hp storage

january 2002

technical white paper

# hp virtual array technology

## virtual storage technology extends the capabilities of fault-tolerant storage subsystems

### executive summary

Storage has become a key asset to today's IT infrastructure. Effective storage solutions must provide data integrity, data availability, performance, and efficient management. This technical white paper outlines the key differentiators in a new line of fault-tolerant storage devices from Hewlett-Packard, and their value for increasing the effectiveness of the corporate IT center. HP recognizes the need to provide IT administrators with storage solutions that increase operational availability while reducing the IT management costs. These features are embodied in the VA7000 series storage arrays. As part of HP's Federated Storage Area Management strategy, the VA family of products provides an industry-unique solution to the complex environment of the IT center.

### introduction

Modern data centers continue to grow in capacity and complexity at a once unimaginable rate. Today's IT administrators are often presented with the daunting task of managing this growth and complexity with decreasing human resources.

Storage is one of the main contributors to this growth and complexity. Once referred to as a peripheral, storage (and its associated data) has been elevated to the absolute center of the IT infrastructure. In many cases, storage and storage management is now the single largest budget item in medium to large data centers.

The basic requirements of storage have not changed significantly in the recent past. Data protection and integrity have always been a key need for any business. However, as the data center has become a 24x7x365 resource for a business, the need for continuous access and protection of data have become critical, and is now the number one requirement of storage systems.

Continuous access to data has two important aspects: high availability and performance. Availability refers to the state of the storage device. This is a two state condition, where the storage device is either functional in a way that allows systems to access the data, or it is not. The means to assure continuous access to data is often through various fault tolerant features in the storage device and / or the Storage Area Network (SAN) implementation.

Performance also has an implication to the accessibility of the data, but it is typically not viewed as a binary (yes or no) condition. However, insufficient storage system performance acts as a restriction of access to the data. Unacceptable performance in a storage system can have a similar effect on a business as a failure that has rendered the storage system unavailable.

Therefore, both fault tolerance as well as a high-performance architecture are key requirements to ensure continuous access of data for online storage devices.

While the need for data protection and continuous access remain the paramount requirement for storage systems, a new requirement has emerged. Fueled by the decreasing cost of basic storage and the competitive needs of business, new applications for storage have caused the rate of change in the data center to grow exponentially. The costs to administrate these changes have become a significant budget item for the IT manager. Compounding the requirement to plan and execute these changes is the shortage of trained and capable storage administrators. It has been predicted that if there are no improvements to the efficiencies of storage administrators, at the current rate of storage growth, the demand for storage administrators in 2010 will equal the population of the state of California.

Today's IT center requires cost effective storage devices that are highly available, have high data integrity, are high performing and require less management. Technologies that reduce time-to-implementation of new applications and reduce the ongoing management costs can be the key to remaining competitive in today business and economic environment.

Hewlett-Packard has a line of storage systems that offer these advantages.

## hp virtual array technology

Hewlett-Packard has recently announced a new line of advanced, highly available storage systems. The HP Virtual Array series offers modular storage solutions for the mid-range to high-end data center applications. Based on the third generation of HP's patented virtual storage technology, the Virtual Array series extends the capabilities of mid-range storage systems.

The HP Virtual Array series offer storage solutions with the following highlights:

- Advanced data integrity and high availability features – greater protection against data loss and down-time
- Modular architecture – low cost, pay as you grow design
- High performance – industry leading price/performance for a highly available storage system
- Powerful management capabilities – reduces both time-to-implementation and ongoing management costs

## data integrity and high availability features

At the heart of a HP Virtual Array is a state-of-the-art custom application specific integrated circuit (ASIC) developed by HP. This ASIC is key to providing high-performance, data protection and low costs within a single architecture. A team of HP engineers with experience in array, disk mechanism, and server design spent three years developing this critical ASIC. This patent-protected implementation is industry leading for mid-range modular storage. In Figure 1, this ASIC is labeled the Data Flow Processor (DFP).

