choosing

the right RAID solution for your business

introduction

Businesses today are increasingly using network servers to create enterprise-wide information technology solutions at a fraction of the cost of mainframe- or minicomputer-based solutions. Unfortunately, many of these systems are not designed for business-critical applications. As a result, disk drive failures can literally bring your business to a standstill — costing you thousands of dollars per hour in lost sales, productivity, and goodwill.

Redundant Arrays of Inexpensive Disks (RAID) technology was developed to address the fault-tolerance and performance limitations of conventional disk storage. While arrays were once considered complex and relatively specialized storage solutions, today they are easy to use and essential for a broad spectrum of client/server applications. This guide provides a highly understandable introduction to RAID technology and its applications, and summarizes the range of array solutions Adaptec offers.

A longtime leader in high-performance, system-based bandwidth solutions, Adaptec has the expertise and experience to help you choose the right array solution for your business. With our flexible, expandable array solutions, you can match performance and cost to a specific set of applications and easily grow over time — ensuring that you always have the high availability and data integrity needed.

At Adaptec, we are committed to helping you choose the right RAID solution for your business because in today's competitive environment, all data is business-critical.

What is an array?

An array is defined as two or more disks grouped together to appear as a single device to the host system. In common usage, the term implies the use of some form of redundancy to increase overall data availability, data integrity, and performance. There are several different RAID "levels" or redundancy schemes, each with characteristics that make it suited to a specific set of applications or network requirements.

scalable array solutions

Adaptec offers a broad selection of high-quality, highly-scalable array solutions that apply RAID technology to keep businesses up and running. This wide range of array solutions ensures that users can obtain the exact price/performance match to meet a variety of application and network requirements — from smaller networks running on entry-level servers to enterprise-wide applications.

Equally important, Adaptec's sophisticated, yet easy-to-use CI/O Array Management software is used for all our array solutions.* This gives businesses a common, user-friendly interface and seamless scalability, making it easy to add capacity or reconfigure arrays as needed.

Today, our array family includes three product lines. First is the AAA-130 series, PCI array adapters that are designed for entry-level Intel-based servers running file and print services, as well as small database applications. These array adapters deliver higher performance, more robust fault tolerance, and are easier to use than embedded network operating system (NOS) RAID software — at a price consistent with an entry-level server's cost structure.

** Adaptec CI/O Array Management software will be supported in future releases of the AEC-4312A and AEC-7312A external array controllers.*

Table of Contents

Next, there's the ARO-1130 PCI RAID*port* card, a high-performance RAID upgrade card that provides an ultra-affordable entry point to an array solution for entry-level servers. The ARO-1130 card uses a RAID*port* connector and embedded Adaptec SCSI ASICs in a server motherboard to deliver the same benefits as the AAA-130 series. *The ARO-1130 card is available exclusively through Adaptec original equipment manufacturers (OEMs).*

Finally, the fully platform-independent AEC-4312A and AEC-7312A high-performance external array controllers enable large-scale, fault-tolerant storage subsystems. These array controllers are optimized for midrange and high-end Intel-based and entry-level UNIX-based servers and workstations, as well as for clustered server environments.

Our flexible solutions allow users to start with one array solution — for example, one of the AAA-130 series of array adapters — and easily migrate to another Adaptec array solution, without reformatting, as needs change over time.

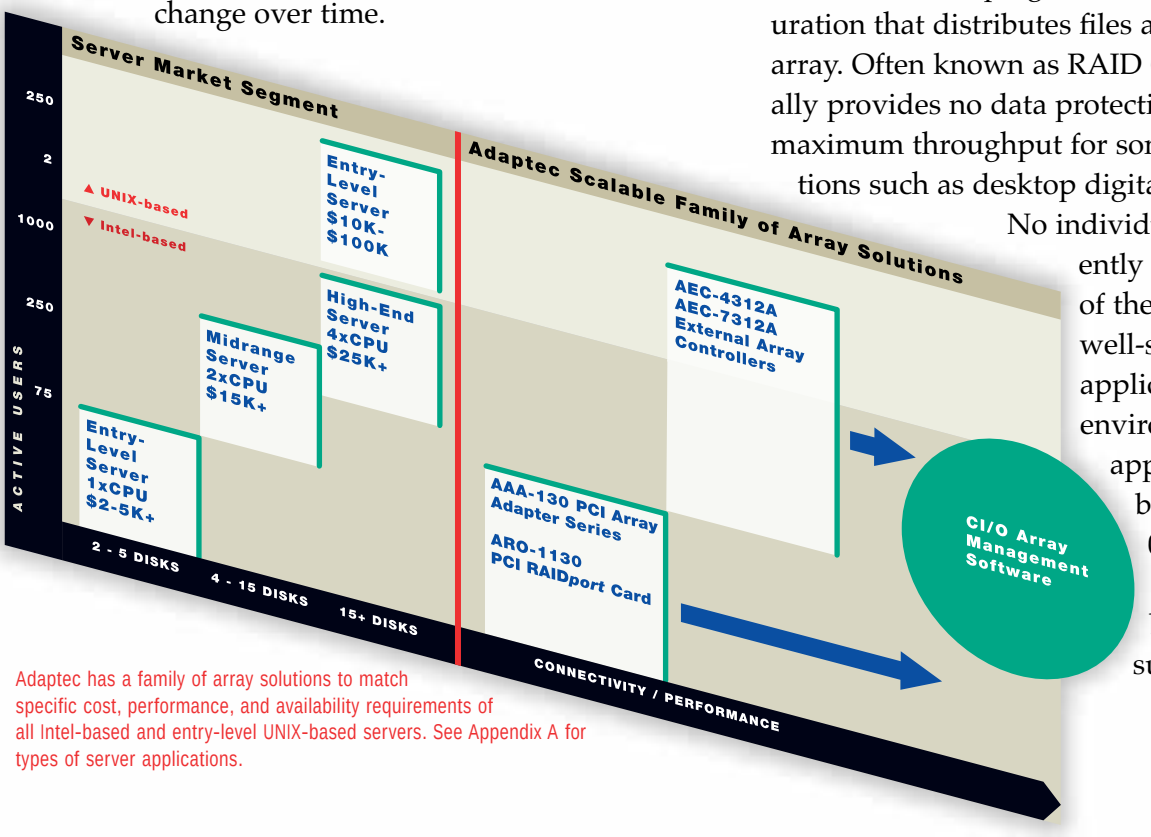

Adaptec has a family of array solutions to match specific cost, performance, and availability requirements of all Intel-based and entry-level UNIX-based servers. See Appendix A for types of server applications.

What is RAID technology?

RAID technology was first defined by a group of computer scientists at the University of California at Berkeley in 1987. The scientists studied the possibility of using two or more disks to appear as a single device to the host system.

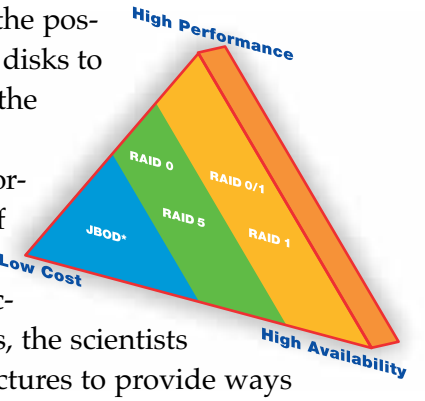

Although the array's performance was better than that of large, single-disk storage systems, reliability was unacceptably low. To address this, the scientists proposed redundant architectures to provide ways of achieving storage fault tolerance. In addition to defining RAID levels 1 through 5, the scientists also studied data striping — a non-redundant array configuration that distributes files across multiple disks in an array. Often known as RAID 0, this configuration actually provides no data protection. However, it does offer maximum throughput for some data-intensive applications such as desktop digital video production.

