HP StorageWorks virtualization

White paper

Table of contents

Executive summary

Storage virtualization has been around for several decades. While technologies have advanced and capabilities have become more sophisticated over time, the motivations for virtualizing resources and the anticipated benefits have remained relatively stable. Storage virtualization continues to be an important topic because it addresses a number of significant challenges. Storage vendors, the media, and other sources expound on virtualization's merits in ways that might lead one to conclude that the term has a single, agreed-upon meaning and value—that perhaps one size fits all. Indeed, the range of storage virtualization technologies and solutions has continued to evolve as storage hardware, the software that powers them, and external storage management applications have matured. This can often create confusion. Customers look to suppliers like HP for guidance and leadership. Thus, vendor virtualization approaches and strategy are important.

Virtualization is the aggregation of physical resources into a unified structure ("pool"), and the presentation of those resources as capabilities that can be consumed by applications or other types of storage clients. The term "virtualization" has really become an umbrella that encompasses a vast array of encapsulation and management technologies and implementation methods. These approaches create pools of sharable resources that enhance utilization and, ultimately, can automatically allocate the resources to match supply to demand.

Business motivations for virtualizing resources vary, and generally fall into broad categories that include:

- Reduce infrastructure costs by improving asset utilization and simplifying the environment
- Reduce ongoing costs by simplifying administration and increasing worker productivity
- Generate incremental revenue by improving application productivity by way of uptime, performance, and other enhancements

Virtualization addresses these needs, often yielding significant economies. For example:

- Virtualized resources create pools that enable data to be consolidated into smaller numbers of storage systems, resulting in higher utilization and management efficiencies.
- Virtualized pools can simplify the sharing of data among multiple systems.
- Data migration among pooled resources can be done with less administrative involvement and little-to-no impact to running applications.
- Virtualization helps to standardize management models and the capabilities delivered by pooled resources, in turn simplifying overall storage management.
- Replication and other virtualization technologies enable higher service levels that directly translate into user and business application productivity and lower downtime costs.

It is important to note that virtualization extends beyond storage. While the focus of this paper is storage virtualization, one can also virtualize server, application, network, and other resources. HP offers a broad range of virtualization technologies and solutions that are not discussed in this paper.

This paper describes the business motivations for storage virtualization and the benefits it can deliver. Basic storage virtualization approaches are described and their benefits and limitations are compared. Finally, HP StorageWorks virtualization offerings strategy are described. The paper also briefly discusses the enterprise context for storage virtualization—a global implementation that can be delivered over time with the HP Adaptive Infrastructure.

Introduction to storage virtualization

Storage virtualization has been around ever since IBM released the first disk drive in 1956, although some might assert that core and magnetic drum memory were also primitive forms of virtualization. In the context of disk drives, the physical location on which a piece of data was placed was obscured to the application writing or reading the data—disk drives keep indices that track data-to-physical block mapping (abstraction). Disk drives also have the ability to change the mapping to accommodate minor failures ("revectoring" or "bad block replacement" mechanisms). Today this is not even thought of as virtualization—it is totally accepted, invisible, and "uninteresting" from an IT perspective.