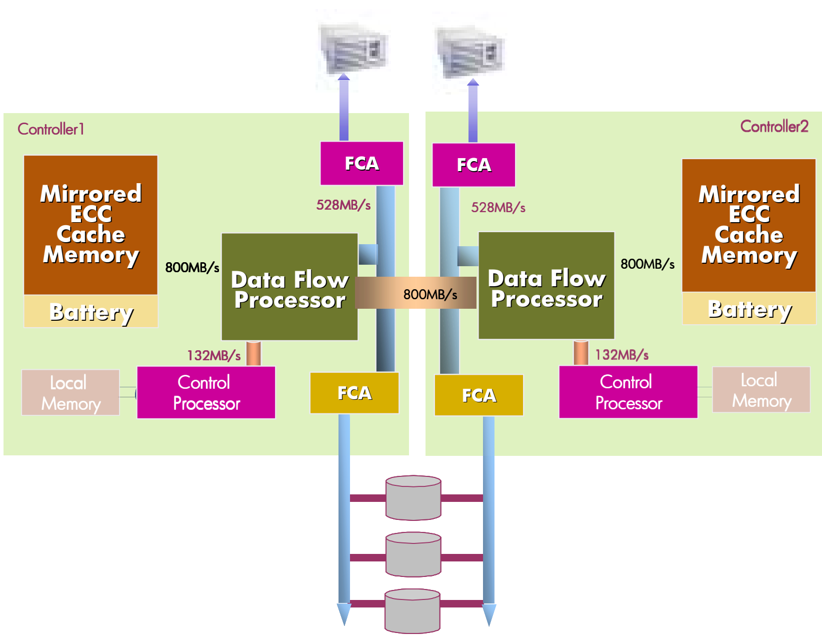


**figure 1 - architecture block diagram**

The hardware architecture of HP's Virtual Array series provides the following key features:

- Redundancy - dual controllers, power supplies, cooling, and mirrored write cache with integrated battery backup.
- Advanced Data protection - ECC and mirrored cache, end-to-end data protection, and Dual Parity RAID 5
- High Performance - 800MB/s tightly coupled array architecture, multiple, independent data paths and high-performance PowerPC based controller

There are a number of excellent white papers from Hewlett-Packard describing all the high availability and performance features provided by the Virtual Array. This paper will highlight just a few of these features. To obtain more detailed feature descriptions and other technical white papers, please refer to the "for more information" section at the end of this white paper.

### end-to-end data protection

End-to-end data protection is a hardware-based process that assures the server will read exactly what was written. For the HP Virtual Array, the end-to-end data protection is implemented directly into the proprietary Data Flow Processor (DFP) hardware. As data is written to the array, the DFP immediately encapsulates each sector with an additional error detection check word.  This check word is propagated along with the data as it flows through the controller and on to the disk drive. When the data is read by the server, the final step of the controller is to verify the correctness of this check word.  If an error is detected, the controller will preempt the server from using the invalid data.  The check words are formed from both data and address information thus can protect the system from errors generated in either the disks or in the array controller.

To illustrate this point, disks do a great job checking that the bits that are read are the same ones that were written to a sector. However, there is a small but finite probability that an unforeseen combination of circumstances will result in the disk returning incorrect sector. In computers, the personal computer on a user's desk probably has only parity checking on its memory; the PC server down the hall has basic error correction, and the very important server running a company-wide database has multiple means of error correction. The more critical a computer system is to the operation of a business, the more important that multiple schemes of integrity checking are in operation. Disk arrays are the same way: the disk in a personal computer only does its own checking, and a modest disk subsystem for non-critical data offers RAID protection against individual disk failures. However, business critical arrays require multiple schemes for data protection and data integrity.

End-to-end data protection is key to providing data integrity for business critical applications.  Similar capabilities are available on all high-end arrays and a few midrange arrays.  Arrays without this feature have a lower class of data integrity.   In the absence of this feature, the array controller could return incorrect data without any indication to the server, and thus the error would be undetected by the application.  Although the occurrence of this failure is rare, its impact can be very disruptive.  The outcome of this failure usually results in an unexplained system crash or database corruption.

