Hewlett Packard/Compaq Storage Area Networks Overview White Paper

By Jim Gursha
June 26, 2002

High Performance System Solutions, Inc.
WWW.SANDISKS.COM
Telephone: 212-831-0291

Introduction

This report discusses the enterprise Hewlett Packard (Compaq) virtual storage array and the Hitachi Data Systems Storage Area Network enterprise products. It delineates a general overview of features and functionality of both Storage Area Network offerings and does not attempt to determine a selection in favor of either vendor's products. This report was prepared by reviewing information from each vendor's respective Storage Area Network documentation that is publicly available.

Since Storage Area Networks (SAN's) are highly complex devices that are composed of multiple hardware and software elements, not all of the discrete pieces that comprise a SAN are discussed individually. For many of these elements that are not discussed in detail, references to the appropriate vendors' web pages have been provided.

Typically when a SAN is deployed some level of consolidation of resources can be achieved. In an environment where servers with direct attached storage (DAS) are being moved into a SAN, large amounts of storage that was "captured" within an individual server may now be available for the general use. In a SAN that does not utilize virtualization aspects, this storage must still be managed. However in a virtualization environment, available storage is made accessible with less management intervention.

Thus potentially large amounts of previously unavailable disk resources are freed up at a reduced overall infrastructure cost. Additionally other benefits such as easier backup/restores, availability of data, reduced management costs, and even more efficient server utilization may be achieved thereby further justifying the use of a SAN based storage paradigm.

Hewlett Packard/Compaq Virtualization

The enterprise level Hewlett Packard[1] (HP) StorageWorks Virtualization products will be all encompassing at every level of the enterprise, i.e. server level, fabric level, and SAN level. The focus of this section of the white paper is to address the capability of HP's StorageWorks enterprise solutions that are available today (not futures) at the SAN level. As such, the scope of this document will be limited to discussion of the HP StorageWorks Enterprise Virtual Array (EVA), a SAN level Virtualization technology that uses asymmetric (out-of-band) hardware and software to manage an EVA implementation.[2]

[1] All references to Compaq StorageWorks products will be referred to in this white paper as HP StorageWorks.

[2] For discussion on other virtualization topics, including Versastor concepts, the reader is referred to the StorageWorks portion of the HP web page, www.hp.com. (Once at this web page, the reader can click on the search of "Both", Compaq and HP, for "StorageWorks and Versastor").

Enterprise Virtual Array

The Enterprise Virtual Array (EVA) uses Versastor concepts to implement the hardware and software functionality available in today's release of the product set. According to HP StorageWorks literature, this product is targeted towards providing solutions at the enterprise client for all levels of implementation (medium through high-end). It is a newer technology offering within the HP (Compaq) enterprise family class of products and can work with a network of other HP SAN older Compaq models that do not offer virtualization functionality, e.g. EMA12000, EMA 16000 one gigabit per second fiber channel models.[3]

The high level architecture of the EVA SAN environment has paired virtual controllers, HSV110's, that support two-fiber channel loops per controller for a total of four fault tolerant loops. In the SAN, fiber channel fabric switches are used to provide for the connection of multiple servers that are typically heterogeneous in nature within an enterprise environment. (Currently a maximum of 20 fiber channel switches may be interconnected to form the fabric).

All major hardware components of the EVA SAN are redundant for fault tolerant operation. With the dual redundant fiber channel loops (each loop pair supports 120 disk drives) load balancing may take place. The largest numbers of disks that can be supported with a single EVA controller set (dual controllers) are 240. In contrast the controllers, HSG80 pairs, for the non-virtual SAN (EMA) supports only 84 SCSI based disks.

Servers are connected to the SAN with fiber channel host bus adapters. At this time, approximately 256 servers are supported. (Each server may have a maximum of four HBA's).

