HP Smart Array Controller technology

technology brief, 2nd edition

Abstract

This technology brief describes specific functions of the HP Smart Array controller family and explains how Smart Array technology meets administrators' requirements for capacity growth, high performance, data availability, and manageability.

Introduction

In today's networking environments, administrators face difficult online data storage problems and ever-increasing demands for performance. While PCI Express (PCIe) provides a high speed interface for data communication, the HP family of parallel SCSI, Serial Attached SCSI (SAS), and Serial ATA (SATA) Smart Array controllers addresses direct attach and storage area network (SAN) requirements. References made in this technology brief to the 'present generation' of Smart Array controllers refers to those SAS-based, PCIe 2.0 compliant Smart Array controllers released in the first half of 2009.

HP Smart Array controllers provide administrators with configuration, management, and diagnostics tools. These tools provide high levels of usability as well as consistency between generations of products. This continuity ensures that administrator can move data between servers and external storage enclosures, and between models of Smart Array controllers.

For complete Smart Array controller compatibility and support information, see: www.hp.com/products/smartarray

Storage trends

Four key trends influence network storage requirements today:

- Application complexity—Network applications are becoming more complex, leading to larger files and the need to maintain more information online for immediate user access.
- Mission-critical migration—More mission-critical data is moving to servers, either down from larger systems or up from paper-based processes.
- Server consolidation—Multiple servers and applications are being consolidated onto fewer servers for increased control and centralized network management.
- Efficiency expectations—Corporate information technology organizations are managing more applications and more data while at the same time reducing the size of their administrative staffs.

These key trends generate four primary data storage requirements:

- Capacity growth—Storage solutions must provide not only adequate capacity for today's applications but also flexibility for future growth.
- High performance—Data storage subsystems must deliver enough performance to accommodate an increasing number of users while maintaining rapid response times. In many environments, the storage subsystem is the most critical determinant of overall system performance.
- Data availability—Because businesses depend on their mission-critical data, it must be accessible at all times to maintain user productivity.
- Manageability—Storage solutions must reduce the total cost of ownership by making network storage management intuitive, informative, and less time consuming.

Meeting data storage requirements

The HP Smart Array controller family has an advanced intelligent architecture developed with a feature set that specifically addresses today's network data storage requirements:

- Capacity growth—Administrators can expand server storage capacity while the server is operating. The HP online Array Configuration Utility (ACU) simplifies storage management for administrators who use fault-tolerant array configurations and experience capacity growth. Refer to appendix A for details about using array creation and expansion to increase capacity.
- High performance—Smart Array controllers supply high levels of throughput and data bandwidth to support demanding workloads. The advanced intelligent architecture contains specialized modules that enhance performance significantly. Online array expansion, stripe size migration, RAID migration, configurable read/write cache sizes, rebuild priority, and configurable write-back caching for each logical drive empowers administrators to tune Smart Array controllers to meet performance demands.
- Data availability—Smart Array controllers include a comprehensive set of fault-tolerance features that guard against disk drive and controller component failures, making online storage administration possible and reducing planned downtime.
- Manageability—All Smart Array controllers use a common set of management software, reducing training requirements and technical expertise necessary to install and maintain HP server storage. The common data format among Smart Array controllers means customers can easily upgrade to future Smart Array products.

High performance

Smart Array controllers offer exceptional performance and reliability characteristics, with support for traditional parallel Small Computer System Interface (SCSI) technology as well as the latest SAS, SATA technology, and advanced RAID capabilities.

PCI Express technology

The present generation of Serial Attached SCSI (SAS) based Smart Array controllers released in the first half of 2009, and all ProLiant G6 servers support the PCIe 2.0 specification. PCIe 2.0 has a per-lane signaling rate of 5 Gb/s —double the per-lane signaling rate of PCIe 1.0 (Figure 1).

Figure 1. PCIe data transfer rates

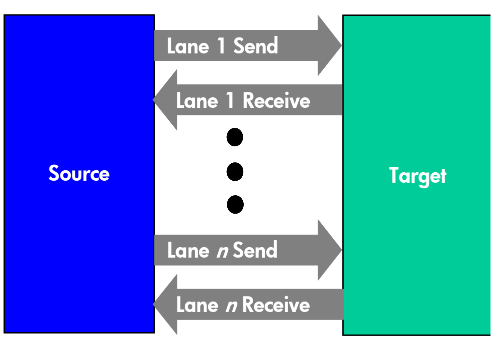

Link size	Max. bandwidth (Send or receive)		Total (Send and receive)	
	PCIe 1.0	**PCIe 2.0**	**PCIe 1.0**	**PCIe 2.0**
x1	250 MB/s	500 MB/s	500 MB/s	1 GB/s
x4	1 GB/s	2 GB/s	2 GB/s	4 GB/s
x8	2 GB/s	4 GB/s	4 GB/s	8 GB/s
x16	4 GB/s	8 GB/s	8 GB/s	16 GB/s

PCIe 2.0 is completely backward compatible with PCIe 1.0. A PCIe 2.0 device can be used in a PCIe 1.0 slot and a PCIe 1.0 device can be used in a PCIe 2.0 slot. Table 2 shows the level of interoperability between PCIe cards and PCIe slots.

Table 1. PCIe device interoperability

PCIe device type	x4 Connector x4 Link	x8 Connector x4 Link	x8 Connector x8 Link	x16 Connector x8 Link	x16 Connector x16 Link
x4 card	x4 operation	x4 operation	x4 operation	x4 operation	x4 operation
x8 card	Not allowed	x4 operation	x8 operation	x8 operation	x8 operation
x16 card	Not allowed	Not allowed	Not allowed	x8 operation	x16 operation

SAS/SATA technology

The newest serial, PCIe 2.0 capable Smart Array controllers use Serial Attached SCSI (SAS) technology, a point-to-point architecture in which each device connects directly to a SAS port rather than sharing a common bus as with parallel SCSI devices. Point-to-point links increase data throughput and improve the ability to locate and fix disk failures. More importantly, SAS architecture solves the parallel SCSI problems of clock skew and signal degradation at higher signaling rates.[1]

The same Smart Array controllers are compatible with Serial Advanced Technology Attachment (SATA) technology and include the following features to enhance performance and maintain data availability and reliability:

- SAS and SATA compatibility — The ability to use either SAS or SATA hard drives lets administrators deploy drive technology that fits each computing environment. HP Smart Array controllers can manage both SAS arrays and SATA arrays. Smart Array configuration utilities help administrators configure arrays correctly so that data remains available and reliable.
- SAS wide port operations — Wide ports contain four single lane (1x) SAS connectors and the cabling bundles all four lanes together. SAS wide ports allow balanced SAS traffic distribution across the links for enhanced performance. In addition, wide ports provide redundancy by tolerating up to three physical link failures while maintaining the ability to communicate with the disk drives. The most common use of wide links is to a JBOD or to an internal server expander connecting to large numbers of drives. No special configuration is required for this functionality.
- SAS expanders — Low-cost, high-speed switches called expanders can combine multiple single links to create wide ports and increase available bandwidth. SAS expander devices also offer higher system performance by expanding the number of hard drives that can be attached to an HP Smart Array controller. SAS expanders are an aggregation point for large numbers of drives or servers providing a common connection. By cascading expanders, administrators can chain multiple storage boxes together. For more information on the HP SAS Expander Card, go to http://h18004.www1.hp.com/products/servers/proliantstorage/arraycontrollers/sas-expander/index.html.

SAS-2 standard

The second-generation SAS (SAS-2) link speed [2] of 6 Gb/s is double the SAS-1 transfer rate. SAS-2 link speeds require SAS-2 compliant hard drives. SAS-2 eliminates the distinction between fanout and edge expanders by replacing them with self-configuring expanders. SAS-2 enables zoning for enhanced resource deployment, flexibility, security, and data traffic management.