## RAID5DP

Another advanced data protection feature for the HP Virtual Array is RAID 5DP.  RAID 5DP is similar to standard RAID 5, but with the addition of a dual parity protection scheme.  Each block of data in RAID 5DP has two blocks of associated redundant information, not just one.  Hewlett-Packard has published two additional white papers that describe how RAID 5DP works, and the statistical advantage to RAID 5DP.   Two pertinent facts emerge from these papers.  First, in a typical configuration RAID 5DP is over **100** times more powerful at protecting data than standard RAID 5 and, in this same configuration, provides 10% greater storage efficiency. Second, RAID 5DP allows the HP Virtual Array to be resilient to almost all simultaneous dual disk failures, in any configuration of the array.

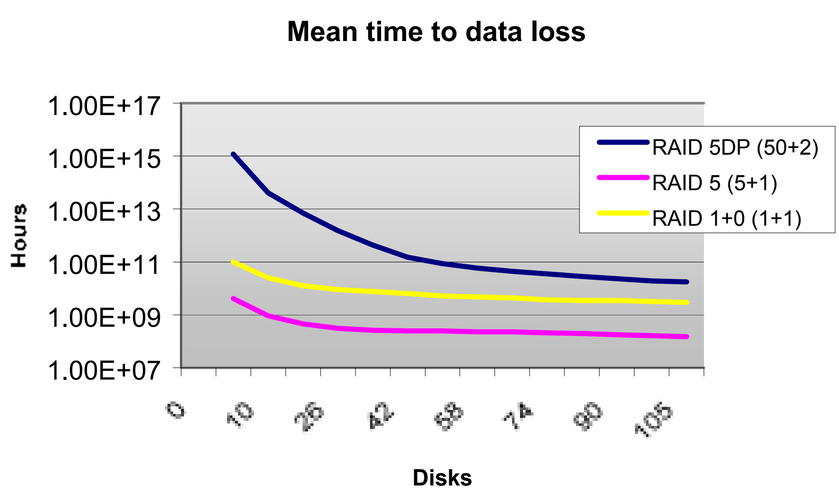


**figure 2 – mean time to data loss**

Let's examine this second point in greater detail.  Independent of the number of disks or the configuration of the array, at the instant of the first disk failure, the Virtual Array has only one or two other critical disks in the entire array[1].  That is, any of the other non-critical disks can fail simultaneously with the first disk, and the array will continue to function without data loss.  The failure of a critical disk would cause data loss. The following example uses a typical configuration of 45 disks. After the first disk failure there are either one or two disks out of the remaining 44, where if they were to fail it would cause data loss.  Failure of any of the other 42 disks will not affect the availability of data.

Although this could be comparable to other arrays configured in RAID 1+0, it exceeds the capability of any other RAID system configured in traditional RAID 5.  In traditional RAID 5, the stripe depth defines the remaining critical disks in the system after the initial disk failure.  For example, with a typical 5+1 disk RAID 5 configuration, there are five critical disks in the system after the initial disk failure.

[1] The number of disks in the redundancy group determines the number of critical disks.  An even number of disks will result in just one critical disk; an odd number will have two critical disks.  Knowing this rule, it is easy to create configurations that minimize the number of critical disks after a disk failure.

The logic encoded within the controller of the Virtual Array extends this capability even farther. Each controller 'knows' what data is exposed to a second simultaneous failure, and will rebuild this data first, before restoring the dual redundancy in the RAID 5DP data. This capability is unmatched in the industry.

These are two examples of high-availability features within the Virtual Array. End-to-end data protection is a feature that matches the capabilities of high-end arrays, and RAID 5DP is a feature that exceeds the capability of high-end arrays. There are many other valuable high-availability features of the Virtual Array series of products described in white papers available from Hewlett-Packard. Each feature adds to the data protection and high-availability characteristic of the HP Virtual Array.

## modular design allows flexibility of growth

Modular arrays allow storage administrators to have more flexibility in initial configurations, as well as the ability to easily expand the system after its initial installation. Modular arrays are a practical component of a Fibre Channel Storage Area Network (SAN). The small footprint of these types of arrays and the corresponding high-density design allow effective scaling using SAN technology. This system architecture results in a lower cost, pay-as-you-grow solution when compared with monolithic, cache-centric architectures.