The EVA SAN is administered using a Management Appliance, which is essentially a separate computer system with its own software, that does not interfere with the SAN's data flow operations. (The management appliance(s) will be fiber channel connected into each fabric switch.) This out-of-band configuration allows the storage manager to use web enabled GUI software to graphically monitor and administer the SAN resources.

[3] A caveat to this functionality was found that states that the HSV and HSG element management software components of the SAN management appliance must, at this time, be on different SAN zones because the management appliance software of the different models is not currently compatible. Until HP StorageWorks indicates that all management appliance software components can be utilized together, the virtualization and other enterprise SAN's must have separate management appliances. According to new information received as of June 4,2002, this problem is fixed and this restriction is no longer applicable. However this resolution is not reflected within the most currently available documentation on the pertinent web pages.

The EVA SAN is configured with a minimum of two EVA HSV110 controllers that have two fiber channel ports per controller for a total of four 200 megabytes per second fiber channel ports. Up to one gigabyte of mirrored cache and support for up to 240 disks for a maximum of approximately 17.5 terabytes (raw space when the 72 gigabyte, 10,000 RPM disks are used) can also be configured. (Smaller 36 gigabyte 10,000 RPM disks are available). A client can configure a single cabinet to hold approximately 12 terabytes (a maximum of 168 disks) within a single cabinet.

According to the specifications, the EVA subsystem can be configured to a maximum of approximately 35 terabytes of raw storage utilizing 480 72-gigabyte disks and two EVA controller pairs across three cabinets. The HP StorageWorks enterprise model is configured with a minimum of two virtualizing controllers for redundancy. (More EVA controller pairs can be added as the solution needs dictate). The fiber channel disks within the EVA are configured between four redundant fiber channel loops (transfer ability of 200 megabytes per second, 2gigabits, per second) that are implemented for maximum high availability between all components.

The EVA has been measured to achieve approximately 168,000 input/output operations per second with a maximum throughput of 674 megabytes per second. The EVA SAN supports 2 gigabits per second fiber channel. However, this speed will only be supported after the initial product release when all SAN components can fully support the two-gigabit per second implementation.

The cache support of the EVA has several modes. For all disk write operations, the cache uses a default “write back” algorithm for the mirrored cache when everything is working normally. When the data is emplaced in the cache, it will not be immediately written to disk. Instead, the I/O operation will be indicated as completed back to the application at the host level, and the data will reside in cache until a point in time is reached when operational efficiency of the write is obtained. This mode is used unless the cache batteries are not sufficiently charged and cannot guarantee the integrity of the write I/O transaction if a power failure occurs. In this case, the write cache algorithm reverts to a write through algorithm, e.g. all writes of the data are written through to the appropriate disk area so that critical updated data is not lost. (The use of the write-through algorithm may involve comparatively stiff performance penalties.)

The read side of the cache algorithms involves various approaches. When the read-cache is turned on for a virtual disk, several different cache mechanisms may be used. The normal cache read approach is that data is moved from the disks into the cache during the read I/O cycle. If the host based application requests the same data again before it has been taken out of the cache, a cache based read will occur. Another feature, called read-ahead, moves data into the cache that is anticipated to be read sequentially after an I/O pattern of sequential reads has been established. By moving sequentially positioned data into the cache in anticipation of sequential I/O requests from the application continuing, better read performance is realized. A more advanced cache feature called “Adaptive Read Caching” uses inherent program algorithms to determine the benefit and efficiency

of moving data into cache based on previous read patterns exhibited by the application requesting the data.

The individual components of the EVA have a pair of HSV controllers that handle a group of physical disks that are housed in disk enclosures. For good performance, the number of disk enclosures should be evenly balanced between controllers. Each pair of controllers is connected to the actual disk drives with pairs of fiber channel arbitrated loops. (Two pairs of loops exist per controller pair for full high availability and redundancy.)