SAS-2 connections have the potential to deliver peak raw data bandwidth of up to 600 megabytes per second (MB/s) per physical link in each direction. SAS-2 devices are capable of sending and

[1] For more information about SAS technology, refer to the HP technology brief titled "Serial Attached SCSI storage technology" available at http://h20000.www2.hp.com/bc/docs/support/SupportManual/c01613420/c01613420.pdf .

[2] Serial Attached SCSI-2 (SAS-2) is an American National Standards Institute (ANSI) standard from the INCITS T10 Technical Committee on SCSI Storage Interfaces. SAS-2 is the successor to SAS-1.1 and SAS-1.

receiving data simultaneously across each physical link, which is known as full duplex. When effectively implemented, full duplex, 6 Gb/s SAS connections can deliver peak raw data bandwidth of up to 1200 MB/s between the controller and storage device. It is important to note that the SAS-2 data bandwidths described here are theoretical speeds identified by the SAS-2 standard. Real world performance will be affected by the performance of storage devices attached to the SAS-2 connection.

Smart Array controllers, with releases beginning in the first quarter of 2009, incorporate SAS-2 connections. The SAS-2 standard is compatible with both Serial SCSI and Serial ATA protocols for communicating commands to SAS and SATA devices. SAS-2 compliant controllers are fully compatible with 1.5 Gb/s and 3 Gb/s SATA technology.

For an up to date listing of HP Smart Array controllers that support the SAS-2 specification, see the Smart Array controller matrix: www.hp.com/products/smartarray

Mini SAS 4x cable connectors and receptacles

Mini SAS 4x connectors and receptacles (Figure 2) are replacing SAS 4x connectors and receptacles in present generation Smart Array controllers. The ground pins in Mini SAS connectors can be used for power in active cables.

Figure 2. The Mini SAS 4i connector (left) and receptacle (right) are replacing SAS 4x connectors and receptacles.

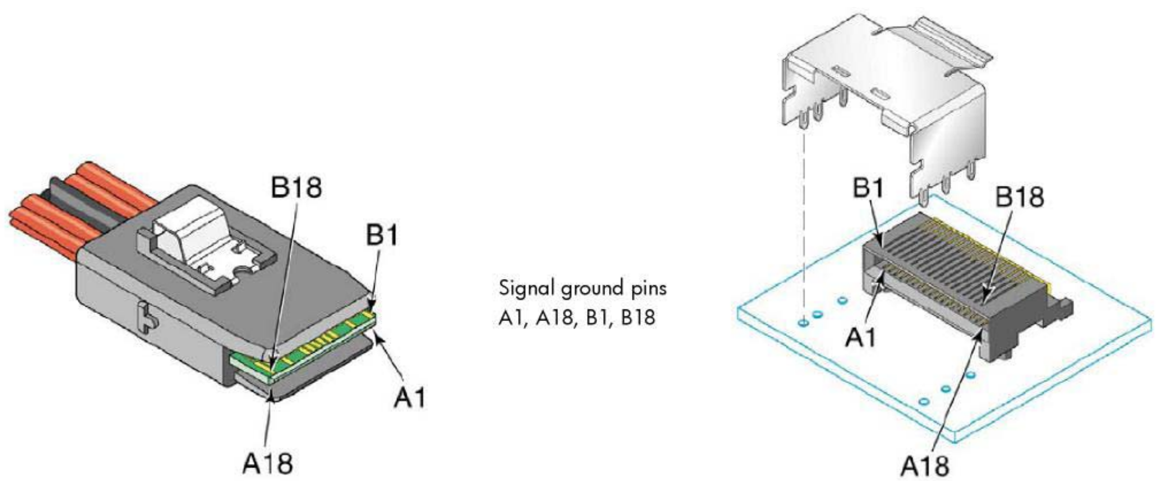

For more detailed information on SAS technology and SAS-2 zoning, refer to the Serial Attached SCSI storage technology brief:
http://h20000.www2.hp.com/bc/docs/support/SupportManual/c01613420/c01613420.pdf

High-performance processor

HP Smart Array controllers use a variety of high-performance processors for managing the RAID system. The Power PC and MIPS processors are the most widely used among Smart Array controllers. The PowerPC processor is a 64-bit processor based on reduced instruction set computer (RISC) technology. The processor connects to its internal peripherals using separate 128-bit-wide read-and-write buses, each running at 133 MHz.

Power PC processors have a highly pipelined architecture that allows dual instruction fetch, decode, and out-of-order issue, as well as out-of-order dispatch, execution, and completion. The Power PC flexibility features include independently configurable data cache arrays, write-back and write-through operation, and performance characteristics that increase cache memory allocation efficiency. Designed for extensive power management and maximum performance, Power PC processors provide increased throughput and a streamlined path to instruction completion, which means faster user access to data.

The Smart Array P410, P411, and P212 controllers use an embedded MIPS storage processor. This reduced instruction set computing (RISC) instruction set architecture (ISA) processor operates at 600 MHz, has 34K instruction and data caches, and a 32-bit multi-threading I/O. The MIPS processor uses a 4-way set associative write-back, and flexible thread policy manager with programmable quality of service (QoS) support.

Cache module benefits

With advanced read-ahead and write-back caching capabilities, the Smart Array controller cache module produces significant performance improvements for I/O operations.

Read-ahead caching

The HP Smart Array controller family uses an intelligent read-ahead algorithm that can anticipate data needs and reduce wait time. It can detect sequential read activity on single or multiple I/O threads and predict when sequential read requests will follow. The algorithm then "reads ahead," or pre-fetches data, from the disk drives before the data is actually requested. When the read request occurs, the controller retrieves the data from high-speed cache memory in microseconds rather than from the disk drive in milliseconds.

This adaptive read-ahead scheme provides excellent performance for sequential small block read requests. At the same time, the system is not penalized by random read patterns because read-ahead functionality is disabled when non-sequential read activity is detected. Thus, Smart Array controller technology overcomes the problem with some array controllers in the industry today that use fixed read-ahead schemes to increase sequential read performance but degrade random read performance.

Write-back caching

HP Smart Array controllers also use a caching scheme that allows host applications to continue without waiting for write operations to complete to the disk. This technique is also called posted-writes, or write-back caching. Without this type of caching, the controller must wait until write data is actually written to disk before returning completion status to the operating system. With write-back caching, the controller can "post" write data to high-speed cache memory and immediately return "back" completion status to the operating system. The write operation is completed in microseconds rather than milliseconds.

Once write data is located in the cache, subsequent reads to the same disk location will be sourced from the cache. Subsequent writes to the same disk location will replace the data held in cache. This technique is called a "read cache hit." This feature improves bandwidth and latency for applications that frequently write and read the same area of the disk.

After write data is located in the cache, the controller finds opportunities to combine adjacent write transfers into a larger transfer that can efficiently move to the hard drive. This technique is called "write coalescing." Hard drives achieve higher throughput when processing fewer large transfers instead of many small transfers. This feature improves write bandwidth whenever data is sequentially written to the controller using write command sizes smaller than the stripe size of the logical drive.

With write data located in the cache, logical drives in RAID 5 and RAID 6 with advanced data guarding (ADG) configurations achieve higher write performance by combining adjacent write requests to form a full stripe of data. This technique is called "full-stripe writes." Write operation for

RAID 5 and RAID 6 with ADG normally requires extra disk reads to compute the data for the parity drives. However, if all the data required for a full stripe is available in the cache, the controller does not require the extra disk reads. This feature improves write bandwidth whenever data is sequentially written to a logical drive in a RAID 5 or RAID 6 with ADG configuration.

Data in the controller's write cache is written to disk later, at an optimal time for the controller. While the data is in cache, it is protected against memory chip failure by error checking and correction (ECC) DRAM technology and against system power loss by the integrated battery backup mechanism. Smart Array controllers avoid the risk of data loss by ensuring that the battery backup is present before enabling write-back cache. Hard drives provide an option to enable write-caching that is not battery backed. HP advises against enabling hard drive write cache because a power or equipment outage could result in data loss.