Figure 1. HP has a long history of storage virtualization innovation

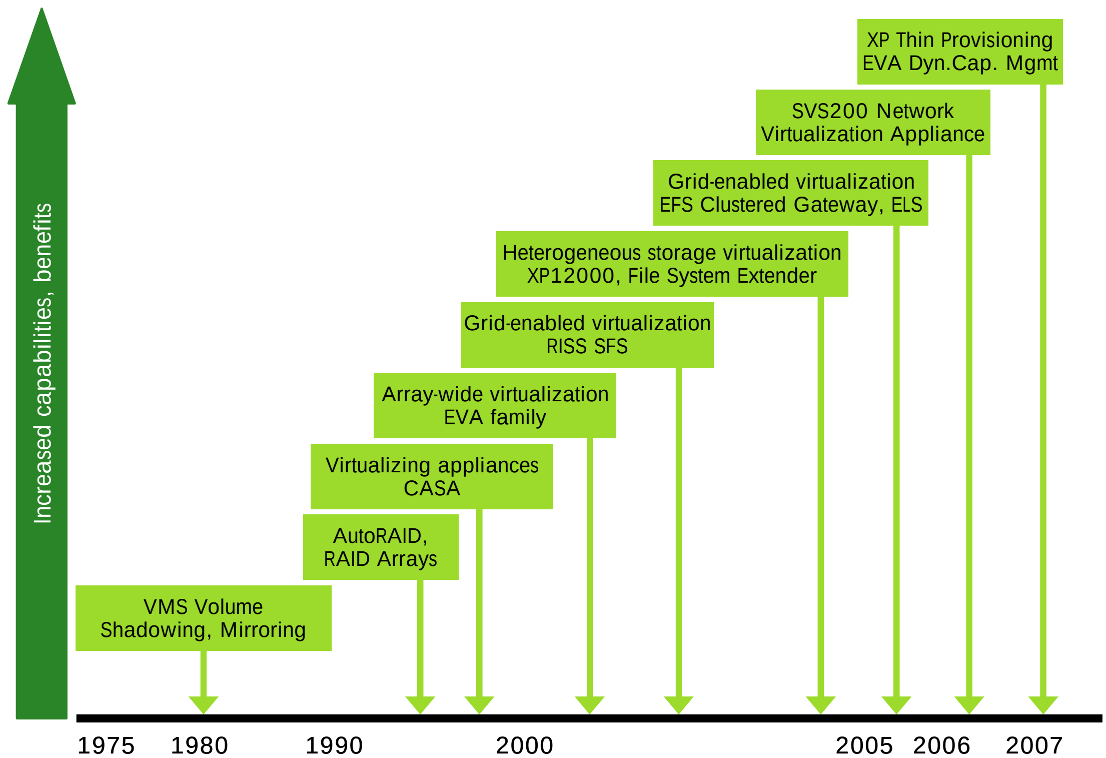

Figure 1 shows a representative evolution of more interesting forms of storage virtualization that HP has delivered. Initially, various forms of RAID implementations were shipped. The first was VMS Volume Shadowing—a host-based (implemented with server-resident software) RAID 1 implementation that was delivered several years before the term "RAID" was invented. After more comprehensive RAID algorithms were developed, HP incorporated them into modular array families like the RA/EMA and SmartArray product lines. Even more interesting, the VA family incorporated a self-adjusting RAID mechanism called AutoRAID—an industry-leading approach that automatically migrated data onto the most appropriate level of RAID within a subsystem, as well as automatic and dynamic RAID selection. Also, the VA7400 incorporated online dynamic capacity expansion—a harbinger of virtualization capabilities that started entering the mainstream in 2007.

All RAID schemes virtualize by creating a single pool consisting of a number of disk drives. Application data is algorithmically deposited across the disks by array controllers. In some cases (RAID 3, RAID 4, RAID 5, RAID 6), one or more disks' worth of parity data is also generated by the array and stored on its constituent disks to provide resilience to disk failures. Although RAID was initially invented to allow inexpensive commodity disk drives to reach the availability characteristics of higher-cost proprietary drives, the enduring benefits of RAID include:

- Higher and more uniform device (disk) utilization
- Better performance, achieved through automatic load balancing among the disks aggregated into the RAID
- Improved data availability, because built-in redundancy enables greater uptime
- Reduced probability of data loss, because loss of a disk drive results in no loss in data (although a performance loss may occur)
- Easier manageability: one RAID is easier to manage than the corresponding number of independent disk drives

In 2001, HP extended RAID implementation to encompass the entire contents of a storage subsystem, as opposed to creating multiple physically discrete RAIDs within a single subsystem. The result was the HP StorageWorks Enterprise Virtual Array family (EVA). The EVA amplified the benefits of RAID to encompass a complete system that could enhance the utilization of many disks simultaneously, and warrant even greater availability through the proprietary application of RAID algorithms.
The system-wide approach also made EVA easier to manage than peer products.

Remarkably, the EVA extended storage-based virtualization (virtualization embedded within the storage system) in interesting and useful ways. It transcends conventional RAID by delivering "advanced storage-based whole-system virtualization" with capabilities like:

- Dynamic expansion of the virtual disk pool: disks inserted into a configured array are automatically added to the existing virtualized pool (by the EVA controller), and the workload is automatically redistributed to all members—without administrative involvement
- Dynamic expansion of the virtual disks ("LUNs") presented to servers, although the ability to dynamically present the new capacity is dependent on host operating system ability to recognize the changed LUNs
- Snapshot and other point-in-time replication technologies: these "in-the-box" replication technologies create virtual images—copies of the LUNs—and can be used for rapid data protection and recovery, and for time-shifting backup operations
- Remote replication—with synchronous and asynchronous options—invokes mirroring and other technologies to provide resilience against site outages, as well as being useful for remote backup and other purposes

While general attention has been focused on disk storage virtualization, it is useful to point out that tape libraries, Redundant Arrays of Independent Tapes (RAIT), and other technologies apply virtualization to tape-based storage.