For redundancy, each HSV pair should be interconnected to two fiber channel fabric switches. With the present SAN guidelines, a maximum of 20 fiber channel fabric switches can be utilized and supported in approved configurations. These fabric switches can be configured in eight, 16, 32 and 64 port models. However the lower level models are indicated to only support one gigabit per second speeds. [4]

The EVA SAN implementation will allow a large number of client servers to be connected. The maximum number of connections is limited by the 256-host connection threshold for each EVA controller pair. (The maximum host connections can be further broken down to 1024 server based host bus adapters (HBA's)[5] of four HBA's per server. These fiber channel HBA's will be able to support up to 200 megabytes per second.)

The server and operating systems that are currently supported are HP/Compaq OpenVMS and Tru 64 Unix on Alpha servers with full cluster support; Sun Solaris (Veritas cluster support. Sun's cluster product is not currently supported) with Sun servers and Microsoft Window 2000/NT on Proliant servers. (Microsoft cluster support although supported by HP has not been officially approved by Microsoft and requires that the Secure Path software be used for host bus adapter high availability.)

When configuring multiple servers with different operating systems, heterogeneous implementation, a server is blocked from accessing information that is owned by a different server by using the concept of LUN masking also known as selective storage presentation. Selective Storage Presentation is the ability to grant or prohibit a server at the controller level from having access to a virtual disk. The ability to implement data zoning is also supported at the SAN level.

Once the SAN is built with its inherent switches, controllers, shelves and physical disks, the storage manager performs configuration functions. In order to consume data resources, the actual disks are formed into pools of available storage called disk groups.

---

[4] The supported switch configurations and the port speeds are constantly being enhanced and updated. The reader is cautioned that this white paper does not thoroughly address the HP StorageWorks approved SAN switch technology available today. The reader is referred to the appropriate section of the HP StorageWorks web page for more in-depth technical information on this subject.

[5] Within the Versastor environment, the software and hardware that allows a host to interconnect with fiber channel and communicate to a SAN is termed a Fiber Channel Adapter (FCA). However throughout the available literature the older term host bus adapter (HBA) is used interchangeably with FCA. For this paper, I have elected to use HBA.

A default disk group is actually set up when the SAN array goes through its initialization functions. More disk groups (a maximum of 16) can be created by the storage manager. The number of real disks within a disk group will delineate how much physical storage will be available for the virtual disks that are to be created. (You may use a minimum of four disks to a storage pool). Once the storage pools are available, the storage manager can create the virtual disks, which will utilize the available storage within the storage pool. Virtual disks can be created against a storage pool until all of the available physical storage has been assigned to virtual disks. When creating virtual disks, there may be from one gigabyte to two terabytes of space assigned for each virtual disk to a maximum of 256 LUN's. (A ceiling maximum that takes the total number of virtual disks times the number of hosts for a total of 8192 is enforced).

The EVA SAN's have implemented raid support as a virtual raid concept. This implementation takes all of the data structures of the user selected raid group, e.g. VRAID-0 (striping of data segments across the Virtual disk's storage pool), VRAID-1 (mirror set), VRAID-5 (striping of data with inherent parity segments across the Virtual disk's storage pool) and attempts to spread the data out evenly in 128 kilobyte disk chunk sizes among the actual disk drives that make up the designated virtual disk. When constructing a virtual disk, one of the VRAID types must be selected. Having a vanilla JBOD (just a bunch of disks) configuration is not allowed.

The storage manager can dynamically expand a storage pool by adding more physical disks. Once a storage pool is expanded, the data elements on the different virtual disks (VRAID-0, VRAID-1, VRAID-5) are automatically evenly distributed, load balanced, across the available disks of the pool according to their inherent VRAID type. This background movement of the data does not conflict with online operations and serves as a methodology to ensure the best performance possible within each designated storage pool. Two modes are available when adding physical disks to a storage pool. An automatic mode ensures that the disk group with the least amount of free available space will receive the newly added drives. A manual mode, SAN default, allows the storage manager to assign the physical storage to the disk groups, storage pools, in one-gigabyte increments.