Balanced cache size

Smart Array controllers allow administrators to adjust how the cache is distributed for write-back and read-ahead operations. Administrators can configure the cache module for optimal performance for any storage need. The default setting configures the cache for 50 percent write-back operations and 50 percent read-ahead operations. Additionally, the cache module capacity can be upgraded to increase caching performance.

RAID performance enhancements

Smart Array controllers use several enhancements to increase RAID performance.

Disk striping

Striping combines several individual disk drives into a larger disk array containing one or more logical drives. Performance of the individual drives is aggregated to provide a single high-performance logical drive. The array controller evenly distributes the logical drive data into small "stripes" of data sequentially located across each member disk drive in the array. Administrators can adjust the stripe size to achieve optimal performance. Performance improves as the number of drives in the array increases.

Parity data

In a RAID 5 configuration, data protection is provided by distributed parity data. This parity data is calculated stripe by stripe from the user data that is written to all other blocks within that stripe. The blocks of parity data are distributed evenly over every physical drive within the logical drive. When a physical drive fails, data that was on the failed drive can be calculated from the remaining parity data and user data on the other drives in the array. This recovered data is usually written to an online spare drive through a process called a rebuild.

RAID 6 (Advanced Data Guarding- ADG), like RAID 5, generates and stores parity information to protect against data loss caused by drive failure. With RAID 6 (ADG), however, two different sets of parity data are used so that data can still be preserved even if two drives fail. Each set of parity data uses a capacity equivalent to that of one of the constituent drives. This method is most useful when data loss is unacceptable but cost is also an important factor. RAID 6 (ADG) provides better protection for data than a RAID 5 configuration because of the additional parity information.

Background RAID creation

When a RAID 1, RAID 5, or RAID 6 logical drive is first created, the Smart Array controller must build the logical drive within the array before enabling certain advanced performance techniques. While the logical drive is created, the storage volume is accessible by the host with full fault tolerance. The Smart Array controller creates the logical drive whenever the controller is not busy; this is called background parity initialization. Parity initialization takes several hours to complete, depending on the size of the logical drive and how busy the host keeps the controller. Before parity initialization completes, normal writes to RAID 5 and RAID 6 logical drives are slower because the controller must read the entire stripe to update the parity data and maintain fault tolerance. These writes during parity initialization are called regenerative writes or reconstructed writes.

RAID 5 and RAID 6 read-modify-write

After parity initialization is complete, writes to a RAID 5 or RAID 6 logical drive are typically faster because the controller does not read the entire stripe to update the parity data. Since the controller knows that the parity data is consistent with all the member drives in the stripe, the controller needs to read from only two hard drives during a RAID 5 write (or three hard drives for a RAID 6 write) to compute the parity data (regardless of array size). This technique is called a read-modify-write or backed-out writes.

Striping across arrays

RAID 50 and 60 methods stripe the data across multiple RAID/JBOD sets with different levels of parity. These nested RAID types allow users to configure arrays across HP Modular Smart Arrays (MSAs).

RAID 50 (RAID 5+0) is a nested RAID method that uses RAID 0 block-level striping across RAID 5 arrays with distributed parity. RAID 50 will tolerate one drive failure in each spanned array without loss of data. RAID 50 configurations require a minimum of six drives and require less rebuild time than single RAID 5 arrays.

RAID 60 (RAID 6+0) is a nested RAID method that uses RAID 0 block-level striping across multiple RAID 6 arrays with dual distributed parity. With the inclusion of dual parity, RAID 60 will tolerate the failure of two disks in each spanned array without loss of data. RAID 60 configurations require a minimum of eight drives.

RAID 6 and 60 are available as an option with the Smart Array Advanced Pack (see section later in this paper) and are not supported on all HP Smart Array controllers.

RAID 1 load balancing

In general, the same stripe and array sizes RAID 0, RAID 5, and RAID 6 with ADG logical drives have the same read performance. RAID 1 logical drives contain two copies of the data. During reads to RAID 1 logical drives, the Smart Array controller issues read requests to either drive in the mirrored set. During a heavy read load, the Smart Array controller balances the number of requests between the two hard drives to achieve higher read bandwidth. This technique is called RAID 1 load balancing.

Hardware RAID

Today's operating systems offer basic RAID 0, RAID 1, and RAID 5 disk management functions, called software RAID, to create logical drives that do not contain the operating system. Software RAID requires a significant amount of the server's resources to perform management functions. However, Smart Array controllers use a separate processor and memory subsystem for management functions. Furthermore, for Smart Array controllers the parity calculations required by RAID 5 and RAID 6 with ADG are performed by specialized hardware engines that maximize data throughput for disk write, rebuild, and regenerate read operations.

Smart Array performance

Figure 3 displays the results of HP testing for sequential read/write performance between Smart Array controller SA-P411, announced in April 2009, and the previous generation SA-P800. The data in Figure 3 is the result of testing on a ProLiant G5 DL380 server in a RAID 5 environment using the default 64KB stripe size. With the two controllers operating under the same constraints of 3Gb SAS and PCI-e Gen 1, the sequential write performance of the SA-P411was approximately 200 percent higher than that of the SA-P800. The significant performance gains displayed the Smart Array P411 are a result of improvements to the Smart Array firmware and hardware.

Figure 3. Sequential read/write performance for the Smart Array P800 and P411 controllers in a RAID 5 environment

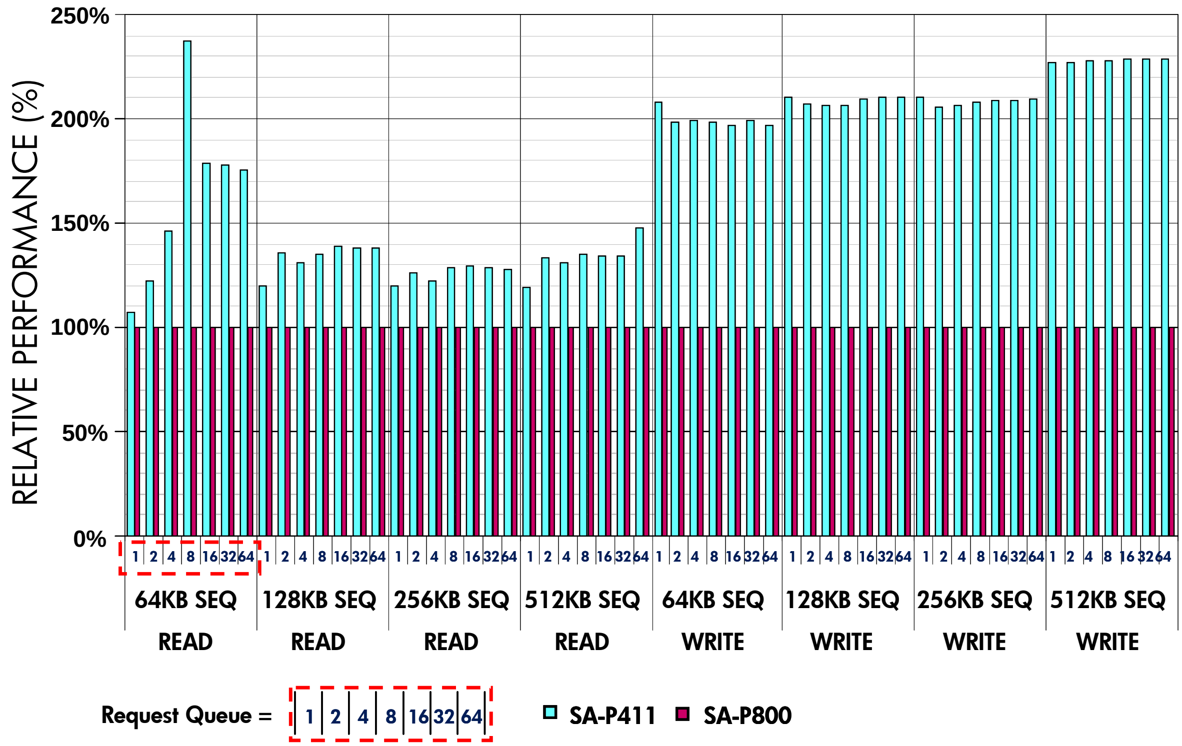

The max I/O performance test results displayed in Figure 4 also show significant performance gains for the Smart Array P411 over the previous-generation P800 controller. Performance was similar in both read and write tests.