As the millennium changed, the industry began delivering network-based virtualization capable of pooling heterogeneous, multi-vendor storage systems using virtualization technologies embedded in the fabric (SAN or LAN) to which the storage resources are attached. HP offered the HP OpenView Continuous Access Storage Appliance (CASA) to aggregate multiple disk subsystems into a single pool for data migration, capacity expansion, management simplification, and other capabilities. Network-based virtualization seemed to be ahead of its time back then, and faded from the limelight until recently—in part for reasons discussed later in this paper.

Around 2003, a new genre of storage systems was introduced with the HP StorageWorks Reference Information Storage System (RISS), which was recently renamed as HP Integrated Archive Platform. The virtualization employed in the HP Integrated Archive Platform can be thought of as a hybrid of storage-based and network-based virtualization, because grid-enabled subsystems are composed of a number of interlinked, networked processor nodes that together present an interestingly useful storage system. HP Integrated Archive Platform leverages grid-computing principles to deliver a virtual pool of storage that scales well, performs consistently over a huge capacity range, and incorporates higher-level intelligence to perform a variety of business application-centric tasks. In addition to HP Integrated Archive Platform, clustered storage applications running on integrated infrastructures—such as the HP StorageWorks Enterprise File Services (EFS) Clustered NAS Gateway—employ virtualization to deliver a variety of scale-out storage capabilities.

Today, storage virtualization technology evolution is largely focused on:

- More efficient storage utilization and provisioning, driven by the desire to reduce energy costs and move to a “pay-as-you-grow” costing model
- Virtualization of heterogeneous, multi-vendor storage assets, motivated by the need to move data among storage systems—for example, when an aging system is replaced with a new one

When old and new storage systems are combined into the same pool, network-based virtualization can be used to:

- Present storage to business applications
- Migrate data among storage systems using mirroring or other techniques
- Logically substitute the new array for the old when data migration is complete
- Accomplish the entire process almost non-disruptively to running business applications (there may be slight downtime). There is no technical need to perform backup operations in conjunction with migration activities, although business practices may require it.

Because of the usefulness of inter-system data migration, HP released the HP StorageWorks 200 Storage Virtualization System (SVS200) in 2006, which was later replaced with the diskless XP20000. Similar capabilities, integrated with traditional storage-based virtualization, are found in the HP StorageWorks XP10000, XP12000, XP20000, and XP24000 Disk Arrays. In both cases, heterogeneous and multi-vendor storage systems are unified into a single pool. The pool can include a range of storage systems representing tiered storage. Application-usable volumes are carved from the pool, and administrators focus their attention on the SVS200 or XP array, rather than the individual systems comprising the pool.

2007 saw the broad introduction of virtualization capabilities that allow individual storage systems to scale as needed, and to dynamically change the amount of capacity they present to applications. Collectively referred to as “thin provisioning”, these capabilities allow a disk array to create and present a logical disk (LUN or LU) of a given size without requiring physical storage capacity equal to the presented capacity. This means that a storage system can be configured with less disk capacity than it is ultimately expected to need, and then all LUNs required are created with whatever capacity is needed at the time. Later, if a LUN needs more capacity, the array controller can dynamically add more capacity (as disks are added) without changing the LUN’s presentation to applications. This is a huge step in storage efficiency, and it allows the system to grow in a very cost-efficient manner. HP introduced this capability in the XP20000 and XP24000 arrays. A similar capability, called Dynamic Capacity Management, was introduced in the EVA family.

It is clear that HP has been on the forefront of storage virtualization for decades. The next section looks at what virtualization is, why it is important, and how it can be implemented.

Virtualization overview

Basic principles

Several concepts pervade the many forms of storage virtualization. At a high level, virtualization is the abstraction of physical resources: storage virtualization technologies create logical views of storage that are distinct from their physical components. Resources are pooled, or aggregated, to encapsulate their resources and capabilities in useful and meaningful ways: ultimately this provides a foundation for deploying, managing, allocating, and delivering storage capabilities as services. One could think of virtualization as producing virtual pools of resources that can be provisioned and managed as needed. Practically, this results in insulating application access and utilization of storage from infrastructure configuration and management. In turn, storage can be manipulated with little-to-no impact on business application productivity.