The storage manager also has the capability to hold back enough disk space from all volumes of the disk group to accommodate the failure of an actual physical drive(s) within the storage pools. Once a physical drive fails, the controller software, VCS, will rebuild the VRAID-1 and VRAID- 5 data using the extra capacity that was reserved from utilization within the storage pool. A VRAID-0 recovery is dependent on the amount of warning that a drive will give the controllers that it is about to fail. Assuming that there is enough time the striped data will be copied to the available space.

Another set of backup/recovery tools of the storage manager is to use different types of snapshots and clones that are available. Under the EVA SAN environment, a snapshot may be taken of a virtual disk. There may be a maximum of 255 snapshots of individual virtual disks outstanding at any one time. A single virtual disk may only have one

snapshot allocated to it at a time. The ability to perform a snapshot that spans storage pools, disk groups, or crosses virtual disk VRAID types is not currently supported.

At present, there are two formats that a snapshot may take. The standard snapshot reserves an equivalent amount of space of the virtual disk. When the master, original, virtual disk's data is updated, the data update is indicated within the snapshot area. The second form of available snapshot, termed a "Virtually Capacity-Free Snapshot", does not utilize space until the data on the disk that is being snapshot changes. This type of snapshot shares the original disks metadata and data pointers until the data is updated.

The storage manager has the ability to make a clone, duplicate, called a "Virtually Instantaneous SnapClone". In this case, the data of the SNAP Clone is available to use because the pointers to the original data on the original virtual disk are available while the clone is being made in the background. (This process does not instantaneously create a cloned volume.) If a portion of the data is updated, the EVA controller keeps track of all changed data elements until the clone is complete. Most standard clone procedures do not allow read and write access to the clone volume until the actual copying of the original volume has been completed.

Today the above tools are the only way to ensure data protection. The data replication to a remote SAN that is available on the EMA arrays is not yet released for the EVA SAN's. It is anticipated and expected that the ability to perform wide area data replication using network extension hardware will be available in a release in the very near future.

Hitachi Data Systems

Hitachi Data Systems (HDS) manufactures a Storage Area Network product family that offers solutions from the departmental through enterprise level of data management. This product family essentially can be broken down into three product groups, i.e. departmental through small scale enterprise customers, large departmental to medium enterprise and large to ultra-large scale enterprise. The lower end targeted client tier can be adequately serviced by the 9200 product line while the middle level is delineated by the larger 9200 configurations and smaller models of the 9900. The higher tier and beyond are serviced by very large SAN configurations. Many HDS models when properly configured and managed offer a 100% data availability guarantee.

In May 2002, HDS announced a new line of medium to exceptionally high-end SAN products as part of an upgrade to the 9900 product line. This announcement included an "open" storage management software product that allows for the administration of multiple vendors' heterogeneous storage platforms within a single infrastructure. This new product set, which is delineated as the 9900V series, introduces storage virtualization into the Hitachi family with the 9970V and the 9980V models. These products, which will eventually replace the older 9900 models when they are placed in the end-of-life

category, significantly expand the delivery and performance capability over the 9910 and 9960 platforms.[6]

All of the HDS SAN's support Network Attached Storage for integration of both SAN and NAS within the same physical storage hardware platform. Although not a specific topic of this white paper, there is a major movement within the storage community to strongly integrate NAS into the overall SAN structure.

The Hitachi Data Systems SAN's support both McData and Brocade fabric switches. The storage designer needs to consider the design and implementation architecture solutions and use the switch most suited for the overall solution. Although important, these concepts are not addressed in this white paper analysis.

In general the architecture of the HDS Thunder 9200 is more comparable to the StorageWorks fiber EVA array. However, the 9200 allows larger cache configuration options. The HDS philosophy is that a larger cache allows the server based applications to perform larger transactions faster because more data can be managed in the cache and the database designer or data manager has the option of designating portions of cache for database tables or large files. (In the 9200, the two controllers can have two gigabytes of cache for a total of four gigabytes maximum. Communication between the caches in each controller is sustained by a 235 megabyte per second bus to maintain the necessary fault tolerant information updates.) By using read-ahead tailored algorithms, the larger cache enables the 9200 to maintain more information within the cache thereby speeding up overall I/O operations within the SAN.