Figure 4. Max I/O read/write performance for the SA- P800 and SA-P411 controllers in a RAID 5 environment

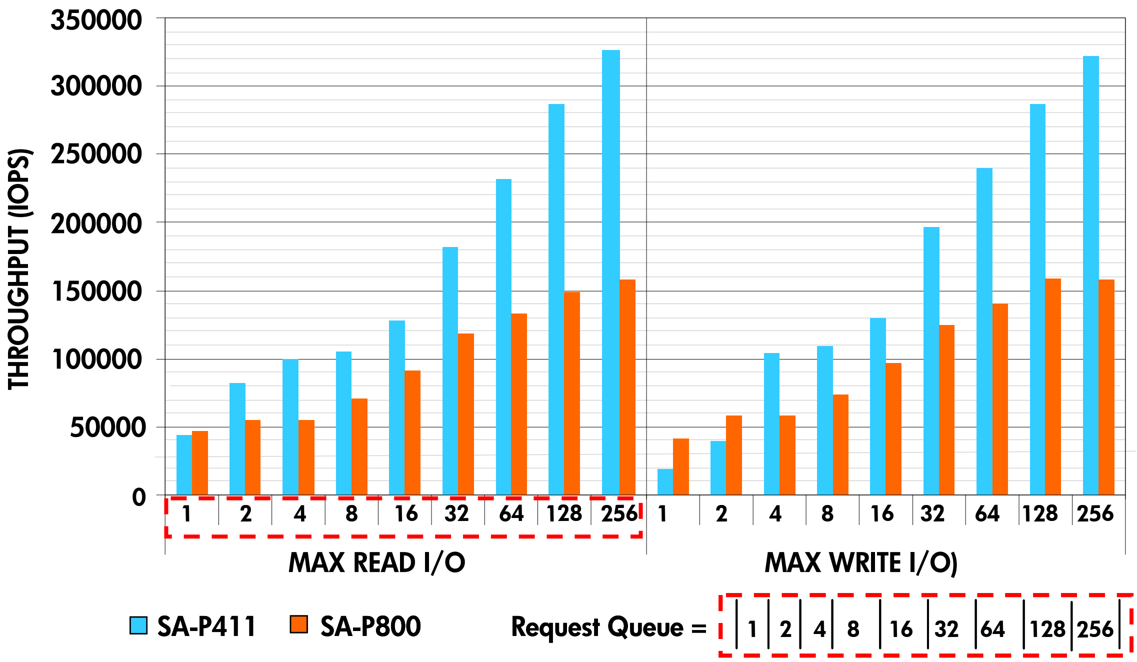

The hardware and software specifications for the test results in Figures 3 and 4 were as follows:

- All testing used ProLiant DL380 G5 servers
- Smart Array P411 Firmware specifications: v1.62; Memory: 512 MB; Driver: HP Smart Array Controller Driver HpCISSs2 6.14.0.32 B25, RAID 5 (64 KB); Internal hard drives: 2 logical drives; External storage: 2 MSA60 with 12+12-DRV; Drives: SEAGATE 15K 146GB large form factor SAS (HPDA- firmware version on the hard drives being tested)
- Smart Array P800 Firmware specifications: v5.26; Memory:512 MB; Driver: HpCISSs 2 6.14.0.32 B25, RAID 5 (64 KB); Internal hard drives: 2 logical drives; External storage: 2 MSA60 with 12+12-DRV; Drives: SEAGATE 15K 146GB large form factor SAS (HPDA)

Data availability

Smart Array controllers support online array expansion, logical drive extension, stripe size migration, and RAID migration. These technologies protect data and allow network administrators to modify the array without interrupting user access. Smart Array controllers can monitor I/O activity, track key parameters, predict potential problems, take corrective actions, provide automatic recovery, and deliver full fault management to protect against downtime.

RAID support

In 1989, HP (then Compaq) was the first company to introduce RAID subsystems in the network server marketplace. Today, RAID is an industry-standard technology and most online network data storage is protected with some level of RAID.

Smart Array controllers support RAID levels 0, 1, 1+0, 5, 6 with ADG, 50, and 60. Today, only Smart Array controllers offer RAID 6 with ADG, which provides higher fault tolerance than RAID 5 with lower implementation costs than RAID 1+0 and greater usable capacity per unit density than RAID 1.

Stripe size migration

Using the ACU, administrators can change the stripe size of a logical drive without any downtime. Stripe size migration can improve logical drive in response to changes in the read/write environment. Each RAID level has a default value designed to provide good performance across many types of applications. Table 2 shows the default values and the range of stripe sizes available for each RAID level. Table 2 provides a generic overview of stripe sizes for the listed RAID levels and is generally applicable to the present generation of Smart Array controllers. These values can change for certain controllers and storage device implementations. Users should consult configuration documentation for the specific controller and storage device in question.

Table 2. Stripe sizes available per RAID level

Fault-tolerance level	Default in KB	Available stripe sizes (KB)
RAID 0	128	8, 16, 32, 64, 128, 256
RAID 1 or 1+0	128	8, 16, 32, 64, 128, 256
RAID 5	64	8, 16, 32, 64, 128, 256
RAID 6 with ADG *	16	8, 16, 32, 64, 128, 256
RAID 50	128	8, 16, 32, 64, 128, 256, 512
RAID 60 *	16	8, 16, 32, 64, 128, 256

* Requires a valid SAAP license key

Certain applications, especially those performing mostly one type of transaction, write transactions for example, may require adjusting stripe size to achieve optimal performance. Table 3 lists recommended stripe sizes for general types of server applications. Administrators can start with these general recommendations and then fine tune to determine the best overall performance for a particular application. The ACU allows administrators to make these changes online without disruption or data loss.

Table 3. Recommended stripe sizes

Type of server application	Suggested stripe size change
Mixed read/write	Accept default value
Mainly sequential read (such as audio/video applications)	Larger stripe sizes
Mainly write (such as image manipulation applications)	Smaller stripes for RAID 5, RAID 6 with ADG Larger stripes for RAID 0, RAID 1, RAID 1+0

NOTE:
Different controllers may have different stripe sizes. Users should consult the controller user guide for more details.

RAID migration

Using the ACU, administrators can also change the RAID level of the logical drive without down time. Administrators can perform RAID migration to increase raw data storage capacity, improve performance by increasing the number of spindles in a logical drive, or change fault-tolerance (RAID) configurations. Table 4 summarizes RAID levels and the amount of space required for each type of fault tolerance.

Table 4. Summary of RAID methods

	RAID 0 (striping)	**RAID 1 (mirroring)**	**RAID 1+0 (striping and mirroring)**	**RAID 5 (distributed data guarding)**	**RAID 6 with ADG (advanced data guarding)**
Usable drive space*	100%	50%	50%	67% to 93%	50% to 96%
Usable drives	N	N/2	N/2	N-1	N-2
Minimum number of drives	1	2	4	3	4
Tolerant of single drive failure?	No	Yes	Yes	Yes	Yes
Tolerant of multiple simultaneous drive failures?	No	No	Yes, if failed drives not mirrored to each other	No	Yes
Read performance	High	High	High	High	High
Write performance	High	Medium	Medium	Low	Low
Relative cost	Low	High	High	Medium	Medium

* The values for usable drive space are calculated assuming a maximum of 14 physical drives of the same capacity (or a maximum of 56 for RAID 6 with ADG) with no online spares. HP recommends not exceeding these maximum figures (excluding any allowable online spares) when configuring a drive array due to the increased likelihood of logical drive failure with more hard drives.

Drive roaming

HP Smart Array controllers support drive roaming, which allows administrators to move hard drives and arrays while maintaining data availability. Drive roaming allows administrators to move one or more hard drives that are members of a configured logical drive to a different bay position as long as the new bay position is accessible by the same controller. In addition, it allows administrators to move a complete array from one controller to another, even if controllers are in different servers.