A key goal of virtualization is to remove the complexities inherent when scores of individual components are gathered within an infrastructure, and to provide efficiencies through standardization. Standardization reduces the visible differences between entities to create efficiencies within an infrastructure, as well as greater utilization efficiency and flexibility. In this context, virtualization ultimately helps to simplify administration.

For IT managers, virtualization creates a "common ground" upon which management and control can be exercised. This can be thought of as using storage management technologies to standardize the infrastructure by amalgamating disparate components. For example, in an EVA, virtualization allows one to manage the EVA as a whole, rather than as a collection of physically distinct RAIDs or FC and Fibre Attached Technology Adapted (FATA) disk drives. Similarly, an XP24000 or XP20000 Disk Array with attached external arrays allows one to manage the tiered storage behind it as a single entity, rather than needing to manage each attached array independently. Conversely, a single XP system can be partitioned into a number of logically distinct virtual arrays if needed for security or other purposes.

At a higher level, virtualization applied at the network level ("network-based") enables unified, common management of heterogeneous arrays, and facilitates data migration and replication among arrays. The concept of neutralizing management differences with virtualization also applies to provisioning, hierarchical storage management (HSM), remote replication, and other operations. This illustrates that virtualization applies to more than disk drives—it can actually encompass storage management and other functionality that extends the abstraction of core physical resources.

Virtualization can occur at many levels in the data processing chain. In addition to virtualizing storage resources and capabilities, higher-level functionality can add automation and adjust provisioning, load balancing, paths through networks, and more. Leaving the storage domain, server resources can be virtualized. This is commonly used to partition today's powerful servers, using virtual machines or other technologies, as a way to enhance server resource (capacity and memory) utilization. Applications and other resources can also be virtualized, as is done by clustering, grid-computing, and other service-oriented architectures. This paper will not cover non-storage virtualization, but it is useful to be aware of the breadth to which virtualization technologies can be applied.

Importantly, virtualization functionality can be thought of as a many-to-many relationship in that there are many virtualization technologies, and they can be implemented in various places in the data processing flow to create highly usable solutions. The technologies are often deployed in a complementary fashion with the ultimate goal being to deliver the greatest benefits to I.T. productivity.

Why virtualize?

Today, IT managers are challenged to create more efficient and productive infrastructures, to deliver more reliable services (thereby enhancing application and business productivity and agility), and to reduce overall costs—both relating to infrastructure and ongoing administration. Clearly, the natural tension between the dual goals of higher efficiency and lower cost is difficult to resolve using conventional approaches. This is where virtualization provides substantial help. Figure 2 summarizes how storage virtualization can impact costs and productivity. This is the value proposition of storage virtualization. By mapping individual business (or IT department) costs into the model, the value of virtualization to the organization can be estimated. This is an important concept: the value of a solution should be greater than the total cost (not price) of the solution. The figure can aid in identifying the relevant components of a solution, and their associated costs to the company. HP can provide tools and expertise to determine the TCO and ROI of a solution: Figure 2 and the following discussion provide a starting point.

Figure 2. StorageWorks virtualization value proposition: storage costs (green boxes), facilities (gray), administration (blue) and business productivity (orange) can all be affected by storage virtualization