From a high level perspective, the 9900 series architecture is also highly "cache-centric". Data is received into the array through microprocessor managed SAN storage ports, and sent directly into the cache. (User data and control instructions are kept in separate caches.) The data is moved into cache through Cache Switches. All data paths between the Cache Switches and the cache are non-blocking for high-speed data movement. All data requests from the application servers' hosts are serviced by retrieving data from the appropriate cache area and not directly from the actual physical disks themselves. If the requested data items are not available in cache, they are retrieved from the appropriate disk locations and staged into cache memory. These data items are accessed through backend paired Array Control Processors, which manage their inherent ports with microprocessors, that connect multiple fiber channel loops with dual ported fiber channel disks.

The newer HDS SAN's can manage more information on their disks with better overall performance because of the availability of larger cache configurations, more powerful front and backend processors, sophisticated data management capabilities, and a larger number of faster data paths. When combined with their virtualization ability to provide

[6] As far as can be determined, the 9900 and the 9910 are slated to be placed in the end-of-life category at the end of 2002. It is anticipated that these models will be serviced and maintained according to Hitachi Data Systems' product and maintenance policy.

larger scale consolidation of servers with all of the inherent benefits, these higher end SAN's offer extremely high returns on investment.

On an informational level, the HDS product family can be broken down as discussed in the following sections of this white paper.

Departmental SAN (Thunder 9200)

This departmental Storage Area Network solution is targeted towards clients that have a need for "open systems solutions" (UNIX, Microsoft 2000/WNT and Novell Netware) including email, relational databases, with small to medium applications and business continuity requirements that are limited in scope and distance. These storage systems are anticipated to have a data growth pattern that will be shared amongst a reasonable number of servers and will not exceed a hardware maximum of 100 disk drives. Generally, the 9200 series can be configured with one or two controllers, four fiber ports (each controller has a maximum of two fiber ports that may be configured at a port speed of either one or two gigabits per second), and up to 10 disk drives per disk container. (Initially disks are placed in the main control unit, which is referred to as a command module or RK unit. Additional drives are added into modules called RKA units where each unit can hold up to 10 disks).

The disk drives are dual ported and fiber attached to four fiber loops (each controller holds two fiber loops). The 9200 array can have a maximum of four gigabytes of mirrored cache (each controller has two gigabytes). Each 9200 controller is managed by a power PC and has two specialized Tachyon microcomputers at the front and back-ends for a total of four per controller. The aggregated throughput can be summarized as approximately 2.4 gigabytes per second. The 9200 also supports multiple Raid configurations, Raid-0, Raid-1, Raid-0+1, and Raid 5.

Although at first glance this medium tiered unit can physically hold approximately seven terabytes of information. (This capacity number is a "raw" estimate derived by assuming the configuration uses the HDS disks drives which are approximately 73 gigabytes and 10,0000 RPM. [7]) However, a fully configured SAN may introduce potential performance issues. This possibility depends on how the customer's selected raid choice has been configured and the storage manager's perceived performance expectations of the overall application workload that is being serviced. However when properly designed and implemented, the total throughput expectation for a 9200 model SAN is very good.

For a departmental to intermediate product, the 9200 has a reasonable amount of layered product availability. Several of these products are summarized and discussed within this document. The storage manager has the ability to make multiple copies of data within a local array for on-line redundancy of data. When properly utilized, the local data copies can be specified as point-in-time copies so that data volumes can be very quickly restored on an as needed basis without a major impact to the overall throughput. (HDS

[7] The user can configure 18GB and 36GB, 15K RPM, or ~73GB, 10K RPM, or ~ 181GB, 7200RPM disks.