No individual RAID level is inherently superior to any other. Each of the five array architectures is well-suited for certain types of applications and computing environments. For client/server applications, storage systems based on RAID levels 1, 0/1, and 5 have been the most widely used. This is because popular NOSs such as Windows NT®

* *Just a bunch of drives.*

Server and NetWare manage data in ways similar to how these RAID architectures perform.

The table in Appendix B summarizes the strengths and weaknesses of each RAID level.

What does RAID provide?

RAID technology does not prevent drive failures. However, RAID does provide insurance against disk drive failures by enabling real-time data recovery without data loss.

← Each RAID level has inherent cost, performance, and availability (fault-tolerance) characteristics designed to meet different storage needs. Most RAID levels can effectively satisfy only one or two of these criteria. RAID 5 is the exception — providing the best overall balance of cost, performance, and availability for most servers.

The fault tolerance of arrays can also be significantly enhanced by choosing the right storage enclosure. Enclosures that feature redundant, hot-swappable drives, power supplies, and fans can greatly increase storage subsystem uptime based on a number of widely accepted measures: Mean Time to Data Loss (MTDL), Mean Time of Data Availability (MTDA), Mean Time to Repair (MTTR), and Mean Time Between Failure (MTBF). For brief definitions of these terms, please refer to the glossary.

The driving factors behind RAID

A number of factors are responsible for the growing adoption of arrays for critical network storage. As today's applications create larger files, network storage needs have increased proportionately. To accommodate expanding storage requirements, users are adding disk drives — raising the odds of drive failures. In addition, accelerating CPU speeds have outstripped data transfer rates to storage media, creating bottlenecks in today's systems.

Arrays overcome these challenges by providing a combination of outstanding data availability, extraordinary and highly scalable performance, and high capacity.

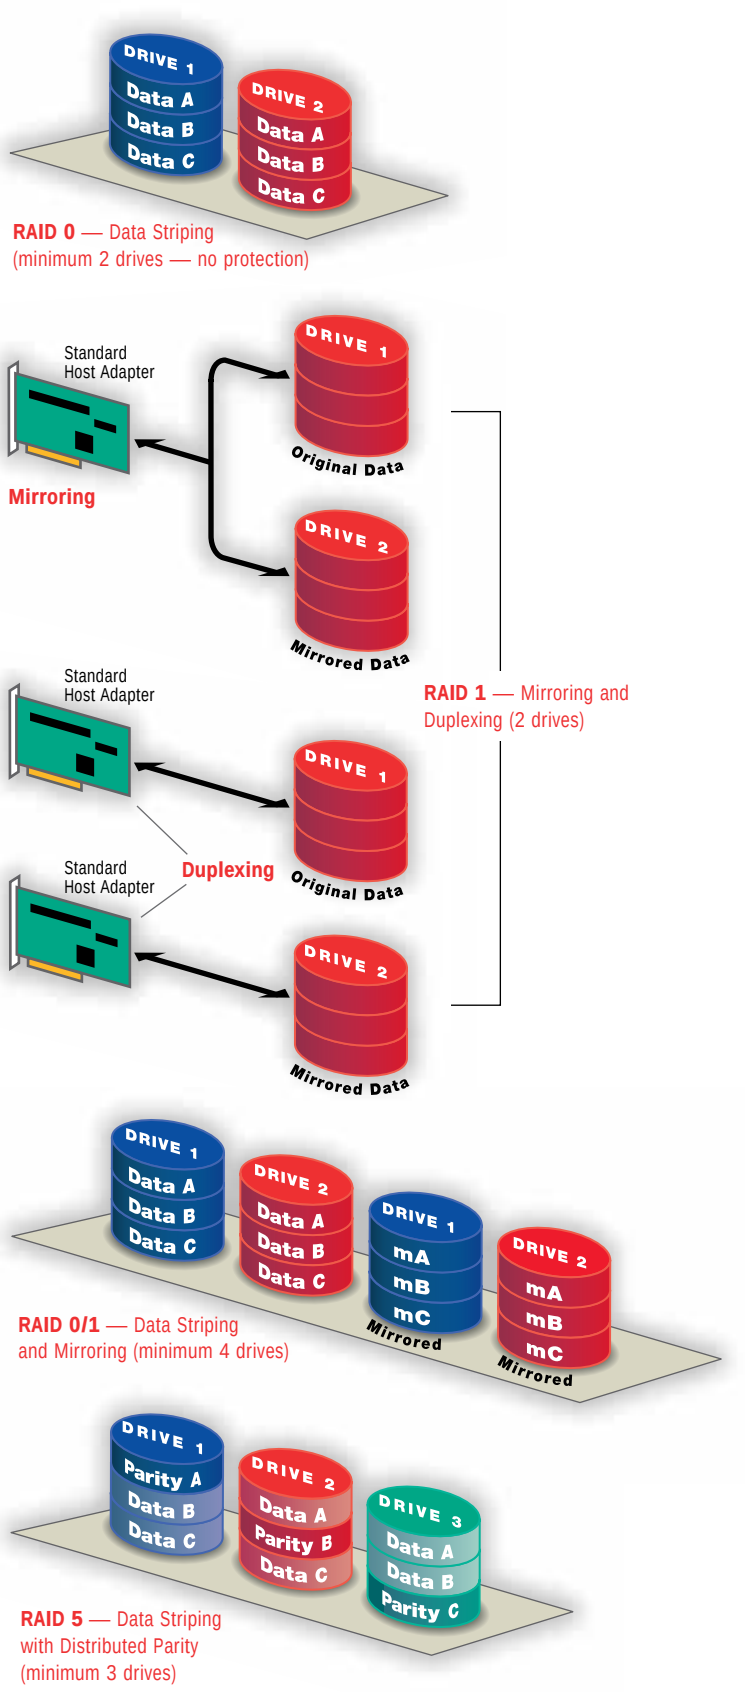

Duplexing — which eliminates the host adapter as a single point of failure — is one of the most popular methods for increasing fault tolerance under RAID 1 (mirroring). Used exclusively with low-performance, software-based RAID solutions, duplexing host adapters alone unfortunately, offers no protection against the failure of other server or storage system components. For example, the failure of power supplies and fans, which have relatively low MTBF ratings, can bring down a system more often than a host adapter failure.

Outstanding availability

More and more organizations have created enterprise-wide networks to improve productivity and streamline information flow. While the distributed data stored on network servers provides substantial cost benefits, these savings can be quickly offset if information is frequently lost or becomes inaccessible. RAID storage solutions provide critical reliability by enabling recovery with no loss of data or interruption of user access.

Scalable performance

While disk drive performance has risen steadily, it has not kept pace with the increase in CPU speeds. As a result, I/O processing time and throughput are limited by the capabilities of any single disk. Array solutions provide a greater range of performance scalability than individual drives as capacity is increased. In arrays, disks work together to handle multiple I/O requests simultaneously. Further, sustainable throughput can be improved because disks can be written and read in parallel.

High capacity

By integrating multiple drives into a single array — which is viewed by the network operating system as a single disk drive — organizations can create cost-effective, minicomputersized solutions of up to a terabyte or more of storage. This is much higher than the capacity of any single disk drive.

Types of arrays

There are three primary array implementations: software-based arrays, bus-based array adapters/controllers, and subsystem-based external array controllers. As with the various RAID levels, no one implementation is clearly better than another — although software-based arrays are rapidly losing favor as high-performance, low-cost array adapters become increasingly available. Each array solution meets different server and network requirements, depending on the number of users, applications, and storage requirements.

It is important to note that all RAID code is based on software. The difference among the solutions is where that software code is executed — on the host CPU (software-based arrays) or offloaded to an on-board processor (bus-based and external array controllers).