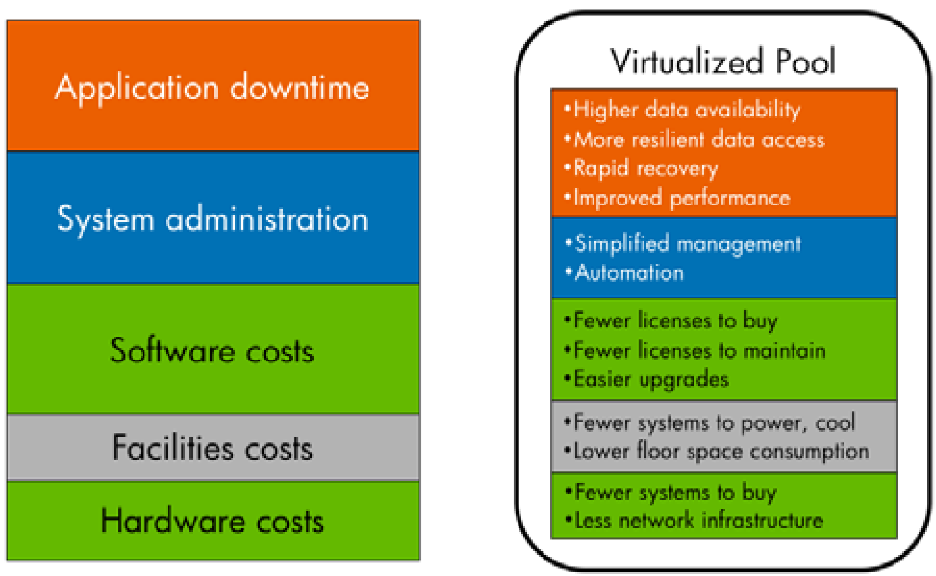

Many factors contribute to efficiency. Asset utilization is a key focus area. Total asset costs are sometimes difficult to precisely determine, as they may be components in a number of categories such as:

- Acquisition—Costs associated with obtaining and preparing the equipment for use. These costs include procurement and installing the equipment, configuring and testing it before use. They may also extend to the cost of software licenses and their deployment.
- Maintenance—After warranty expires, service and other costs continue
- Facilities—Storage subsystems consume floor space, power, and cooling on an ongoing basis. While footprint, power consumption, and management advances are expected to reduce facilities' requirements, virtualization can help even more by reducing the total number of storage systems needed to do a given amount of work.
- Management—Administrative costs associated with provisioning and fine-tuning the system on an ongoing basis. These costs may include storage, server, network, and application administrators.
- Disposal—When the useful life ends, cost is incurred during the disposal process

Virtualization technologies reduce administrative costs while improving business application and user productivity. For example:

- Workload balancing can enhance capacity distribution among pooled elements. This can significantly increase capacity utilization compared to manual provisioning—with a resulting reduction in the amount of physical storage, and associated software licenses, that must be deployed. Capacity utilization may increase significantly when workload distribution is improved among disk drives, arrays, and so on.
- Workload performance distribution results in enhanced load balancing among pooled assets, and can reduce the number of array controllers and network infrastructure needed to satisfy application demands. Workload distribution and optimization throughout the storage network—which includes paths through multiple HBAs in hosts—can contribute toward reduced LAN/SAN/WAN networking administrative overhead. Also, when multiple host connections and appropriate software are deployed, data access (availability) is improved. This has a direct positive impact on application and user productivity.

Operational efficiencies are also important. While these can be difficult to quantify, there are at least two key considerations:

- Reduce administrative costs, for example, by increasing the amount of storage an administrator can effectively manage
- Improve service delivery levels to increase overall application productivity

Storage virtualization can be thought of as providing a level of standardization among the pooled assets. For example, a pool may contain several different storage arrays—perhaps from different vendors. In a non-virtualized environment, each type of array has its own operational characteristics, management software (configuration utility and interface), and requirement for specialized operator understanding of the array's operating characteristics. When these assets are virtualized, a software layer intercedes to manage the arrays ("pool their capacity"), and presents a single interface for the administrator to interact with: it reduces the number of objects to manage and effectively brings all of the arrays to a common level of functionality. Also, application access to virtual disks within the pool may be simplified. This can be a huge simplification for storage provisioning management—by reducing the number of administrative tasks, for example. This sort of simplification can significantly extend the amount of storage an administrator can manage. As shown later, this type of virtualization can be found in the HP StorageWorks XP array family.

From a business perspective, a central focus for IT is to keep applications running (business productivity), and the quality of service (QoS), or service level expectations (SLE) of the storage farm plays a significant role. Thus Service Level Agreements (SLAs) are created, and a sufficiently robust infrastructure to enable service level objectives (SLOs) to be met—along with appropriate policies and practices—is put in place to try to meet the requirements. Storage virtualization helps by providing a unified foundation—the storage pool—that distributes and balances resource utilization in predictable ways. Upon the common pool, higher levels of management software can work to improve SLO attainment through failover, data replication, data recovery, and other mechanisms that operate across the pooled assets. The result: higher probability of meeting SLOs, with resulting measurable application productivity. Productivity measures include reduced downtime and higher, predictable performance for all applications using resources from the pool.