Mirror splitting and recombining

Mirror splitting is a task that splits an array with one or more RAID 1 or RAID 1+0 logical drives into two identical new arrays with RAID 0 logical drives. This is useful for administrators who want to replicate a configuration or need to build a backup before performing a risky operation. Using the ACU, administrators can also recombine a split mirrored array. Beginning with the present generation of SAS-based Smart Array controller controllers, the Smart Array Advanced Pack (SAAP) is required to use mirror splitting and recombining. For support information regarding mirror splitting on specific controllers, see the controller Quickspecs. More information about mirror splitting can be found in the technology brief "RAID 1(+0): breaking mirrors and rebuilding drives" at:

http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00378986/c00378986.pdf

NOTE:
An array cannot be split if it contains logical drives in RAID 0, RAID 5, or RAID 6 with ADG configurations. An array can be split or re-mirrored only when the server is offline and operating in the standard configuration mode of the ACU. When a split mirrored array is recombined, all data on the second array is destroyed. For more information, refer to the product documentation.

Online drive flash

Serial-based Smart Array controllers support online drive flashing, which saves time when updating firmware. Instead of taking the hard disk drive (HDD) offline before loading a new firmware image, administrators can download an updated HDD firmware image to the supported Smart Array controller and update all of the HDDs the next time the server is rebooted. This greatly reduces the time involved in updating disk drive firmware.

Recovery ROM

HP Smart Array controllers store a redundant copy of the controller firmware image to protect against data corruption. If the active firmware image becomes corrupt, Smart Array controllers use the redundant firmware image and continue operating. The recovery read-only memory (ROM) provides protection against power outages during firmware flashing.

Pre-Failure Warranty using S.M.A.R.T technology

HP (then Compaq) pioneered failure prediction technology for hard disk drives by developing monitoring tests run by Smart Array controllers. Called monitoring and performance (M&P), or drive parameter tracking, these tests externally monitor hard drive attributes such as seek times, spin-up times, and media defects to detect changes that could indicate potential failure.

HP worked with the hard drive industry to help develop a diagnostic and failure prediction capability known as Self-Monitoring Analysis and Reporting Technology (S.M.A.R.T.). Over the years, as S.M.A.R.T. matured, HP used both M&P and S.M.A.R.T. to support hard drive failure prediction technology for Pre-Failure Warranty replacement of hard drives.

S.M.A.R.T. has matured to the point that HP relies exclusively on this technology for hard drive failure prediction to support Pre-Failure Warranty. Since 1997, all HP SCSI, SAS, and SATA server-class hard drives have incorporated S.M.A.R.T. technology. S.M.A.R.T. hard have the capacity to inform the host when a hard drive is experiencing abnormal operation likely to lead to drive failure.

S.M.A.R.T. improves failure prediction technology by placing monitoring capabilities within the hard disk drive. These monitoring routines are more accurate than the original M&P tests because they are designed for a specific drive type and have direct access to internal performance, calibration, and error measurements. S.M.A.R.T. uses internal performance indicators and real-time monitoring and analysis to improve data protection and fault prediction capability beyond that of the original M&P tests. In addition, HP Smart Array controllers proactively scan the hard drive media during idle time and repair, or report, any media defects detected.

S.M.A.R.T. can often predict a problem before failure occurs. HP Smart Array controllers recognize S.M.A.R.T. error codes and notify HP Systems Insight Manager[3] (SIM) whenever a potential problem arises. HP SIM, in turn, immediately notifies administrators of drive failures. The drive parameter tracking feature allows Smart Array controllers to warn of potential drive problems before they occur.

[3] For a more detailed discussion of HP Systems Insight Manager, refer to the "HP Systems Insight Manager" section.

HP drives that fail to meet the expected criteria may be eligible for replacement under the HP Pre-Failure Warranty.

Automatic data recovery with rapid rebuild technology

When a hard drive in an array is replaced, Smart Array controllers use the fault-tolerance information on the remaining drives in the array to reconstruct the missing data and write it to the replacement drive. This process is called automatic data recovery, or rebuild. If fault tolerance is compromised, the data cannot be reconstructed and is likely to be permanently lost.

The latest generation of HP Smart Array controllers includes rapid rebuild technology for accelerating the rebuild process. Rapid rebuild technology comprises several features that result in substantially quicker rebuild time. Faster rebuild time reduces the risk of logical drive failure by restoring logical drives back to full fault tolerance before a subsequent drive failure can occur.

Generally, a rebuild operation requires approximately 15 to 30 seconds per gigabyte for RAID 5 or RAID 6 with ADG logical drives. However, actual rebuild time depends on several factors, including the amount of I/O activity occurring during the rebuild operation, the number of disk drives in the logical drive, the rebuild priority setting, and the disk drive performance. The ACU allows administrators to view progress of the rebuild and set the priority for the rebuild operation.

Online spare

Smart Array controllers enable administrators to designate an unlimited number of drives as online spares to arrays containing one or more fault-tolerant logical drives. The same spare drive can be assigned to multiple arrays as a global spare. Smart Array configuration utilities ensure that SAS hard drives can only be assigned as spares for SAS arrays (and likewise, SATA hard drives for SATA arrays). During system operation, these spare drives remain up and running but not active; that is, no I/O operations are performed to them during normal array operation. Spare drives are held in reserve in case one of the active drives in the array should fail, and then an online spare drive is selected as the replacement hard drive.

If an active drive fails during system operation, the controller automatically begins rebuilding each fault-tolerant logical drive onto the online spare; no administrator action is required. Once the rebuild operation is complete, the system is fully fault-tolerant once again. The failed drive can be replaced at a convenient time. Once an administrator installs a replacement drive, the controller will restore data automatically from the failed drive to the new drive. At that point, the original online global spare, as shown in Figure 5, will return to standby mode.

Figure 5. RAID 5 with an online spare drive

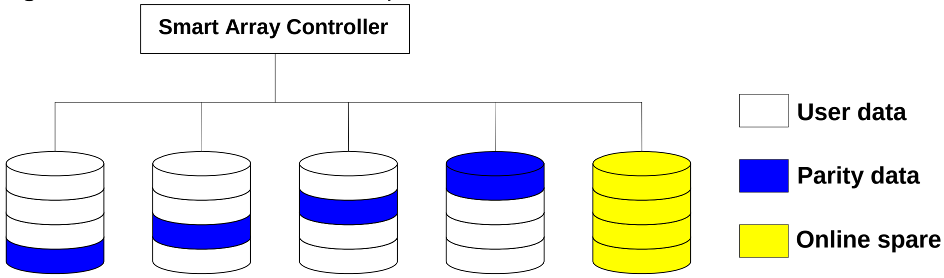

Dynamic sector repair

Under normal operating conditions, disk drive media can often develop defects over time, caused by variances in the drive mechanisms. To protect data from media defects, HP built a dynamic sector repair feature into Smart Array controllers.

During inactive periods, Smart Array controllers configured with a fault-tolerant logical drive perform a background surface analysis, continually scanning all drives for media defects. During busy periods, Smart Array controllers can also detect media defects when a bad sector is accessed. If a Smart Array controller finds a recoverable media defect, the controller automatically remaps the bad sector to a reserve area on the hard drive. If the controller finds an unrecoverable media defect and a fault-tolerant logical drive is configured, the controller automatically regenerates the data and writes it to the remapped reserved area on the hard drive.

ECC protection

HP Smart Array cache modules use ECC technology to protect cache data. DRAM is accessed by transferring read-or-write data 32 or 64 bits at time, depending on the type of cache module. The ECC scheme generates 8 bits of check data for every 32 or 64 bits of regular data. This check data can be used not only to detect data errors, but also to correct them. Data errors could originate inside the DRAM chip or across the memory bus.

Battery-backed write cache

HP Smart Array controllers ensure that data is protected in a location isolated from server failures before acknowledging that the data transfer has completed. The write-back cache allows the controller to acknowledge transfer completion before the data is physically stored in the hard drive. To improve disk write performance, data is temporarily stored in battery-backed write cache (BBWC), which uses DRAM and is substantially quicker when compared to the disk drive.