Software-Based Arrays Primarily used with entry-level servers, software-based arrays rely on a standard host adapter and execute all I/O commands and mathematically intensive RAID algorithms in the host server CPU. This can slow system performance by increasing host PCI bus traffic, CPU utilization, and CPU interrupts. Some NOSs such as NetWare and Windows NT include embedded RAID software. The chief advantage of this embedded RAID software has been its lower cost compared to higher-priced RAID alternatives. However, this advantage is disappearing with the advent of lower-cost, bus-based array adapters such as the AAA-130 series and the ARO-1130 RAID*port* card from Adaptec.

Bus-Based Array Adapters/Controllers Unlike software-based arrays, bus-based array adapters/controllers plug into a host bus slot [typically a 133 MByte (MB)/sec PCI bus] and offload some or all of the I/O commands and RAID operations to one or more secondary processors. Originally used only with mid- to high-end servers due to cost, lower-cost bus-based array adapters are now available specifically for entry-level server network applications (i.e., AAA-130 series from Adaptec).

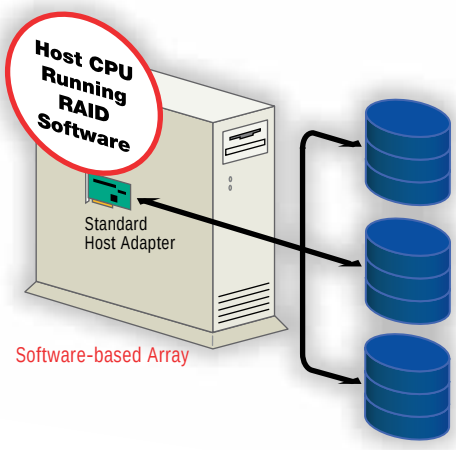

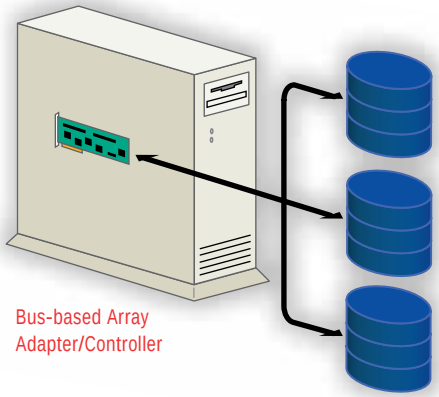

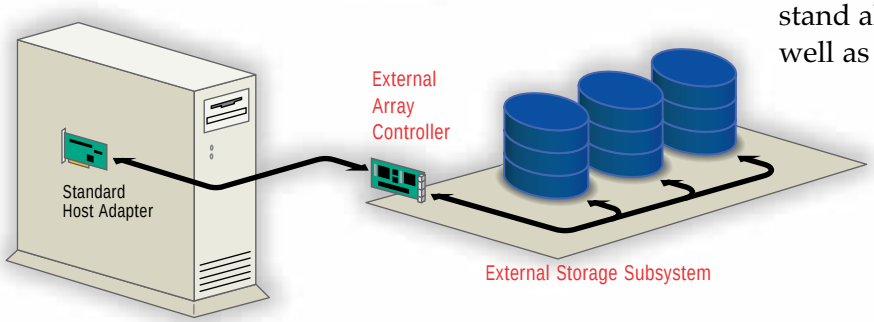

In addition to offering the fault-tolerant benefits of RAID, bus-based array adapters/controllers perform connectivity functions that are similar to standard host adapters. By residing directly on a host PCI bus, they provide the highest performance of all array types. Bus-based arrays also deliver more robust fault-tolerant features than embedded NOS RAID software.

As newer, high-end technologies such as Fibre Channel become readily available, the performance advantage of bus-based arrays compared to external array controller solutions may diminish.

External Array Controllers Intelligent external array controllers "bridge" between one or more server I/O interfaces and single- or multiple-device channels. These controllers feature an on-board microprocessor, which provides high performance and handles functions such as executing RAID software code and supporting data caching.

External array controllers — such as the AEC-4312A and AEC-7312A array controllers from Adaptec — offer complete operating system independence, the highest availability, and the ability to scale storage to extraordinarily large capacities (up to a terabyte and beyond). These controllers are usually installed in networks of stand alone Intel-based and UNIX-based servers as well as clustered server environments.

Basics of simple parity

The concept behind RAID is relatively simple. The fundamental premise is to be able to recover data on-line in the event of a disk failure by using a form of redundancy called parity. In its simplest form, parity is an addition of all the drives used in an array. Recovery from a drive failure is achieved by reading the remaining good data and checking it against parity data stored by the array. Parity is used by RAID levels 2, 3, 4, and 5. RAID 1 does not use parity because all data is completely duplicated (mirrored). RAID 0, used only to increase performance, offers no data redundancy at all.

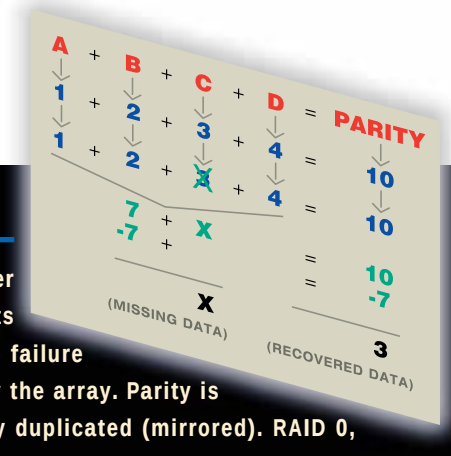

building an array storage subsystem

Adaptec array adapters and controllers are the cornerstones of a fault-tolerant storage solution. However, additional hardware components can greatly enhance the reliability and data availability of an array. These components include the storage enclosure, drives, terminators, and cables.

Storage enclosures

To create an array solution, users can choose from two types of storage enclosures: internal or external (see right). Many entry-level, PCI-based servers now include internal hot-swappable drives for array applications. Both the AAA-130 series and the OEM-offered ARO-1130 RAID*port* card are well-suited for these types of servers.

The second option is to buy an external array enclosure (see Appendix C) and add drives. When choosing the enclosure, it is important to balance cost with the need for storage system reliability. Various features found in array enclosures — such as redundant, hot-swappable power supplies and fans — can dramatically improve data availability, data integrity, and system serviceability.

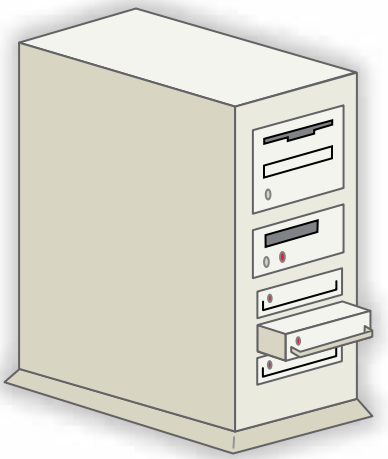

An internal (within the server) RAID solution usually incorporates removable drives to provide an inexpensive compact, real-time data recovery storage system.

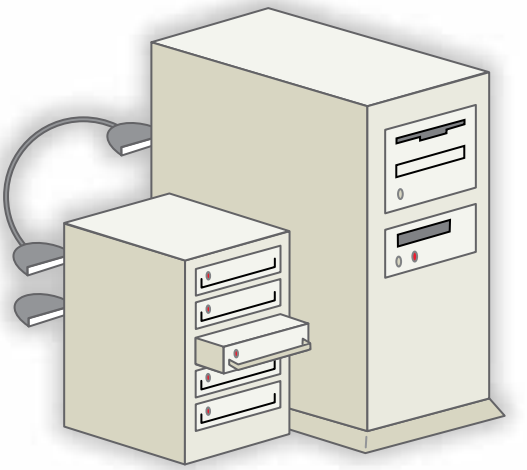

In addition to increased storage capacity, an external RAID solution may provide non-stop features beyond the server's capabilities, such as redundant power and cooling systems.