The discussion so far has focused on the direct results of virtualizing storage today. Derivative factors also could be considered. For example, with fewer arrays in the environment (because of higher capacity utilization), backup and recovery tasks can become simpler and may be less time consuming. Also, matching storage usage to data usage can further reduce costs. HSM and HP Information Lifecycle Management (ILM) solutions are relevant examples. In the future, automated service-oriented storage provisioning methods will further reduce infrastructure and administrative costs while improving business application productivity. One can think of these as providing application- and data usage-oriented virtualization, which can take the form of allocating storage from tiered physical resources (different QoS sub-pools within a larger aggregated pool, for example). Note that the ability to dynamically link all IT resources to business requirements is part of the overall Adaptive Infrastructure strategy from HP.

Virtualization taxonomy

Storage virtualization has historically been described in terms of **where** it is being implemented: either within storage subsystems (storage-based virtualization), in the storage network fabric (network- or fabric-based virtualization), or in a server (server- or host-based virtualization). At another level, virtualization describes **what** is being virtualized: physical resources (arrays, tape libraries, and so on), data "containers" (LUNs, filesystems, files, and so on), storage access; and **how** the virtualization is accomplished (in or out of the data path). Figure 3 shows a simplified infrastructure stack that is helpful for understanding where virtualization can be implemented and the resources that can be incorporated in a pool.

Figure 3. Virtualization is traditionally implemented by software residing in storage, network components, or servers

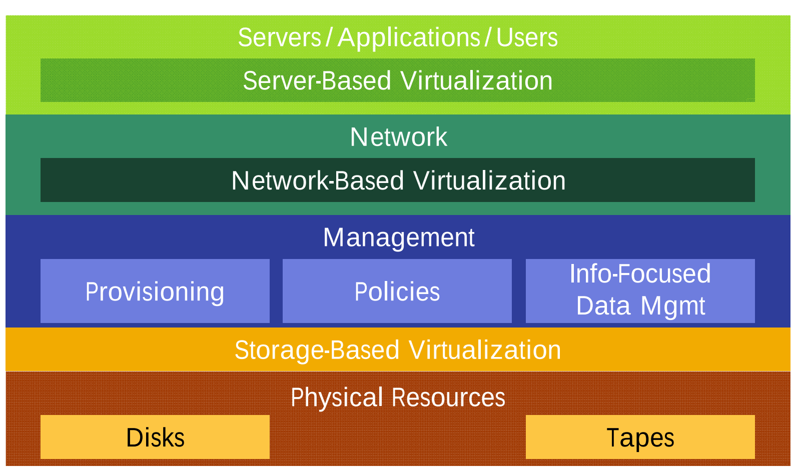

Implementing storage virtualization

Storage-based virtualization is applied to elements within a single storage subsystem or “frame.” In this approach, the virtualization technology (implemented with software or firmware) resides within the subsystem. Virtualized objects (disks, filesystems, tape drives, or other objects) are available to any host that can access the subsystem—through either a direct or network connection. RAID, snapshot, partitioning, and controller failover software are examples of basic storage-based virtualization functionality. Storage vendors continue to evolve storage-based virtualization with advanced resilience, provisioning, data protection, and other capabilities. Storage-based virtualization is generally host-neutral—it can be presented equally well to supported computers and operating systems. It also tends to have low latency (although it might impose high requirements for controller processors and memory), and is provided to applications as securely as non-virtualized storage. Storage-based virtualization techniques and software tend to be unique to each storage system (or family). It should be pointed out that for virtualized storage presentation to be host-neutral, there must be no reliance on OS- or other host-specific capabilities. It is possible to build interesting virtualization capabilities into an array—or other system—that require driver or other host-based features, which can make these virtualization features host-specific.

In addition to unifying disk drives into an array pool, storage-based virtualization can be used to reduce or remove complexities that occur when an application outgrows the disks provisioned to it. This is accomplished by techniques collectively called thin provisioning. Normally when a virtual disk is created, it includes a volume map and all of the physical storage capacity represented by that map. In other words, a 100 GB virtual disk presented to an application can occupy around 100 GB of physical disk. If the application’s needs grow beyond the disk’s capacity, a new disk needs to be created, the original disk’s data needs to be migrated to the new bigger disk, and then the new disk needs to be logically substituted for the old disk—typically an application-disruptive and time-consuming process. With thin provisioning, such as is implemented in the HP StorageWorks XP20000 and XP24000 arrays, virtualization is called into play. With thin provisioning, the virtual disk size does not need to match the physical capacity that supports it—only the immediately required physical capacity needs to be allocated to the disk. Thus a 500 GB virtual disk could be created and presented to an application. Initially, the disk may have only 100 physical gigabytes actually assigned to it. Over time, as the application’s needs grow, the array would automatically add physical capacity in anticipation of higher capacity requirements—and the virtual disk size as viewed by the application would not change. In practice, an administrator would typically set a threshold that would trigger the array to take the required actions. Importantly, the virtual disk’s application-visible attributes do not change as physical capacity is added. This means that thin provisioning enables the easy acquisition and deployment of physical storage as it is needed. Also, changes in physical storage allocations to applications can happen transparently—with no downtime or administrative complexity.