Battery power is required for RAID controllers to perform operations such as write-back cache, array expansion, logical drive extension, stripe size migration, and RAID migration.

Recovering data from battery-backed cache

If an unexpected server shutdown occurs while data is held in battery-backed cache, Smart Array controllers automatically signal the memory chips to enter a self refresh state and the controller initiates battery power, or system auxiliary power if present. An amber LED, available either on the cache module or battery pack, begins flashing to indicate that data is trapped in the cache. Smart Array controllers automatically write this data to disk when power is restored to the system. If power is not restored within the specified backup duration, the batteries may become drained of power. If that happens, posted-write data in the cache will be lost. Once system power is restored, the batteries will automatically recharge if needed. Battery recharge takes between 30 minutes and 2 hours, depending on the remaining capacity level.

In the event of a server failure, the Smart Array controller and all of the drives can be moved to another server to allow writing the data in the write cache to the drives.

In the event of a controller failure, the cache module containing posted-write data can be moved to a new Smart Array controller. However, to preserve the cached data, the new Smart Array controller must be attached to the original drives for which the posted-write data is intended.

Administrators should be aware of a special concern when using an embedded RAID controller with battery-backed cache. If the server board fails, the replacement board must be the same model server board so that the controller type and drive bays are the same. The cache module, battery pack, and drives must be moved to the replacement system to extract the data from the battery-backed cache.

Selection criteria for battery-backed cache

HP Smart Array battery cells, battery enclosures, and contacts are custom designed to preserve the integrity of business-critical information beyond the minimum specified backup duration. HP Smart Array battery cells were selected to achieve the specified three-year backup life in typical server environments.

A dedicated battery microcontroller continuously monitors the HP Smart Array battery pack for signs of damage, including an open battery terminal, partial battery short, charge timeouts, and over discharge conditions. Battery status information is indicated with an LED, power-on self-test (POST) messages, event messages to the host, ACU information pages, ADU, and within HP SIM.

The battery microcontroller automatically disables the battery-backed cache features any time it detects battery damage or the charge level falls below the required limits to achieve the specified backup duration. The battery microcontroller automatically restores battery-backed cache features when the microcontroller detects a replacement battery or when battery recharging is complete. High-end HP RAID controller designs contain two batteries to protect against a single battery cell failure while data is held in cache.

For detailed technical information on all HP cache options and controller compatibility, go to www.hp.com/products/smartarray.

Types of batteries

HP Smart Array controllers use rechargeable Nickel Metal Hydride (NiMH) button cell batteries specifically designed for longer life at the temperatures found inside rack-mounted servers. Typical capacity reduction for the HP Smart Array battery pack is 5 to 10 percent over a 3-year period, depending on server temperature and number of discharge cycles.

NiMH cells are also environmentally friendly, since they do not contain harmful lead, mercury, or cadmium material. Additionally, NiMH chemistry does not suffer capacity memory effects that can lower battery capacity. For example, memory capacity of Nickel Cadmium (NiCD) batteries is reduced when the batteries are exposed to short discharge cycles. Lithium Ion (Li-Ion) batteries are typically smaller than NiMH batteries, but their capacity is permanently reduced at high temperatures and they are usually limited to 100 full discharge cycles.

The HP 650 mAh P-Series battery has the same form factor as previous versions and extends battery life up to 48 hours before recharging is necessary.

Battery replacement

HP Smart Array controllers include serviceable battery packs which allow tool-free battery pack replacement with no need to replace either the Smart Array controller or the detachable cache module.

Alternatives to battery replacement

If it is not possible or desirable to replace the batteries, administrators have three options to disable write-back cache and avoid losing critical data:

1. The ACU provides a method for adjusting the read-and-write cache ratio to 100 percent read cache.
2. The ACU provides a method to disable the array accelerator for each logical drive, which disables both read-ahead and write-back cache.
3. Administrators can replace an existing RAID controller with a newer Smart Array controller model.

NOTE:
If the write cache is turned off, some write performance degradation may occur.

For more information about Smart Array controller battery configuration and specifications download the "HP Smart Array Controllers for HP ProLiant Servers User Guide": http://h20000.www2.hp.com/bc/docs/support/SupportManual/c01608507/c01608507.pdf?jumpid=reg_R1002_USEN

Entry level RAID solutions

Software RAID

HP has developed a software RAID solution based on the Smart Array firmware. The B110i SATA Software RAID supports the ACU, ACU-CLI (command line tool), SNMP agents, and Web-Based Enterprise Management (WBEM) providers.

The B110i features an OS-specific driver from HP that uses the embedded ICH10R controller. The B110i supports RAID 0, 1, and 1+0 and a maximum of two logical drives. The B100i supports up to four 1.5G or 3G SATA drives. Administrators can migrate drives to a Smart Array controller in a seamless procedure that maintains the user data and RAID configuration.

The ProLiant DL320 G6 server supports the B110i. For a listing of the complete feature set and support information for the B110i SATA Software RAID, download the B110i user guide at http://h20000.www2.hp.com/bc/docs/support/SupportManual/c01706551/c01706551.pdf

Zero Memory RAID

Using Zero Memory RAID (ZMR), administrators can create a RAID 0-1 configuration without additional memory. ZMR uses memory embedded in the controller, approximately 1K in size, and supports limited configurations. ZMR supports up to eight drives in Zero Memory Mode, or seven drives and one tape drive. ZMR does not include caching; however, all systems can be upgraded to a BBWC memory module that can significantly increase performance.

ZMR is supported on present generation Smart Array controllers for internal, direct connections only. This includes the Smart Array P410, P410i, P212, and P712 controllers. Consequently, the P212 controller does not include ZMR on the external connector. Modular Smart Array (MSA) products are not supported in ZMR mode.

NOTE:
The P212 controller can only be upgraded to 256 MB BBWC, so it does not support 512 MB BBWC.

NOTE:
Smart Array Advanced Pack is not available on Zero Memory configurations.

Storage support and pathway redundancy

HP Smart Array controllers support Native Command Queuing and Dual Domain providing increased performance and redundancy on the storage network. Smart Array controllers continue to support tape back up devices.

Native Command Queuing

Native Command Queuing (NCQ) increases SATA HDD performance by internally prioritizing read/write command execution. This reduces unnecessary drive head movement which results in increased performance, especially in server or storage-type applications where multiple simultaneous read/write requests are common. Without NCQ, the drive can process and complete only one command at a time. NCQ must be supported and turned on in both the controller and the drive.

Dual Domain support

Dual Domain SAS creates redundant pathways from servers to storage devices. These redundant paths reduce or eliminate single points of failure within the storage network and increase data availability. Dual domain SAS implementations can tolerate host bus adapter (HBA) failure, external cable failure, expander failure, or failure in a spanned disk (JBOD).

Dual Domain support is available for the HP Smart Array P800 attached to an MSA60/70 with the HP StorageWorks Dual Domain I/O Module Option (AG779A). Dual Domain support requires HP Smart Array firmware v5.10 or higher and dual-port SAS drives.

Dual Domain supports Multi-Initiator JBOD Clusters. I/O Module Option with the HP SC44Ge Host Bus Adapter supports HP-UX and OpenVMS on select HP Integrity servers only.

Dual Domain is also supported by HP Serviceguard software on HP-UX with HP-UX 11iv2 and 11iv3 only.

Tape device support

Smart Array controllers support tape back-up devices. The One Button Disaster Recovery (OBDR) ProLiant server/controller compatibility matrix for currently shipping HP products is available at www.hp.com/go/obdr.

Smart Array Advanced Pack

HP Smart Array Advanced Pack (SAAP) firmware provides advanced functionality within Smart Array controllers. This firmware further enhances performance, reliability, and availability of data. SAAP is hosted on the Smart Array controller hardware firmware stack. It can be enabled beginning with the present generation of Smart Array controllers.