Array enclosures also differ in the way they allow drive replacement. Many server and storage manufacturers claim to have array storage systems that have removable, hot-swap drive bays. Over the past few years, the term "hot swap" has been applied to everything from being able to replace a drive while on-line and under-load to simply permitting the system power to remain on when adding or removing a drive. The latter feature is more correctly referred to as "warm swap."

For true on-line, under-load, hot-swap functionality, users should select a server or an array enclosure that has a SCSI backplane printed circuit board (PCB) or connector on the drive tray with power and ground pins that are longer than the data pins. This design ensures that power will always be disconnected after the data pins are removed from the SCSI bus and vice-versa. The result is that no power spike "glitches" can occur on the active SCSI bus, thereby avoiding data corruption or loss. For "warm-swap" enclosures, network managers

must see to it that all activity on the SCSI bus has ceased before replacing a drive. To verify whether a particular server or storage enclosure offers hot-swap support, check with the manufacturer.

Any Adaptec AAA-130 series adapter, ARO-1130 RAID*port* card, AEC-4312A controller, or AEC-7312A controller will provide an effective array solution when choosing an external storage subsystem. With a warm-swap enclosure, Adaptec's CI/O Array Management software gives network managers the added flexibility of pausing activity on the SCSI bus during drive replacement, without affecting the NOS.

Disk drives

The choice of disk drives can greatly impact the overall performance and capacity of array storage. Drives with relatively high MTBFs and high RPM speeds will fail less often and provide better performance than inferior drives. Keep in mind that mixing drives of different storage capacities will limit the capacity of all drives in the array to that of the lowest-capacity drive. This ensures that there is enough disk space available for redundancy information on all stored data. Appendix D provides a listing of some popular drive manufacturers.

In addition, many disk drives now include Self-Monitoring, Analysis, and Reporting Technology (S.M.A.R.T.). These drives can be set up to report predicted failures, providing managers the opportunity to replace a drive before problems arise.

Adaptec array adapters and controllers support mixing of different capacity drives within the same array. However, Adaptec strongly recommends using drives with similar spindle speeds and capacities for maximum performance and capacity.

Termination and cabling

Proper termination and high-quality cables should always be used with SCSI devices, particularly when building an array storage subsystem. When implementing Fast, Wide, or UltraSCSI technology, it is a good idea to use active terminators and impedance-matched cables to avoid signal reflections and data corruption. In addition, the maximum cable lengths outlined in the SCSI specification should be strictly followed.

Termination must be at the end of the cable or on the storage enclosure's backplane, not at the device. External array enclosures typically require an external active terminator that attaches to the SCSI connector on the back of the enclosure. Appendix E provides a listing of SCSI cable kits available from Adaptec.

Both the AAA-130 series of array adapters and AEC-4312A* and AEC-7312A* array controllers are available in kits that include high-quality cabling and active termination.

* *Kits available in subsequent releases.*

adaptec array solutions

Clearly, there are many compelling reasons to employ array solutions. The challenge in the past has been to find the right solution to get the required data availability and protection without paying a disproportionately high price or buying more of a solution than needed.

Adaptec's high-performance, high-availability array solutions are designed to meet a broad spectrum of business requirements and effectively scale as needs change over time.

AAA-130 PCI array adapter series

Smaller businesses, departments, and workgroups searching for a cost-effective RAID solution have often had to settle for the mirroring/RAID software native to an NOS. The easy-to-use Adaptec AAA-130 series of array adapters provides a better, highly-affordable alternative that outperforms embedded NOS RAID software, while offering superior fault tolerance for enhanced data availability and data integrity.

The AAA-130 series improves performance over NOS RAID software through a combination of two leading-edge hardware technologies (see above right). First, it creates a secondary on-board, 133 MB/sec PCI bus to offload and shield data traffic — such as disk reads/writes and RAID functions — from the host's main PCI bus. This unique PCI-to-PCI bridge architecture is critical because most entry-level, Intel-based servers have only one main PCI bus that supports three to four PCI slots. Add-in boards such as array adapters, host adapters, network cards, and video cards can greatly increase host main PCI bus traffic.

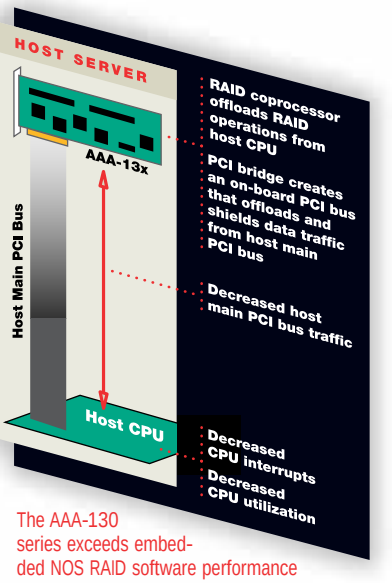

The AAA-130 series exceeds embedded NOS RAID software performance by using an on-board RAID coprocessor and unique PCI-PCI bridge design.

Second, a separate RAID coprocessor with 1 MB of memory boosts performance by offloading RAID 1 (mirroring) and RAID 5 write operations from the server CPU. With the AAA-130 series of adapters, the server CPU only has to pass data once across the main PCI bus. Then, the RAID coprocessor executes all RAID write operations on the board using the on-board PCI bus to talk directly with the disks in the array. When the RAID write operations are completed, a single interrupt is issued to the server CPU.

Beyond performance improvements, the AAA-130 series enhances data availability by enabling hot swapping of drives in enclosures that are designed with that capability. It also supports hot sparing, which allows users to configure spare drives for automatic recovery in the event of a failure — without having to manually swap drives. Further, because the AAA-130 series Array*Config*™ set-up utility makes it possible to boot systems from a disk array, the NOS itself can be protected from drive failure.

AAA-130 adapters support both disk and non-disk SCSI devices and are designed to meet a variety of connectivity, performance, and cost requirements. The AAA-131 single-channel model is perfect for systems with an installed add-in host adapter

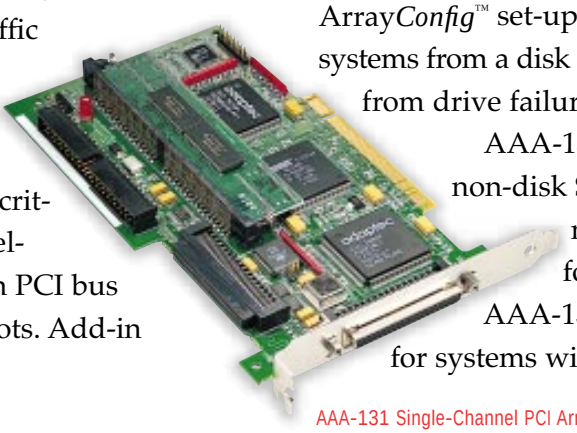

AAA-131 Single-Channel PCI Array Adapter

or embedded SCSI capabilities for low-speed tape and CD-ROM drives. However, if these slower-speed SCSI devices are only accessed during off-peak hours, then it is not necessary to connect the devices to a separate host adapter — thus, further reducing system costs. The AAA-133 adapter features three channels, enabling arrays to operate over multiple channels and allowing slower peripherals — such as CD-ROMs and tape drives — to communicate over their own, separate channel. This multichannel feature also preserves valuable PCI slots and reduces the need for a separate host adapter.