The virtual technology used within the HP Virtual Array gives added value and flexibility in modular array solutions. The following are highlights of the growth capabilities of the HP Virtual Arrays:

- Quickly expand the usable capacity from the minimum supported disk configuration to the maximum configuration without any interruption in service.
- Use any configuration of supported disks as well as any count or capacity – without complicated management processes
- The flexibility to locate add-on storage shelves anywhere in the data center via an optical Fibre Channel interface
- Compatible with the full line of Fibre Channel infrastructure devices from HP – topologies for either direct connect or Storage Area Networks
- Broad list of non-HP operating systems support, including Solaris, AIX, Linux, Netware and more – providing both the opportunity for storage consolidation and an investment protection

These features combined offer greater flexibility and reduced management than any other disk array in the industry. Let's examine these features in greater detail. The following will compare a typical storage growth scenario and compare the Virtual Array to a traditional array.

### easily add capacity

One dramatic difference between HP VA series arrays and traditional storage arrays is initial purchase of capacity. With traditional array technology, disks must be purchased in groups according to the desired RAID configuration. The initial configuration of a traditional disk array represents a static trade-off of performance, storage efficiency and availability. With the HP VA series arrays, there is no need to purchase disks in groups for either RAID 1+0 or RAID 5DP configurations. Array-based storage virtualization eliminates this requirement; disks are treated as a pool of available storage. Adding more disks simply increases the capacity of this pool. In addition, HP offers an intuitive capacity planning tool that shows the minimum number of disks needed to meet the capacity requirements. The combination of storage virtualization plus HP's capacity planning tool makes it easy to purchase only the needed capacity with an HP Virtual Array.

Now let's compare the growth phase. The value of modular storage is that it allows a pay-as-you-grow model. However, traditional arrays limit either the configuration flexibility or the system performance after the growth. As with the initial configuration, traditional arrays require that disks be purchased in groups reflecting the RAID level and stripe size. The pragmatics of the upgrade in a traditional array necessitate that these disks are installed as an island of storage, isolated and independent from the other disks in the array. This isolation results in both poor storage efficiency and poor usage of the performance value of these additional disks.

In comparison, with HP's Virtual Array technology, capacity growth can be accomplished with any count or capacity disks. Therefore, it is very easy to purchase only the needed disks to meet the capacity requirements. The new capacity can be installed and configured in just a few minutes – not the hours required for traditional arrays. Once installed, the controller will reorganize the data in the background to efficiently include the new disks, thus assuring the maximum storage efficiency and performance of the new disks.

These differences not only minimize the product budget of the storage solution, they also minimize the management time required to implement the solution. The benefits are lower costs and a faster time to solution.

## array performance

Array performance can be as critical to the business value of online storage as fault tolerance features. A poorly performing storage system is a potential source of lost revenue or dissatisfied IT customers. The key to a successful solution implementation is the correct optimization of the storage device. Unfortunately, fault-tolerant disk arrays are very complex devices, and successful performance tuning can be very difficult to accomplish. Compounding the difficulty of this goal are any ongoing changes and demands on the storage device. As the applications change or additional capacity is added to the array, re-tuning is often omitted due to management time constraints.