SAAP requires a license key for activation. After activation, administrators can use several standard capabilities:

- RAID 6 with Advanced Data Guarding (ADG) protects against failure of any two drives. It requires a minimum of four drives, but only two will be available for data. ADG can tolerate multiple simultaneous drive failures without downtime or data loss and is ideal for applications requiring large logical volumes, because it can safely protect a single volume of up to 56 disk drives. RAID ADG also offers lower implementation costs and greater usable capacity per U than RAID 1.

- RAID 60 allows administrators to split the RAID storage across multiple external boxes. It requires a minimum of eight drives, but only four will be available for data.

NOTE:
HP BladeSystem servers do not support RAID 6 and 60.

- Advanced Capacity Expansion (ACE) automates higher capacity migration using capacity transformation to remove logical drives by shrinking and then expanding them online. Standard drive migration and expansion remain unchanged.
- Mirror Splitting and Recombining in Offline Mode breaks a RAID 1 configuration into two RAID 0 configurations. This is similar to a scaled down rollback functionality that requires two disk drives.
- Drive Erase completely erases physical disks or logical volumes. This capability is useful when decommissioning, redeploying, or returning hard drives.
- Video On Demand Performance Optimization optimizes performance of video on demand and improves latency during video streaming.

More information about SAAP is available at www.hp.com/go/SAAP.

NOTE:
At a minimum, a 256 MB cache and battery kit is required to enable the SAAP license key. SAAP is not available on Zero Memory Configurations.

Storage management

To improve data storage management, HP Smart Array controllers include built-in intelligence that makes it easier for administrators to configure, modify, expand, manage, and monitor storage.

HP provides five utilities for managing an array on a Smart Array controller:

- ACU (Array Configuration Utility)—ACU is the main tool for configuring arrays on HP Smart Array controllers. It exists in three interface formats; the ACU GUI, the ACU CLI, and ACU Scripting. All three formats have separate executables:
  - Starting with version 8.28.13.0, ACU Scripting is now a standalone application that is distributed with the ACU CLI application. In ACU versions prior to 8.28.13.0, the scripting executable was provided with the ACU GUI component.
  - Users familiar with the previous versions of ACU Scripting must now install the ACU CLI application to obtain the scripting executable. The new ACU scripting executable (hpacuscripting) replaces the former executable (cpqacuxe) in all scripts.

- ORCA (Option ROM Configuration for Arrays)—A simple ROM-based configuration utility
- CPQONLIN—A menu-based configuration utility specifically for servers using Novell NetWare
- HP System Insight Manager— Extensible, secure management suite for server and storage environments
- ADU (Array Diagnostic Utility)—A diagnostic and reporting utility for Smart Array controllers. With the introduction of version 8.28.13.0 or later of ACU or HPACUCLI, ADU is no longer provided as a separate utility but now has its functionality integrated with ACU and HPACUCLI.

For more information about storage management and using these utilities, see the "Configuring Arrays on HP Smart Array Controllers Reference Guide," available at:

http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00729544/c00729544.pdf

Table 6 summarizes the capabilities of ACU, ORCA, and CPQONLIN utilities.

Table 6. Summary of HP configuration utility features

Utility features	ACU	ORCA	CPQONLIN
Uses a graphical interface	Yes	No	No
Available in languages other than English	Yes	No	No
Executable at any time	Yes	No	Yes
Available on disc	Yes	No	No
Uses a wizard to suggest the optimum configuration for an unconfigured controller	Yes	No	Yes
Describes configuration errors	Yes	No	No
Creating and deleting arrays and logical drives	Yes	Yes	Yes
Assigning RAID level	Yes	Yes	Yes
Sharing spare drives among several arrays	Yes	No	Yes
Assigning multiple spare drives per array	Yes	No	Yes
Setting stripe size	Yes	No	Yes
Migrating RAID level or stripe size	Yes	No	Yes
Configuring controller settings	Yes	No	Yes
Expanding an array	Yes	No	Yes
Creating multiple logical drives per array	Yes	No	Yes
Setting of boot controller	No	Yes	No

Array Configuration Utility

The ACU[4] is a browser-based graphical application that helps configure Smart Array controllers. The ACU is also supported on the HP MSA family of entry-level SAN products. This provides administrators with a seamless set of tools to use with both HP Smart Array direct-attached storage (DAS) and HP SAN-attached storage.

The ACU runs online on Microsoft® Windows® Server and Linux® operating systems. Because the ACU is available on a bootable disc, an administrator using other operating systems[5] can run the utility offline by booting the system from the SmartStart or Server Storage Support Software disc.

[4] For more information about the ACU, refer to http://h18004.www1.hp.com/products/servers/proliantstorage/software-management/acumatrix/index.html.

[5] Refer to product documentation to verify operating system support.

The ACU also contains a command line interface (ACU-CLI), offering a quicker way to deploy multiple servers by automating creation of arrays and logical drives. Configuration information for the ACU is available at
http://bizsupport1.austin.hp.com/bc/docs/support/SupportManual/c00729544/c00729544.pdf

Option ROM Configuration for Arrays

Option ROM Configuration for Arrays (ORCA) is an alternative method of viewing, creating, and deleting multiple arrays and logical drives during system power up. ORCA is designed for administrators who have minimal configuration requirements. The ACU is recommended for more advanced array configurations.

CPQONLIN

CPQONLIN is a configuration utility that runs online on Novell NetWare. It functions like the ACU in this environment. Refer to the HP Smart Array controller documentation for more information about using CPQONLIN.

HP Systems Insight Manager

HP Systems Insight Manager is a client/server tool for integrated server environment management. Based on SNMP, it is capable of monitoring more than 1,200 system-wide parameters for performance and other operational characteristics of Smart Array controller storage. The program displays configuration information, operating system device driver version numbers, controller firmware version numbers, Pre-Failure Warranty information, and operating statistics.

Performance monitoring

HP SIM gives administrators a window to look at low-level performance characteristics of Smart Array controllers in the environment. It monitors three basic Smart Array controller performance parameters for proactive storage subsystem management:

- I/O commands per second
- Average command latency
- Local processor utilization

Analyzing these key parameters can assist administrators in fine tuning configurations for performance. HP SIM can graphically chart performance over time for each of these parameters.

Fault prediction

A background task monitors several key drive parameters and notifies HP SIM if a drive fails to meet certain factory-preset criteria. HP SIM alerts the administrator to the potential problem.

Array Diagnostics Utility

The HP Array Diagnostics Utility (ADU) is an in-depth diagnostic and reporting utility for all Smart Array controllers. The ADU quickly identifies problems such as incorrect versions of firmware, improperly installed drives, inappropriate error rates, and failed batteries on the array accelerator board.

The ADU displays a detailed analysis of the system configuration. If the cause is not apparent, the ADU can generate a full report for administrators to fax or e-mail to HP customer service for phone support.

Summary

HP Smart Array controllers are powerful I/O solutions for today's most demanding network data storage requirements. Smart Array controller technology combines an advanced feature set and full-spectrum, hardware-based fault management with exceptional performance characteristics.

Smart Array controllers provide solutions for all four primary data storage requirements: capacity growth, high performance, data availability, and manageability. Smart Array controllers represent patented HP technology and continue the HP commitment to providing the most complete and advanced storage solutions for server customers.

HP is the only server provider with a seamless storage solution set that spans the range from embedded Smart Array controllers in servers to plug-in PCI Smart Array controllers to SAN-attached MSA storage. The tools used for managing and configuring storage are the same in all of those environments. Data sets are compatible across all of those environments, The HP universal drive strategy allows customers to easily migrate data from DAS to SAN.

Appendix A: Capacity growth technologies

The Smart Array controller family includes a standard toolset that administrators can use to configure array controllers, expand an existing array configuration by adding disk drives, or reconfigure an array by extending logical drive sizes. Before this innovation, expanding the storage capacity attached to an array controller required a time-consuming backup-reconfigure-restore cycle.