With the AAA-130 series, businesses can create up to four arrays in any combination of RAID levels, including RAID 0, RAID 1, RAID 0/1, and RAID 5. Up to 15 Fast/Wide SCSI or UltraSCSI devices can be connected per channel. The AAA-130 series is fully compatible with leading NOSs such as Windows NT, NetWare, OS/2 Warp,* SCO OpenServer,* and UnixWare.* This gives users more choices in systems and peripherals, while protecting a business' hardware and software investment.

AAA-130 array adapters are supplied in a complete high-value kit, including the array adapter with Adaptec Array1000™ Family Manager Set driver software, easy-to-use Array*Config* initial boot array installation software, Adaptec CI/O Array Management software, 50- and 68-pin internal cables with active termination, and complete documentation. Simplified procedures make installation fast and painless.

** Available in subsequent releases.*

ARO-1130 PCI RAID*port* card (OEM only)

Available exclusively from selected Adaptec OEM customers, the ARO-1130 PCI RAID*port* card is a cost-effective, high-performance array solution that delivers the same features and benefits as the AAA-130 array adapter series, but in a different implementation. The ARO-1130 card upgrades embedded Adaptec SCSI channels and connectors on a motherboard to SCSI RAID channels. OEMs can easily install a RAID*port* connector to make a motherboard with embedded Adaptec ASICs RAID-ready. Then, by plugging in the ARO-1130 RAID*port* card, the motherboard is automatically upgraded from standard SCSI to SCSI RAID. Using on-board SCSI controllers as RAID channels reduces the cost of the card, while leveraging and preserving the investment in motherboard SCSI.

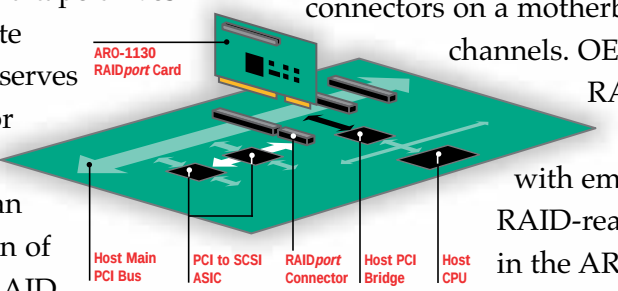

Once in place, the ARO-1130 RAID*port* card provides benefits that are similar to those delivered by the AAA-130 series array adapters — making it an ultra-affordable, high-performance RAID solution for entry-level servers. For more information about the ARO-1130 RAID*port* card, or for a current list of OEMs contact Adaptec, or visit Adaptec's Web site at www.adaptec.com/RAID.

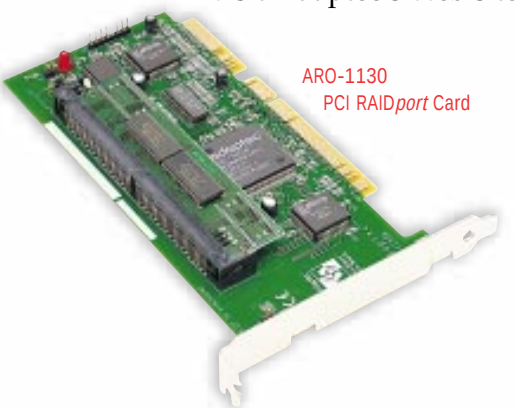

Adaptec
RAIDport
On Board

AEC-4312A and AEC-7312A external array controllers

Adaptec's AEC-4312A SCSI-to-SCSI and AEC-7312A Fibre Channel-to-SCSI external array controllers are ideal for high-availability, high-capacity client/server applications and clustered environments such as midrange and high-end Intel-based servers, as well as entry-level UNIX-based servers. Relying on a dual-channel Wide UltraSCSI disk interface and a Wide UltraSCSI host interface, the AEC-4312A controller delivers data transfer rates of up to 40 MB/sec per channel. The AEC-7312A controller features a blazing 100 MB/sec Fibre Channel host interface and a dual-channel Wide UltraSCSI device interface.

AEC-7312A External Array Controller

The two controllers offer a number of key benefits, such as increased performance, outstanding fault tolerance and scalability, easy field upgrade capability for investment protection, and simplified array configuration and management. Both the AEC-4312A and AEC-7312A controllers are fully operating system (OS) independent, connecting to any host without special OS drivers or software. The controllers can support up to 30 external hard disk drives, in up to four RAID arrays, to provide as much as 270 Gigabyte (GB) of RAID storage on a single controller. Since each controller is presented as a single SCSI ID to the host CPU, up to four controllers can be installed to deliver more than a terabyte (1000 GB) of storage capacity.

The performance advantages of the AEC-4312A and AEC-7312A controllers start with Adaptec's high-performance PCI-to-PCI bridge architecture. Each controller has two independent, on-board PCI buses — one optimized for microprocessor activities, the other for data transfers. This dual PCI bus design increases overall controller performance compared to controllers with a single bus and memory, by allowing multiple operations to be performed concurrently.

The controllers' innovative design also takes full advantage of multiple on-board processors to maximize data throughput by separately executing RAID commands, parity operations, and data caching. An on-board 133 MHz 5x86 microprocessor — equivalent to a 75 MHz Pentium — with its own memory subsystem makes quick work of controller I/O commands. In addition, two powerful on-board RAID coprocessors, supporting up to 32 MB of memory, accelerate performance by calculating the mathematically intensive parity and caching operations independently. Performance can be further enhanced for specific applications by utilizing the AEC-4312A and AEC-7312A controllers' advanced read-ahead, write-back, and write-through caching algorithms.

Both array controllers provide fault tolerance by supporting RAID levels 0, 1, 0/1, 3,* 4,* and 5. Redundancy is also supplied through support for hot-swap drives, hot sparing, ECC (error correction code) SIMMs, and battery backup of cached write-back data. For the highest data availability and performance, redundant controllers can be used to achieve active/active failover capability.*

Because host technologies tend to change more rapidly than storage technologies, both the AEC-4312A and AEC-7312A array controllers feature field-upgradable daughterboards. This designed-in upgrade path protects an organization's investment in controllers and storage by allowing simple, cost-effective migration to new host interface technologies as requirements change.**

** Support for RAID levels 3 and 4 and for active/active failover will be available in future releases.*

*** Subsequent releases of both controllers will also be available in a "cable-ready" canister for easy installation into a standard drive bay of any external RAID storage enclosure.*

adaptec CI/O array management software

Arrays provide powerful solutions delivering high availability and data integrity for today's business-critical applications. Unfortunately, the complexity of current array solutions can create difficult manageability challenges for network managers.

Adaptec CI/O Array Management software is the interface to all Adaptec array solutions,* simplifying array management and providing seamless scalability through a built-in upgrade path.

With Adaptec's CI/O Array Management software, network managers can monitor and manage storage either locally or remotely from any PC or workstation on the network. In Windows® operating systems, information is presented in an easy-to-use graphical user interface.

The management software allows network managers to see at-a-glance both physical and logical array configurations and other SCSI peripherals for any server using Adaptec array adapters and controllers. Alerts, color-coded by their severity, keep network personnel informed about all events — such as drive failures or additions to the array configuration — as they happen across the network. Reconfiguration of arrays is easy. Network managers simply point and click at the number of drives desired.

- Pick and click to reconfigure arrays
- Zero-in on individual servers and arrays
- Keep track of everything that happens

Without leaving their desks, network personnel can use the software to:

- Determine if any drives have failed
- Initialize arrays
- Reactivate offline drives
- Verify parity information
- Blink individual drive lights or the lights of all the drives in an array
- Pause I/O to permit the removal of any drives that aren't hot-pluggable
- Schedule management activities to be carried out automatically any time of the day or night.