ShadowImage utility). In addition to the local copy utility, data replication can be used in either synchronous or asynchronous mode. (TrueCopy Utility) If synchronous data replication is selected, the standard distance limitations apply (approximately 60 miles). For longer distances, asynchronous replication is used to move data. When employed in conjunction with network extension capabilities such as Nishan's multi-protocol storage switch, data can be moved very large distances. (Nishan has tested in the lab with distance simulators data replication to 12000 KM. They have in concert with Qwest and Hitachi Data Systems performed a cross-country proof of concept, which used the 9910 SAN, designated as the Promontory project).[8]

Achieving acceptable performance from the application and end user's perspective when doing asynchronous data replication requires a thorough understanding of all elements involved. A successful implementation will require configuration of the proper high-speed data lines that will allow the amount of data to be replicated within time frame parameters that are acceptable.[9] With all of this technology properly integrated, remote data replication to another physical array may be used to support a redundant backup facility or disaster tolerant data center implementation.

The Flash Access layered product allows for data (determined by LUN) to be locked into the cache. A database designer can tell the storage manager that a certain element of data will need to be frequently accessed very quickly. By making the data memory resident in the data cache, fast access can be ensured.

The 9200 SAN has the ability to report back to HDS corporate diagnostics about the condition of the SAN and field service may be dispatched to perform proactive repairs. Additionally LUN security and expansion are supported as well.

The 9200 may be managed by multiple software products. One product called the Hitachi Resource Manager can be used for managing a small number of 9200's. The other product, Hi-Command, is used for large scale, heterogeneous storage management on a more sophisticated scale. [10] This layered product allows the storage manager to use a web enabled GUI application to perform SAN necessary management functions such as utilization of supported layered product software, i.e. local volume copies (ShadowImage), locking of data elements into cache (Flash Access), LUN security (Santinel – control of data access using world wide names and port id's), and dynamic, non-disruptive LUN expansion, and management of important SAN maintenance

---

[8] For more information please refer to
http://www.nishansystems.com/promontory/Promontory_White_Paper_NSWP_11.pdf
09/23/01, 138506 bytes
11/08/01, 3120 bytes

[9] For detailed information please refer to the Hitachi Data Systems web page, www.HDS.com, and the Nishan web page, www.nishansystems.com.

[10] A more detailed explanation of Hi Command can be referenced on the HDS web site.

functions, e.g. non-disruptive micro-code updates, fabric topology and LUN creation and mapping.

Entry Level to Middle Tier Enterprise SAN's (Lightning 9910)

The middle tier of the SAN product set consists of a low end 9900 model called the 9910. This model supports the IBM mainframe (S/390), email, relational database activity, data mining capability, enterprise and departmental applications, local, and wide area business continuity services, and comes with an uptime guarantee, 99.999%, and a data guarantee. The 9910 ships in a single cabinet and supports up to 16 gigabytes of mirrored cache, 24 ports, 48 disks and approximately three terabytes of raw disk capacity when the 73 gigabyte disks are used. (A total of approximately 2.4 terabytes of useable disk space is available for RAID-5 implementations). The 9910 has a Hi-Star switch implementation (similar to a crossbar switch) like the 9960 and its aggregated throughput performance is rated at approximately 3.2 gigabytes per second. (The Hewlett Packard XP48 is the version of the 9910 model that is manufactured by and OEM'd[11] from Hitachi Data Systems.)

The 9910 is a lot of package for a small footprint and represents good performance verses its price/performance/functionality comparisons. However if the price and functionality are compared to a 9200, it can be shown that the 9200 model SAN might be a better choice for specific application solutions. Since the 9910 can not be easily expanded into a 9960, a storage architect considering the 9910 needs to review and answer several very important questions a few of which are represented below:

1. Will the long-term requirements of the data storage growth pattern require an upgrade before this implementation can be cost justified?

