration and administration; the integration of Mac OS X into heterogeneous environments; web application server support; support for technical issues requiring the use of command-line tools for resolution; and workflow issues in your environment. Three levels of Mac OS X Server Software Support are available—Select, Preferred, and Alliance—depending on the number of incidents supported, desired response time, and technical account management required. 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.

Order number	M9721LL/A	MA208LL/A	MA209LL/A
Total available storage—RAID 0	1TB[1]	3.5TB[1]	7TB[1]
Usable storage—RAID 1	500GB[1]	1.5TB[1]	3TB[1]
Usable storage—RAID 3 and 5	750GB[1]	3TB[1]	6TB[1]
Ultra ATA Apple Drive Modules	Four 250GB drives[1]	Seven 500GB drives[1]	Fourteen 500GB drives[1]
On-drive cache	8MB per drive	8MB per drive	8MB per drive
Controller cache	512MB per controller	512MB per controller	512MB per controller
Expansion	Fourteen drive bays with independent Ultra ATA channels for up to 7TB 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 telephone support and one-year limited warranty; optional extended service and support products		

Note: 250GB drives are also available in 7- and 14-drive configurations.

System Requirements

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

- Xserve G5 or Xserve with Mac OS X or Mac OS X Server v10.3.9, v10.4.2, or later
- Power Mac G5 or G4 (800MHz or faster) with Mac OS X or Mac OS X Server v10.3.9, v10.4.2, or later

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

Related Products

- Apple Fibre Channel PCI-X Card with two copper SFP-to-SFP cables
- 500GB Ultra ATA Apple Drive Module
- 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
  - Fibre Channel switches
  - Tape backup
  - Uninterruptible power supply (UPS)

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 supporting up to 7TB of total storage[1] using Apple Drive Modules, available in the following capacities:
  - 250GB 7200-rpm Ultra ATA with 8MB disk cache
  - 500GB 7200-rpm Ultra ATA with 8MB disk cache and rotational vibration safeguard
- 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
- 512MB of cache memory 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, and 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 rebuilding[5]; 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-X Card (sold separately) or compatible third-party PCI and PCI-X cards
- Support for point-to-point, loop, and switched fabric topologies
- Dual 10/100BASE-T Ethernet interfaces for remote management

Apple Fibre Channel PCI-X Card (sold separately)

- 64-bit, 133MHz 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 SFP transceivers; compatible with short- and long-haul SFP transceivers and fiber-optic cables

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: 3U rack-optimized, 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

AppleCare Service Parts Kit for Xserve RAID

- Includes RAID controller module, power supply, and cooling module
- Choice of Apple Drive Module:
  - 250GB Apple Drive Module,[1] order number M9383G/A
  - 500GB Apple Drive Module,[1] order number MA269G/A

For More Information

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

[1]For hard drive capacity measurements, 1GB = 1 billion bytes and 1TB = 1 trillion bytes; actual formatted capacity less. Maximum capacity of 7TB achieved through use of fourteen 500GB Apple Drive Modules. Usable capacity depends on drive configuration and RAID level. [2]See www.apple.com/xserve/raid for information on third-party certifications and qualifications. [3]Actual rates will vary depending on drive configuration and RAID level. [4]Host operating system limitations apply. [5]Automatic rebuild requires a designated hot spare; not available for RAID level 0. [6]Testing conducted by Apple in August 2005 using preproduction Xserve RAID systems. Iometer (version 2004.07.30) testing of raw disk throughput on Xserve RAID in both Mac OS X Server v10.4.2 and Windows XP environments has shown that Xserve RAID is capable of delivering up to 192MB/s on the standard shipping 4 x 250GB disk configuration utilizing a single controller, and an average of over 380MB/s on standard shipping 7 x 500GB and 14 x 500GB raw disk configurations utilizing both RAID controllers. Mac OS X Server v10.4.2 Xserve RAID testing conducted using directly attached Xserve G5 dual-processor 2.3GHz systems; Windows XP Xserve RAID testing conducted using directly attached Dell Precision 670 dual-processor 3.6GHz Xeon systems. Since MBR disks support partition sizes up to only 2TB, sliced arrays were used for Windows XP 14-disk tests. [7]For additional information on Xserve RAID performance in an Xsan environment, please refer to the Xsan Deployment and Tuning Guide, which provides examples for configuring and optimizing Xsan storage volumes: www.apple.com/server/documentation. [8]The expansion process is a background operation that may take several hours and affects overall throughput during the build. [9]A 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.