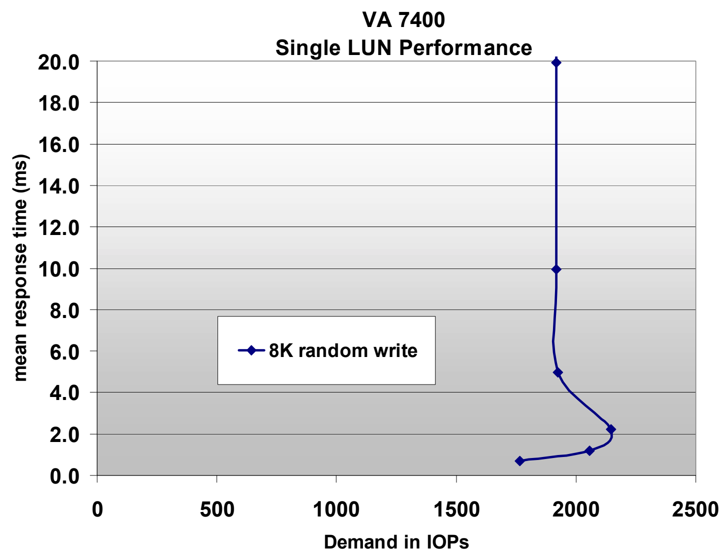


**figure 3 – random write performance[2]**

The HP VA series array has a number of features and characteristics that decrease the complexity of performance tuning and increase probability of a continuously well-tuned system. These are intrinsic components to the design of the array, and they are based on well-known rules of storage, server and database experts.

Three key features stand out as enhancements to performance:

- Write cache response time
- Back-end disk performance
- Large stripe size

Let's review the value to the customer of each of these.

### write cache response times

The VA series array has very fast write cache response times. In general, fast write cache responses result in less wait time for the application. This is fundamental for applications to run at full performance potential. The three features that allow the VA series array to accomplish this are: fast internal processors, a performance tuned firmware path and proprietary high performance write mirroring hardware. While many arrays can claim fast processors and firmware tuning, the architecture of the Data Flow Processor gives the VA series array a distinct advantage in mirroring the write cache. The Data Flow Processor uses a message based, high speed, point-to-point network design. This results in both reduced controller overhead and assured data consistency.

With very few instructions from the internal processor, each mirrored write can be accomplished very quickly and with assurance that the data will not be corrupted either in content or in logical ordering. Figure 3[2] depicts the performance result of this capability. This graph shows the single LUN random write response time as a function of demand. Note that these are random writes to the disks, not just a contrived 100% write cache workload. The no-queue response time measurement is at the lowest point of the curve, at approximately 700 microseconds. This represents the service time for a write cache hit, which includes the time to copy the data to the mirrored cache.

In comparison, disk arrays that use a back-end disk channel for mirroring will have difficulty in achieving this level of performance. With these arrays, a back-end channel read or write requires far greater controller overhead to initiate and the transfer rate is far less than with the VA series array's tightly coupled design. Unfortunately for customers attempting to evaluate alternative array performance, vendors typically publish array specifications that indicate the cache read performance, in operations per second, but this is practically meaningless to real life workload performance.

An application well tuned for performance is far more dependent on write cache response time than read cache IOPs (operations per second). First, a tuned application will use system buffer cache (System Global Area for an Oracle data base) for read cache, not disk based caches. A read from a system memory based cache is far faster, 100's of thousands of IOPs vs. 10's of thousands of IOPs for disk-based caches. Second, although the server buffer cache can also be used as

[2] This test was conducted with a 4-way HP N-Class HPUX system. The storage system under test was a HP VA7400 with 45 15K RPM disks, RAID 1+0, 1GB cache.

a write cache, an application that is well tuned for performance must still direct some IOs to durable media, and therefore cannot use the system buffer cache. Array write cache response time is critical for these IOs, because the application is typically waiting for the IO to complete before proceeding. Faster write cache response time means the application waits less.

Therefore, tightly coupled controller architecture provides the best real-world performance.

## back-end disk performance

The VA series array has very efficient back-end disk performance; it is easy to demonstrate over 200 IOs/second per disk in a benchmark scenario. This performance level is at the practical limit of the disk performance. This is important because the performance value of an array is derived fundamentally from both the array write cache and disk performance. Open systems (Unix and Windows) are very unfriendly to array caches; this is because these systems have an efficient disk cache integrated into both the database and/or IO path. The resulting workload that is presented to the array system has very little locality, and thus is unlikely to benefit from an array-based cache[3]. This does not imply that array based caches are not important. A moderate amount of array cache can be very valuable to increasing system performance. However, the array must have the back-end disk performance to support the front-end cache.

Therefore, when selecting disks arrays for performance, one should compare cache write response times and back-end disk performance per disk. These are good indications of OLTP, file system or NFS performance. Thus, the efficient backend disk performance of the VA series makes it an excellent choice for these applications.