2. Will a similarly configured 9200 SAN with layered products perform at least equal to the 9910?

3. Will a connection with a mainframe, IBM S/390, be required at present or in the future?

In answering the first question, it is quite apparent that a storage requirement of close to the useable 2.4 terabytes (Raid-5 configured) will not suffice to cover a reasonable growth expectation. In this case, a storage manager would either have to size a higher end SAN or consider the 9200 implementation. When the first two questions are taken together, it is apparent that if a reasonable performance expectation can be seen in the

[11] OEM – Original Equipment Manufacture. At present, information is not available that details if HP makes modifications to this HDS model.

less expensive SAN solution (9200) and the client's applications will not require interconnection with a mainframe, an easily expandable configuration should be chosen. Additionally if we consider the manner in which the mirroring has been implemented on the 9200-product line, we see that although more disks are required over the Raid-5 implementation the Raid-0+1 performance should be better. This performance improvement can be seen in the cases where the application throughput benefits from HDS's implementation of its mirroring functionality on the 9200 series.

Upper Middle to Ultra High End Enterprise Tier (Lightning 9960, 9970V and 9980V)

The middle to upper end of the enterprise space can be said to have some degree of overlap with the previous SAN models depending on the definition of what an enterprise application is and the minimum number of terabytes needed to qualify for an enterprise title. However what is very clear is that the HDS 9960 and 9970V fit in to the high end of the enterprise definition while the large scale 9980V is the SAN within the product line mix to service what will be defined as the ultra-high end enterprise level storage applications.

Quickly reviewing these different models we find that the 9960 supports approximately 37 terabytes of raw disk storage (assuming 73 gigabyte drives are configured), over 4096 logical devices, 32 ports, up to 512 dual-ported fiber attached drives, Hi-Star (crossbar) switches that interconnect cache with front end and back-end systems. Data and instructional control sequences are moved via separate paths. A 1.5 gigabyte control store instruction cache is available for an aggregated throughput of approximately 6.4 gigabytes/second. (The Hewlett Packard XP512 is the version of the 9960 model that is manufactured by and OEM'd[12] from Hitachi Data Systems.)

The previously discussed layered products are also available plus different software options along with a very sophisticated management application. Prior to the recent HDS announcement, the 9960 was considered the flagship of the HDS SAN line and was used to configure complex heterogeneous SAN solutions. With the introduction of the new virtualization features on the 9970V and 9980V, the 9960's importance although not diminished will no longer be considered the ubiquitous solution.

With the introduction of the V-series (virtualization assisted), HDS has introduced two new products that provide "virtualization assist" along with the ability to connect multiple servers at the logical level. By pooling the servers and using Host Storage Domains that have their own logical fiber channel port and an independent set of LUN's, this approach allows more than one server to logically share a port simultaneously. The Host Storage Domains in concert with several layered HDS products that provide security and give specified applications a priority access at the port level is termed a virtualization assist. With this new technology, HDS has provided a solution that when correctly

[12] OEM – Original Equipment Manufacture. At present, information is not available that details if HP makes modifications to this HDS model.

implemented at the client site will be able to increase the ability to consolidate more servers than the older product line was capable of efficiently handling.

Both the 9970V and 9980V benefit from a re-designed and enhanced Hi-Star (crossbar) switch architecture. This new implementation has doubled the number of discrete point to point connections within the Hi Star switch; has increased the aggregated bandwidth to approximately 7 and 15 gigabytes per second respectively; has significantly increased the speeds of both the front and backend associated processors (64 bit MIPS processors) within the SAN; and has enhanced many of the software offerings to take advantage of the new architecture.

The 9970V supports 24 ports, 32 gigabytes of data cache, three gigabytes of instruction control cache, 128 disks drives (approximately nine terabytes of raw space when using 73 gigabyte disks), 8192 logical devices and at a minimum is promised to at least match the performance of the 9960 within a smaller physical footprint. (The Hewlett Packard XP128 is the version of the 9970V model that is manufactured by and OEM'd[13] from Hitachi Data Systems.)