In addition, disk drives equipped with S.M.A.R.T. can be set up to report predicted failures, giving network managers the opportunity to activate a spare drive before problems arise.

* *Adaptec CI/O Array Management software will be supported in future releases of the AEC-4312A and AEC-7312A external array controllers.*

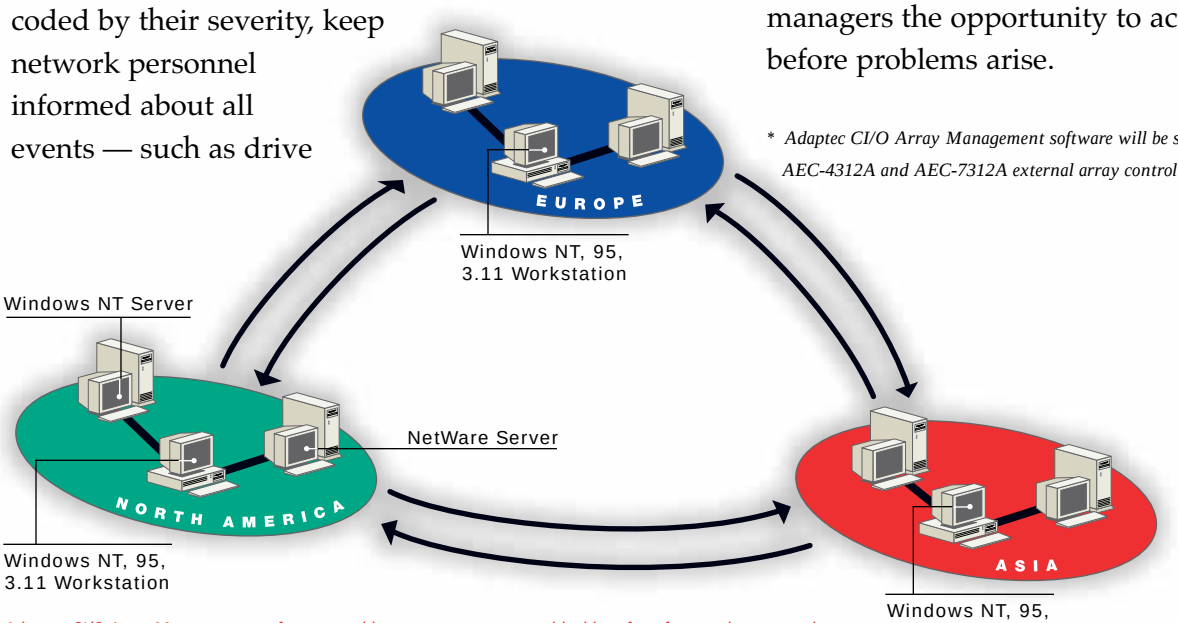

Adaptec CI/O Array Management software provides an easy-to-use graphical interface for creating, managing, and monitoring Adaptec array solutions from across the building, the country, or around the world.

count on adaptec for quality and reliability

glossary

Adaptec array solutions are backed by the same commitment to quality and reliability that has made Adaptec the recognized leader in high-performance system bandwidth solutions for many years.

Trusted for quality and reliability, all Adaptec array adapters and controllers must pass comprehensive functional and mechanical tests in two separate laboratories: the Adaptec Product Test Lab and the Adaptec Compatibility Test Lab. Adaptec has also earned the coveted ISO-9002 international quality certification, a distinction that ensures the highest level of quality in product design and production. Adaptec products are incorporated into the products of nearly all major computer and peripheral manufacturers around the world.

In addition, the company's products are designed to be compatible with a wide variety of hardware and software configurations, for lasting investment protection.

The result is array solutions that users can count on to deliver exceptional availability, data integrity, and performance for their business.

AAA - Adaptec Array Adapter For example, the AAA-130 array adapter series.

AEC - Adaptec External Controller For example, the AEC-4312A and AEC-7312A external array controllers.

ARO - Adaptec RAID Option For example, the ARO-1130 PCI RAID*port* card.

ASIC - Application Specific Integrated Circuit Often used in the context of motherboard embedded SCSI host adapters. *See also Embedded SCSI ASIC.*

Active–active array controllers Multiple external array controllers, each independently accessing storage devices to provide fault tolerance and improved performance. In the unlikely event of a controller failure, the remaining controller is able to provide uninterrupted access to the entire data storage subsystem. Active controllers offer redundancy as well as nearly twice the performance of a single controller (including active–passive controllers) because each one can access data independently. *See also active–passive array controllers.*

Active–passive array controllers Provides external array controller failover functionality similar to that of a hot-spare drive (or other redundant component). An active-passive array controller is powered up (but not used) and waiting to take over the tasks of the primary controller in the event of a failure. If the primary controller fails, the secondary controller comes online and takes over for the entire system. Active-passive array controllers do not provide any performance benefit, only redundancy (fault tolerance). *See also active–active array controllers.*

Array Two or more disks grouped together to appear as a single disk to the host system.

Array adapter A bus-based (usually PCI) hardware device — such as an add-in card, group of motherboard ASICs, or a combination of both — that converts the timing and protocol of a host's memory bus and an I/O bus. Usually used in entry-level servers, an array adapter also includes an on-board RAID coprocessor to offload most of the RAID operations — for example, secondary RAID 1 writes and RAID 5 parity calculations — from the host CPU. This is in contrast to the microprocessor-based array controllers used in midrange and high-end servers, which also offload I/O commands. Array adapters improve performance over software RAID solutions embedded within network operating systems such as NetWare and Windows NT. These adapters provide the same connectivity functions as a standard host adapter.

Bootable array An array which includes system disk files and allows a server to boot from the array while protecting the network operating system disk — and other data on the array — from drive failure.

CI/O - Comprehensive Input/Output Refers to Adaptec's CI/O Array Management Software.

Clustering The concept of using independent computer systems working together as a single logical system. Clustering is usually implemented to address both availability and scalability.

Cold swap Power must be switched off before the removal or insertion of a component.

Disk array *See Array*.

Disk/data striping Spreading data evenly over multiple disks to enhance performance. Sometimes referred to as RAID 0, data striping actually has no redundancy scheme and, therefore, does not provide any fault tolerance (data protection).

Downtime A time interval when a network cannot be used due to equipment failure or another cause.

Drive Synonym for disk, hard drive, hard disk, disk drive.

Duplexing Mirroring across two host adapters. Used only with software-based RAID storage systems (usually the embedded network operating system RAID software such as NetWare and Windows NT).

ECC (Error Correction Code) Refers to parity error detection and correction within memory or cache (for example, SIMM). Depending on the ECC SIMM, single-bit or double-bit parity errors can be detected but not corrected, or detected and corrected automatically.

Embedded SCSI ASIC (host adapter) A "chip" that is attached directly to the host motherboard and performs the same functions as a standard add-in host adapter.

Exclusive OR (XOR) A process based on a mathematical algorithm that is used by RAID levels 2, 3, 4, and 5 to compare computer data (binary 0s and 1s) created by a write request or by a read request during a drive failure. The result of the XOR process is parity information that will be stored along with data for real-time recovery in the event of a disk failure.

External array controller In contrast to bus-based array adapters and microprocessor-based array controllers, external array controllers reside in the external RAID storage enclosure. They connect to the host through a standard SCSI or serial (such as Fibre Channel) host adapter interface. These external controllers are similar to bus-based, microprocessor-based array controllers, in that they include an on-board microprocessor to offload all RAID functions (I/O commands and RAID operations) from the host CPU. They are usually used in midrange and high-end servers — especially in clustering environments.

Failed-drive mode A mode of reduced-performance operation that a disk array is in after a drive failure.