## large stripe size

This single capability of the VA series array is both the most valuable and controversial feature to enhance real-world performance. Large stripes effectively distribute the performance demand to all the disks in the array, thus increasing the overall system performance. This has been demonstrated numerous times by HP customers and HP's internal systems benchmark lab. The controversy comes from old rules of database configuration, where data base administrators were guided to isolate different components of the database to separate spindles in the array, thus avoiding spindle contention. This rule is no longer valid for two reasons. First, it is very difficult to achieve a perfectly balanced system. Any imbalance in the distribution of the workload will leave performance capabilities unused. Second, efficient write caches have changed the performance behavior of disk subsystems.

Let's look at a bank teller analogy to better understand this statement. Would you rather go to a bank that has a single queue for the tellers, or a bank with separate queues for each teller? It's obvious that a single queue provides the best service time for everyone needing the bank's services. However, the old school of thought from the spindle contention group is that some of the customers are more important than others, and they should have their own group of tellers so they could get extra prompt service. This is fine, except for two issues. First, if the bank sized the number of special tellers incorrectly, an important customer could be waiting in the special teller queue even if there are standard tellers available – not good for service times. Second, write caches are like a proxy for customers to stand in the teller queues for them. Just think if you arrived at the bank and gave your transaction to someone who would stand in the queue for you and execute the transaction with the teller, enabling you to immediately leave the bank. The only thing you would care about is that there is an available proxy for you when you arrived at the bank. After that, the proxies just need to be processed by tellers in the most efficient order, and that's a single queue to the whole population of tellers.

This is a simplified description of why striping across a large set of disks is the best configuration for performance in modern disk arrays. Not only does it make intuitive sense, it really works. Oracle now recommends this approach for data base configuration. Within HP, the performance benchmark teams for TPC-C, SAP, and NFS are all using the VA series arrays. Why? Because the VA series has excellent cache write response times, an efficient high-performance back-end, large stripe sets, *and* the array is inexpensive. It's an unbeatable combination.

Figure 3 is also a good demonstration of this effect. This graph depicts the saturation demand curve for a singe LUN. In this case, the LUN is striped across 20 disks. For this single LUN, the performance curve shows 2000 8K random write IOs per second with 1 millisecond response time to the disks. Not only is this great 'real world' performance, it is easy and fast to configure. In addition, it's easy to maintain, even after system capacity growth.

---

[3] An array-based cache must be many multiple times the size of the system buffer to yield an appreciable hit rate. It is far more cost effective to increase the capacity of the system buffer cache than to invest in a very large array-based cache.

### price-performance

Deciding which storage platform is right for an application is difficult. There are often many choices available, each with a unique combination of features. For the analytics among us, the process usually becomes a specifications comparison – large spreadsheets with rows of features and columns of products. Unfortunately, the most important performance criteria are often left out.

Price-performance must be considered along with raw performance. While the absolute performance may be an indication of the capabilities of the array architecture, the price-performance (and price-capacity, $/MB) is the business value of the storage device's performance.

A competitive analysis should include $/GB, $/IOPs and $/MB/s. The inclusion of these characteristics could lead to un-obvious conclusions. Let's examine the VA series itself. The 'entry level' VA7100 offers less absolute performance than the mid-range VA7400. Many customers would select the VA7400 because it offers the highest performance. However, if the analysis includes the specific price for the applications, capacity requirements, the price-performance may give the VA7100 the edge.

This is exactly what a well-known Internet-based retailer discovered. They chose the VA7100 for an unobvious application – data warehousing. Their best solution was 32 VA7100s directly connected to an HP SuperDome processor! Their business value was price-performance for the solution, not the absolute performance specification for a single array.

## management software

Device management is key to receiving continued value for the array, and for addressing the ongoing problem of continually growing data centers. The VA series of arrays is a key component of the Federated Storage Area Management (FSAM) strategy for HP. As well as providing stand-alone device management, the VA series arrays integrate into the OpenView family of software products to provide enterprise wide storage management. For the purposes of this paper, the discussion will center on the unique device management characteristics for the VA series arrays, they are:

- Topology flexibility
- Initial configuration
- Performance tuning

Let's examine each of these to see how they contribute to the business value of the VA series arrays.