Array expansion

Array expansion is the process of adding physical drives to an array that has already been configured. The logical drives that exist in the array before the expansion takes place are unchanged; only the amount of free space in the array changes.

For example, suppose an existing array consists of four physical drives and the administrator wants to expand the array to six physical drives. This is like having four glasses full of water and dividing that water from the original four glasses among six glasses (Figure A1). The amount of water (the size of the logical drive) has not changed—it has merely been redistributed, or expanded, into a larger number of containers (drives).

Figure A1. Array expansion redistributes an array into a larger number of physical drives. The size of the logical drive does not change.

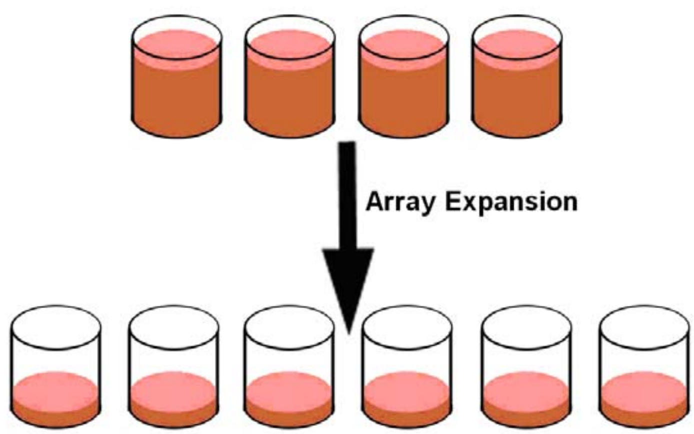

Thus, if an existing array is nearly filled with logical drives, an administrator can add new physical drives and initiate an array expansion using the ACU. The ACU automatically checks the drive hardware configuration to ensure the array expansion. Then, the existing logical drive is distributed across all physical drives in the expanded array without affecting any existing data. If the array being expanded contains more than one logical drive, data is redistributed one logical drive at a time.

The expansion process is entirely independent of the operating system. For example, if a 10-gigabyte (GB) logical drive is expanded from four drives into six drives, the operating system is unaware of the change.

The amount of time it takes to complete an online array expansion depends on subsystem I/O activity, the size of the logical drives on the array, fault-tolerance level, the number of disk drives in the new array, disk drive performance, and the priority level of the array expansion. Administrators can monitor array expansion completion progress within the ACU.

The ACU allows administrators to set the priority of array expansion operations. A high-priority setting prioritizes the array expansion operation over I/O requests. A low-priority setting is the ACU default setting and gives I/O commands precedence over array expansion. During idle time, when no I/O commands are active, the array expansion operation runs at full speed regardless of the priority setting.

Logical drive creation

Once the array capacity is expanded, the added capacity can be used to create a new logical drive (Figure A2) or to extend the size of an existing logical drive (described in the "Logical drive extension" section).

Figure A2. After array expansion, the administrator can use the free space by creating a new logical drive.

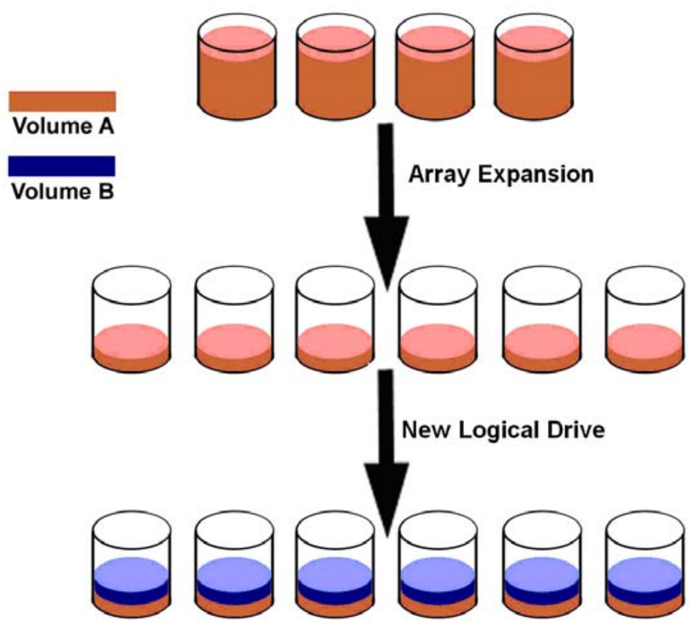

For example, if every department in a company has its own logical drive and a new department is created, the administrator might need to create an entirely new logical drive for that department.

However, many administrators wish to extend the size of their logical drive after array expansion. Logical drive extension lets administrators increase logical drive size if the existing drive is running out of storage space.

Logical drive extension

Logical drive extension increases the storage space of a logical drive (Figure A3). During this process, an administrator adds new storage space to an existing logical drive on the same array. An administrator may have gained this new storage space either by array expansion or by deleting another logical drive on the same array.

Figure A3. Logical drive extension grows the size of a logical drive.

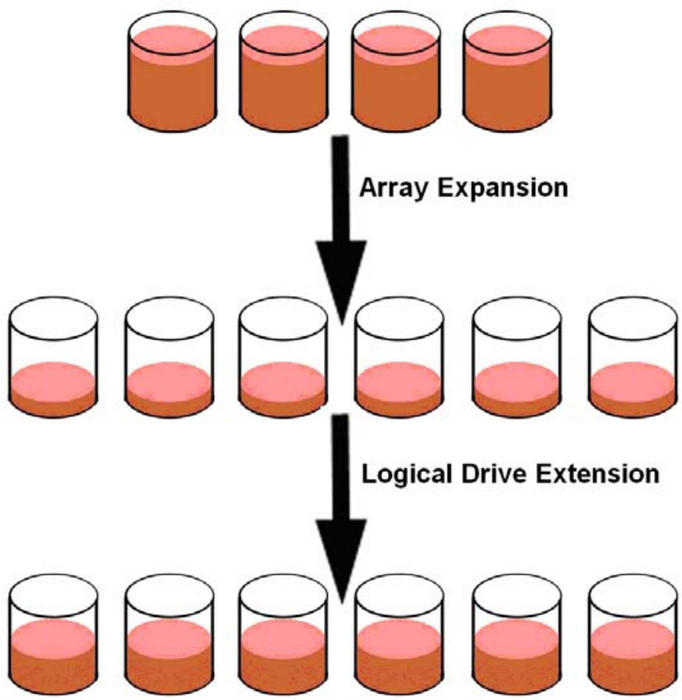

The operating system must be aware of changes to the logical drive size. Refer to operating system documentation for additional information.

For more information

For additional information, refer to the resources listed below.

Resource description	Web address
Smart Array controllers	www.hp.com/products/smartarray
HP hard disk drive products	www.hp.com/products/harddiskdrives
HP Modular Smart Array external storage systems	www.hp.com/go/msa
HP ProLiant servers	www.hp.com/go/proliant
HP SAS Technology	www.hp.com/go/serial
Configuring Arrays on HP Smart Array Controllers Reference Guide	http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00729544/c00729544.pdf
RAID 6 with HP Advanced Data Guarding technology	http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00386950/c00386950.pdf
Smart Array RAID Controllers Support Matrix	www.docs.hp.com/en/SM_SA3/SA_Support_Matrix.htm
Smart Array Advanced Pack	www.hp.com/go/SAAP
Serial Attached SCSI storage technology brief	http://h20000.www2.hp.com/bc/docs/support/SupportManual/c01613420/c01613420.pdf.
Redundancy in enterprise storage networks using dual-domain SAS configurations	http://h20000.www2.hp.com/bc/docs/support/SupportManual/c01451157/c01451157.pdf
HP Smart Array controllers for HP ProLiant Servers User Guide	http://h20000.www2.hp.com/bc/docs/support/SupportManual/c01608507/c01608507.pdf?jumpid=reg_R1002_USEN
HP One-Button Disaster Recovery (OBDR) Solution for ProLiant	http://h18006.www1.hp.com/products/storageworks/drs/index.html

Call to action

Send comments about this paper to TechCom@HP.com.

TC090807TB, August 2009

invent