Failover The automatic replacement of a failed system component with a properly functioning one. Most often used in the context of redundant external array controllers. If one of the controllers fails, failover enables the second controller to take over the failed controller's I/O load. *See also active–active array controllers and active–passive array controllers.*

Fault tolerance The ability of a system to continue to perform its functions, even when one or more components have failed.

Fibre Channel High-speed, serial interface capable of supporting up to 100 MB/sec.

Host adapter A bus-based (PCI, EISA, ISA) hardware device, such as an add-in card or ASIC, that converts the timing and protocol of a host's memory bus and an I/O bus. *See also embedded SCSI ASIC and array adapter.*

Hot spare RAID storage feature that allows a spare drive (or other component) to be configured for automatic (in contrast to hot-swap) replacement and reconstruction in the event of a disk failure. Users can remain on-line and continue to access data. *Also see hot swap and warm swap.*

Hot swap A storage system's ability to allow the removal and replacement of a disk drive (or other component) while users are on-line and accessing data. In contrast to hot spare, this is a manual operation. Hot swap requires that the storage (or server) enclosure drive tray connectors be designed so that when a drive is removed, power is disconnected before the ground connection, and that the ground is restored before the power is reconnected upon reinsertion of the drive. This is usually accomplished by making the ground pin(s) in the drive tray connector slightly longer than the data pins. *See also hot spare and warm swap.*

I/O - Input/Output Refers to network user data requests and host overhead such as swapping and file system activity.

JBOD Just a bunch of drives. Refers to an array of drives without data redundancy.

MTBF Mean time between failure. Used to measure computer component average reliability/life expectancy. MTBF is not as well-suited for measuring the reliability of array storage systems as MTDL, MTTR or MTDA (see below) because it does not account for an array's ability to recover from a drive failure. In addition, enhanced enclosure environments used with arrays to increase uptime can further limit the applicability of MTBF ratings for array solutions.

MTDA Mean time between data access (or availability). The average time before non-redundant components fail, causing data inaccessibility without loss or corruption.

MTDL Mean time to data loss. The average time before the failure of an array component causes data to be lost or corrupted.

MTTR Mean time to repair. The average time required to bring an array storage subsystem back to full fault tolerance.

Member (disk) A disk that is in use as a member of a disk array.

Microprocessor-based array controller In contrast to an array adapter, a microprocessor-based array controller includes an on-board microprocessor (for example, an Intel i960) to offload I/O commands and RAID operations from the host CPU. Usually used in midrange and high-end servers. *See also array adapter.*

Mirroring Also known as RAID 1 or duplexing (when using two host bus adapters). Full redundancy is obtained by duplicating all data from a primary disk on a secondary disk. The overhead of requiring 100 percent data duplication can get costly when using more than two drives.

NOS - Network Operating System For example, NetWare, Windows NT Server, OS/2, SCO, and UnixWare.

PCI - Peripheral Component Interconnect Industry-standard specification that refers to a high-speed (133 MB/sec) host bus commonly used for host adapters, Ethernet adapters, and video cards.

Parity A form of data redundancy used by RAID levels 2, 3, 4, and 5 to recreate the data of a failed drive in a disk array.

RAID Redundant array of inexpensive disks. The term coined in 1987 by researchers at the University of California at Berkeley to describe a series of redundant architectures used in fault-tolerant disk arrays.

RAID Advisory Board (RAB) Industry organization of manufacturers and users of disk systems and related products whose mission is to educate users regarding all aspects of storage technology, and in particular, RAID technology.

RAID levels Numbered 0 through 5, RAID levels refer to different array architectures that offer various advantages in terms of data availability, cost, and performance.

Read-ahead cache A performance caching technique in which data is anticipated and read into the cache before it is actually requested. *See also write-back cache and write-through cache.*

Redundant A duplicate disk or component that provides a recovery path in case of a failure.

SCSI Small computer system interface (pronounced scuzzy). The fast, intelligent input/output parallel bus used by high-performance peripherals.

SIMM Single in-line memory module.

SLED Single large expensive disk. Refers to older mainframe hard disks that were used as a basis for comparison during the initial UC Berkeley RAID studies.

S.M.A.R.T. - Self-Monitoring, Analysis and Reporting Technology Drives equipped with this feature report predicted failures based on threshold values determined by the manufacturer. This allows the network manager to replace a drive before it fails.

Software-based array An array in which all management functions including parity calculation (XOR) are performed by the host server CPU. These products are low priced but have high CPU utilization and limited fault-tolerant features. High-performance, low-cost array adapters are quickly replacing these inferior software-based arrays.

System disk The disk (or array) on which a system's operating system is stored and from which it is initially loaded into system memory. *See also bootable array.*

Usable storage capacity Disk array capacity that is usable for data storage (vs. for mirroring or parity data). For example, under mirroring (RAID 1 and 0/1), usable storage remains a constant fifty percent (half of storage is always used for redundancy). This is in contrast to other RAID levels such as RAID 5, in which usable storage capacity is determined by the formula of "n-1". "n" is the total number of disk drives and "1" is the number of disks worth of capacity used for parity (redundancy) overhead. So, as the number of disks in the array grows, the usable storage capacity percentage increases in relation to parity (redundancy) information.

Warm swap The ability to remove and replace a disk drive while the power is on. All bus activity must be paused (usually done through a utility within the array management software) to maintain data integrity during removal or replacement. Typically used when hot swap is not supported by the server or storage enclosure drive tray. *See also hot spare and hot swap.*

Write-back cache A performance caching technique in which the completion of a write request is signaled as soon as the data is in cache; Actual writing to the disk occurs at a later time. Since the operating system is "fooled" into thinking that the write has actually been written to disk, there is a risk of losing or corrupting data in cache should an error or power failure occur. Therefore, use of a battery-backed cache is recommended to prevent such data loss/corruption. *See also read-ahead cache and write-through cache.*

Write-through cache A caching technique in which the completion of a write request is not signaled until data is safely stored on disk. Performance of write-through is essentially the same as in non-cached systems. *See also read-ahead cache and write-back cache.*

appendix a

Matching Server Type and Application with Adaptec Array Solutions

Server Type	Ideal Applications	Adaptec Array Solutions			
		AAA-130 Series	ARO-1130 RAID*port* Card	AEC-4312A	AEC-7312A
Entry-level • up to 75 active users • up to 5 disks • 1 or 2 CPUs • up to $5K+	• Small company enterprise • Large company workgroup/department • File and print services • Small databases • Internet/intranet	✓	✓		
Midrange • up to 500+ active users • up to 15 disks • 2 or 4 CPUs • up to $15K+	• Medium-large company enterprise • Medium-large databases • File and print services • Clustered servers • Video servers • Internet/intranet			✓	✓
High-end/Entry-level UNIX-based • up to 1000+ active users • > 15 disks • 4 CPUs • up to $25K+	• Large company enterprise • Large databases • Clustered servers • Video servers • File and print services • Internet/intranet			✓	✓

Adaptec has a scalable family of array solutions to match specific server types and their ideal application.