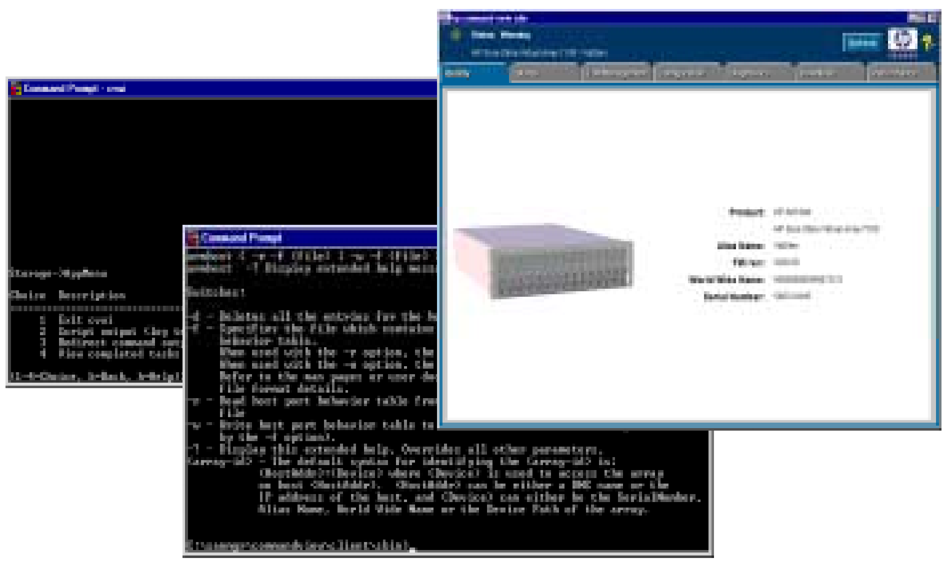

**figure 4– command view sdm options**

### management topology flexibility

Command View SDM has a varied set of operating topologies and interfaces, that, when combined, offer industry leading flexibility and ease of use. First, there is the choice of three interfaces to interact with the array: a graphical interface (GUI), Command Line, and Command Menu. The GUI is an easy to use, point-and-click interface that allows configuration, fault monitoring and performance analysis. The Command Line is a full featured, scriptable, text based management tool. The Command Menu is an easy to use menu driven interface that generates Command Line commands. Each of these modes allows management in the style that matches the need of the administrator and the task. The GUI works great for the inexperienced storage administrator or casual, one-off management tasks. The Command Line is great for scripting the configuration of an array, thus allowing self-documentation of the array configuration. Most Unix storage administrators prefer this mode of interaction. The Command Menu is an easy to use, low bandwidth interface, thus making this the interface of choice for either learning the Command Line commands or for interacting with the array over a low bandwidth interface, such as a modem.

The real topology flexibility comes from the ability to run any of these interfaces, either on the server directly attached to the array, on a dedicated management workstation, or from any computer in the network. Added to this is the flexibility to easily use the GUI interface directly from any browser in the network without installing any special software on that computer system.

Command View SDM takes this flexibility one step farther. These three interface environments are identical in all of the supported OS environments. The GUI, the Command Line and the Command Menu are the exact same commands on Unix and Windows. What does this mean for array administration? It's easier to learn, as the commands are this same for all the systems in the data center. It's flexible; the Unix administrator can manage the VA series array attached to a Windows server. It's server and location independent; a script written to configure an array can be run from any computer in the network and configures the same array. This is an unmatched feature for storage management. In totality, these features give administrators of new set of options to increase the flexibility and lower the cost of device management.

### initial configuration

Out of the box, HP VA Series arrays are simple and easy to setup. The vast majority of the setup time is the physical installation and cabling of the array. Once powered on, any scripted configuration can be completed in less than five minutes. This includes installation of Command View SDM software, configuration of LUNs, Secure Manager VA, and Business Copy VA. This saves time and money during the initial installation as well as during disaster recovery, when down time can be costing tens of thousands of dollars an hour.