The 9980V supports 32ports, 64 gigabytes of data cache, three gigabytes of instruction control cache, 1024 disk drives (approximately 75 terabytes of raw space when using 73 gigabyte disks), 8192 logical devices and is promised to at least double the overall performance throughput of the 9960. (The Hewlett Packard XP1024 is the version of the 9980V model that is manufactured by and OEM'd[14] from Hitachi Data Systems.)

Each of the physical fiber ports of the "V" model SAN's can have 128 virtual ports (Host Storage Domains) associated with them. Each fiber port has a physical speed of two gigabytes per second, approximately 200 megabytes per second. If you carve up the ports using an assumption that most relational database I/O will get a maximum throughput average under five megabytes per second, you can then determine how to best allocate the servers amongst the unassigned Host Storage Domains (virtual ports). There are 4096 Host Storage Domains (virtual ports) total for the 9980V and 3072 HSD's for the 9970V. It is apparent that the storage manager will have to completely understand the applications and server distribution across these virtually assisted arrays to get the best performance ratios.

With the security associated with each HSD and its LUN's, multiple operating systems may now be associated with each physical port without inherent conflict. This ability is different from many other SAN vendor's implementations that do not allow this type of access. Thus the new Hitachi Data Systems virtually assisted SAN's give the storage manager the ability to consolidate servers, storage and even other SAN's onto a single, highly reliable, 99.999%, storage hardware. If this implementation is then designed with a highly available data-center implementation where multiple virtually assisted SAN's

[13] OEM – Original Equipment Manufacture. At present, information is not available that details if HP makes modifications to this HDS model.
[14] OEM – Original Equipment Manufacture. At present, information is not available that details if HP makes modifications to this HDS model.

are interconnected across a wide area network (perhaps using a Storage Over IP implementation), a truly disaster capable architecture can be realized.
The major utilities available on these newer arrays still function similarly to their counterparts on the 9900 series, although some of them have been augmented to take advantage of the higher density of logical devices and better internal speeds of the new equipment.

The new virtualization assisted SAN's support the following operating systems. (For an up to the minute support list, please refer to the HDS web page.

Mainframe

IBM® OS/390®
z/OS®
MVS/ESA™
MVS/XA™
TPF

Open Systems

HP®	HP-UX® (10.2, 11.0, 11i, 11.20)
Sun™	Solaris™ 2.6, 7 and 8
Microsoft®	Windows NT® 4.0/Windows® 2000
IBM	AIX® (4.3.3 and 5.1)
Compaq®	Tru64™, OpenVMS is currently supported on the older 9900 product line. It is anticipated that the 9970V and 9980V will support the latest OpenVMS versions in the near future.
SGI™	IRIX®64 (6.5.13)
Red Hat®	Linux® 7.2
Sequent®	DYNIX/ptx® (later release)
Novell®	NetWare® (5.1 and 6.0)

Summary

The HP EVA and its associated SAN based elements allow a virtual presentation of storage to be efficiently utilized and managed by a web based GUI interface. A high performance SAN solution can be designed and implemented that favorably compares with the Hitachi Data Systems 9200, 9960 and 9970V SAN offerings. Both vendors' products have GUI based management and offer layered product utilities for ease of use, performance improvement and heterogeneous SAN administration. These products also include backup/restore and server-less backup options to high performance fiber channel connected tape libraries for offline storage.

The HP integration strategy appears to be oriented towards keeping the strongest of the storage models and consolidating them into a cohesive product line. They are focused on building a NAS/SAN technology integration where the resulting product will offer the customer seamless access to either storage paradigm. According to HP, both the OEM'd HDS and HP StorageWorks models have their place in the overall architecture strategy for the near future. [15]

[15] For further information regarding an HP roadmap overview, please refer to the following link: http://www.encompassus.org/news/roadmap.html