appendix b

Comparing RAID Levels

RAID Level	Minimum Number of Drives	Description	Strengths	Weaknesses
RAID 0	2	Data striping without redundancy	Highest performance	No data protection; One drive fails, all data is lost
RAID 1	2	Disk mirroring	Very high performance; Very high data protection; Very minimal penalty on write performance	High redundancy cost overhead; Because all data is duplicated, twice the storage capacity is required
RAID 2	Not used in LAN environments	No practical use	Previously used for RAM error correction (known as Hamming Code) and in disk drives before the use of embedded error correction	No practical use; Same performance can be achieved by RAID 3 at lower cost
RAID 3	3	Byte-level data striping with dedicated parity drive	Excellent performance for large, sequential data requests	Not well-suited for transaction-oriented network applications; Single parity drive does not support multiple, simultaneous read and write requests
RAID 4	3 (Not widely used)	Block-level data striping with dedicated parity drive	Data striping supports multiple simultaneous read requests	Write requests suffer from same single parity-drive bottleneck as RAID 3; RAID 5 offers equal data protection and better performance at same cost
RAID 5	3	Block-level data striping with distributed parity	Best cost/performance for transaction-oriented networks; Very high performance, very high data protection; Supports multiple simultaneous reads and writes; Can also be optimized for large, sequential requests	Write performance is slower than RAID 0 or RAID 1
RAID 0/1	4	Combination of RAID 0 (data striping) and RAID 1 (mirroring)	Highest performance, highest data protection (can tolerate multiple drive failures)	High redundancy cost overhead; Because all data is duplicated, twice the storage capacity is required; Requires minimum of four drives

Each RAID level has strengths and weaknesses. RAID 1, 5, and 0/1 are the most widely-used for transaction-based network server applications. RAID levels 0 and 3 are more commonly used in single-user environments like CAD/CAM or desktop digital video applications.

appendix C

Partial Listing of Array (RAID) Storage Enclosure Manufacturers

JMR Electronics, Inc.
20400 Plummer Street
Chatsworth, CA 91311
USA
Tel: (818) 993-4801
Fax: (818) 993-9173
Internet: www.jmr.com

Kingston Technology Corp.
17600 Newhope Street
Fountain Valley, CA 92708
USA
Tel (U.S.): (800) 435-0642
Fax (U.S.): (714) 438-1847
Tel (Intl): (714) 437-3334
Fax (Intl): (714) 438-1820
Internet: www.kingston.com

Trimm Technologies
350 Pilot Road
Las Vegas, NV 89119
USA
Tel: (800) 423-2024
Fax: (702) 361-6067
Internet: www.trimm.com

appendix d

Partial Listing of SCSI Disk Drive Manufacturers

Fujitsu Computer Products of America, Inc.
2904 Orchard Parkway
San Jose, CA 95134-2009
USA
Tel (U.S.): (800) 626-4686
Tel (Intl): (408) 432-6333
Internet: www.fujitsu.com

IBM (International Business Machines) Corp.
1 Old Orchard Road
Armonk, NY 10504
USA
Tel: (914) 765-1900
Internet: www.ibm.com

Quantum Corp.
500 McCarthy Boulevard
Milpitas, CA 95035
USA
Tel: (800) 624-5545
Internet: www.quantum.com

Seagate Technology, Inc.
920 Disc Drive
Scotts Valley, CA 95066
USA
Tel: (408) 438-6550
Fax: (408) 429-6356
Internet: www.seagate.com

appendix e

Cables and Terminators from Adaptec

Adaptec Cable Part #	Cable Type	Cable Length
Internal Cables		
ACK-F2F-5IT	Standard 50-pin internal SCSI ribbon cable with active termination — supports up to four 8-bit devices.	5 ft.
ACK-W2W-5IT	High-density internal Wide (68-pin) cable with active termination — supports up to four 16-bit disk drives.	1 meter (approx. 3 ft.)
External Cables		
ACK-H2L	External 50-pin low-density to 50-pin high-density (also called SCSI-2) cable.	3 ft.
ACK-W2W-E	High-density external Wide (68-pin) SCSI cable shielded with 68-pin male connector at both ends. Cable provides a screw attachment for a more secure connection.	1 meter (approx. 3 ft.)
Converter Cables		
ACK-GCH2L	This device connects low-density external cables to SCSI host adapters with 50-pin high-density connectors.	Not applicable
ACK-68P-50P-E	Fully electrically terminated external adapter to convert an external 68-pin connector to 50-pin connector. Attaches Fast SCSI devices to a Wide SCSI connector. Cable provides a screw attachment for a more secure connection.	Not applicable
ACK-68I-68E	68-pin internal to 68-pin external converter cable. Allows users to port an internal SCSI connector to an external vacant slot in a PC. This enables users to utilize an internal connection to attach external peripherals including disk arrays.	15 in.
ACK-50I-50E	50-pin internal to 50-pin external converter cable. Allows users to port an internal SCSI connector to an external vacant slot in a PC. This enables users to utilize an internal connection to attach external peripherals including disk arrays.	15 in.

These high-quality SCSI cable kits are certified and guaranteed by Adaptec. High-quality cables help reduce cable-related problems.

To Order Cables:

Adaptec, Inc. (USA and Canada)
Tel: (800) 442-7274, or
(408) 957-7274

Adaptec Europe
Tel: 32 11 300 379

appendix f

Adaptec Array Solutions:
Major Features

Features	AAA-130 PCI Array Adapter Series		ARO-1130 PCI RAID*port* Card	AEC-4312A External Array Controller	AEC-7312A External Array Controller
	AAA-131	AAA-133			
Array solution type	Bus-based	Bus-based	Bus-based	External	External
Ideal server type	Entry-level	Entry-level	Entry-level	Midrange and High-end	Midrange and High-end
Host interface	PCI	PCI	PCI	SCSI	Fibre Channel
Device interface	SCSI	SCSI	SCSI	SCSI	SCSI
SCSI, Fast/Wide SCSI, Ultra SCSI, Wide Ultra SCSI support	YES	YES	YES	YES	YES
Number of SCSI channels	1	3	#	2	2
Both 50-pin and 68-pin connectors on internal channels	YES	YES	#	Not applicable	Not applicable
Disk and non-disk support	YES	YES	YES	YES	YES
On-board RAID coprocessor	YES	YES	YES	YES	YES
Multiple operating system support	YES	YES	YES	YES	YES
Data striping (RAID 0)	YES	YES	YES	YES	YES
Mirroring (RAID 1)	YES	YES	YES	YES	YES
Data striping with distributed parity (RAID 5)	YES	YES	YES	YES	YES
Hot-swap drive support	YES	YES	YES	YES	YES
Hot-spare drive support	YES	YES	YES	YES	YES
Array*Config* initial boot array installation software	YES	YES	YES	YES	YES
Array1000 Family Manager Set software	YES	YES	YES	OS independent, not needed	OS independent, not needed
Adaptec CI/O Array Management software	YES	YES	YES	##	##

motherboard dependent
available in subsequent releases

The Adaptec family of array solutions have different hardware implementations to meet specific server/network requirements, but share many of the same advanced features.

Adaptec, Inc.
691 South Milpitas Boulevard
Milpitas, California 95035
Tel: (408) 945-8600
Fax: (408) 262-2533

Adaptec Asia — Singapore
Tel: (65) 278-7300
Fax: (65) 273-0163

Adaptec Europe — Belgium
Tel: (32) 2-352-34-11
Fax: (32) 2-352-34-00

Adaptec Japan — Tokyo
Tel: (81) 3-5365-6700
Fax: (81) 3-5365-6950

Additional Sales Offices

France
Tel: (33) 1-3452-3434
Fax: (33) 1-3452-3432

Germany
Tel: (49) 89-4564060
Fax: (49) 89-45640615

United Kingdom
Tel: (44) 1276-854500
Fax: (44) 1276-854505

Latin America (Miami)
Tel: (305) 265-0399
Fax: (305) 265-1387

Literature:
1-800-934-2766 (USA and Canada) or (510) 732-3829

Pre-sales Support:
1-800-442-7274 (USA and Canada) or (408) 957-7274

Interactive Fax:
(408) 957-7150

World Wide Web:
www.adaptec.com/RAID

Internet ftp server:
ftp.adaptec.com

Adaptec USA Bulletin Board Service (BBS):
(408) 945-7727
(up to 28,800 bps, using 8 bits, 1 stop bit, no parity)

P/N 950062-011 7/97 Printed in U.S.A.