This is a dramatic improvement in comparison with traditional disk arrays. For example, the initial configuration of a large HP FC60 system can take 7 hours or more to configure and format the array. In comparison, the VA7100 and VA7400 are almost immediately available for data access after installation.

### performance tuning

Disks arrays are very complex devices. System, application and array performance tuning can be very difficult to optimize. The traditional array philosophy is to bring forth every operational variable of the array to the storage administrator, which assures the flexibility to configure the array for any possible performance demand or workload. Unfortunately, rarely does a storage administrator have both the intimate knowledge of an application's specific workload and the time to monitor the performance sufficiently to take advantage of this myriad of combinations of controls presented by traditional arrays. Compounding this dilemma is the ever-changing demand on the array. New applications, or even new versions of the same application are added, changing the demand upon the storage. Tracking the performance impact of these is often an unaffordable management task. Traditional arrays are flexible, yes – but pragmatic, no.

The VA Series array offers two practical choices: one configuration for the 'performance-is-paramount' group, and a second self-managed configuration option for the 'I want the best performance I can get, but I don't have an unlimited budget' group. There are a few additional configuration options that guide the storage administrator of the VA series array in making decision about performance versus data integrity, such as accepting the risks associated with a write cache. In other words, HP's VA series arrays are both practical and easy to use.

The basic configuration for the 'performance-is-paramount' group is RAID 1+0. This RAID level offers the best possible performance for almost all applications, but has the lowest storage efficiency, thus is the most expensive in terms of usable $/MB. For the performance-value group, the VA series offers an industry unique self-managed storage option, AutoRAID.

## AutoRAID – a component of hp's virtual storage technology


For those applications that are sensitive to performance and cost of storage, the VA Series provides a self-managed storage optimization mode. Programmed with the same rules that experienced RAID storage administrators use to optimize price-performance, HP's AutoRAID technology component in its Virtual Storage Technology monitors and tunes the data placement for optimum performance.

AutoRAID technology monitors the use of the data stored in the array and determines the best RAID level for that data. Sequentially accessed data is stored in RAID 5DP, which provides the best performance and storage efficiency. Randomly written data is stored in RAID 1+0, thus providing the best performance for highly modified data. The management of the randomly written data is analogous to a hierarchal storage manager: active data is maintained in RAID 1+0; as it becomes inactive, it is automatically migrated to RAID 5DP to provide the best storage efficiency.

Like a good storage administrator, HP AutoRAID technology works 24 hours a day studying the workload, but only makes changes during low usage periods to minimize the impact on the application performance. While AutoRAID is working, it is keeping a plethora of statistics of its actions that can be easily reviewed by the human storage administrator. AutoRAID works weekends and nights so you don't have to!

Management of AutoRAID is also easy. Workloads that require more RAID 1+0 capacity can be accommodated by simply adding more disks to the system. Other options allow the storage administrator to dictate when data movement will take place.

## summary

HP's Virtual Array Storage technology provides a powerful means for today's IT center to cost effectively manage its increasing storage needs, while maintaining access to data. The patented data integrity algorithms, in concert with high-availability features, provide greater protection against data loss and the associated down time. The modular architecture allows addition of drives and JBODs over time, without down time. The configurability of the VA family arrays allows a sophisticated storage administrator to customize a solution based upon individual application needs, while the simplicity of the solution allows a naïve user to easily install and manage the array.

## glossary

| | |
|---|---|
| DFP | Data Flow Processor |
| ECC | Error Correcting Code |
| FSAM | Federated Storage Area Management |
| IOPs | Input/Output operations per second |
| IT | Information Technology |
| JBOD | Just a Bunch Of Disks |
| NFS | Network File System |
| OLTP | Online Transaction Processing |
| RAID | Redundant array of independent/inexpensive disks |
| SAN | Storage Area Network |
| SDM | Storage Device Manager |
| TPC-C | Transaction Processing Council benchmark of OLTP applications |
| VA | Virtual Array |

## for more information

For more information on HP's VA7100 or VA7400, contact any of our world wide sales offices, or visit our Web site at:

http://www.hp.com/go/storage.


11/2001