HP offers a capability similar to thin provisioning in the EVA array family. **Dynamic Capacity Management** (DCM) software goes beyond conventional thin provisioning by allowing the EVA to dynamically shrink volumes as well as growing them. Reducing the physical capacity allocated to a virtual disk is useful when application usage changes—for example when a formerly critical application becomes less used, and some of its data can be retired, thus freeing disk capacity. For utilization efficiency, it is desirable to move physical storage from the less used volume to one that requires more actual capacity.

It should be obvious that shrinking a volume could pose the risk of data loss if in-use physical capacity was removed from a virtual disk. Because of this, DCM requires a host agent in order to gain knowledge of how allocated virtual disk capacity is actually being used by the host—this is how safe volume shrinkage can happen. Since DCM requires host agents to work, it is operating system-dependent, and the supported operating systems can always be found at the HP StorageWorks web site.

Network-based storage virtualization is applied to resources presented by storage subsystems attached to a storage network (typically a SAN)—usually to join a number of disk arrays into a single pool from which virtual (logical) disks are created and presented to hosts. These virtual disks may be provided from virtualized storage subsystems like the HP StorageWorks Enterprise Virtual Array. Notice that in this example, different virtualization approaches may be combined to achieve the desired results (availability, performance, and physical location, for example). Virtual disks created within the network can be presented to any host connected to the network, and they are securely presented in a host-neutral manner. Importantly, network-based virtualization can have functionality beyond pooling storage. For example, network switches often have the ability to partition a network into seemingly independent virtual networks or segments. For purposes of this paper, network-based virtualization that is not specifically concerned with direct operations on storage systems is not discussed.

Network-based storage virtualization can be implemented either within switches or with self-contained devices—sometimes called appliances—attached to the network. Switch implementations may be accomplished within the switch operating software, with more interesting switch-resident functionality often being delivered as software that runs on modules, or blades, installed within the switch. Appliance-based implementations are classified as "in-band" or "out-of-band." In-band appliances, like the HP StorageWorks SVS200, connect to the network and serve as intelligent logical bridges between the resources they pool and the hosts that access them—their functionality becomes part of the data path, and in this sense they function similarly to an array controller. Out-of-band appliances work indirectly, with reduced direct data path involvement—they act more as resource aggregation and connection directors. They perform their virtualization tasks, present the resulting volumes to hosts, and remain out of the data path unless needed for particular operations.

When combined with network-based mirroring, and other capabilities, network-based virtualization can provide a useful means for migrating data among storage systems—for example, moving data from a retiring array to its replacement, or moving data among tiers of storage. Network-based virtualization can make these and other data movements almost invisible to running applications. Essentially, the application is paused while the array servicing it is placed into the virtualized pool. Once in the pool, the application resumes access while the network virtualization technology performs the migration (typically a mirroring operation) and swaps logical disk identities when the migration is complete. Accessing applications may see a slight performance impact during some operations, depending on how the mirroring is actually accomplished. Data migrations performed this way are far less disruptive and labor intensive than traditional "backup and replace" array swaps.

Network-based storage virtualization mechanisms tend to differ among vendors. This is important to keep in mind, because like other virtualization approaches, network-based virtualization relies on algorithms (formulas that dictate how the virtualized pool will be created and data deposited within it), and metadata that tells the system what was done (what data is where, for example). This means that data on arrays that were virtualized with one switch or software/appliance engine may not necessarily be directly usable by, or interchangeable with another virtualization mechanism. The importance: network-based virtualization should be viewed either as a very long-term solution (where your chosen virtualization technology may be in place for a long time), or as a solution for short-term projects (like moving data from one array to another). Changing network-based virtualization products may require significant backup, data migration, or other efforts.

Uses for network-based virtualization include:

- One--time data migration among arrays.
- Creating pools of dissimilar arrays for use as integrated tiered storage environments. This can be done in conjunction with array-based functionality, and may include higher-level storage management software.
- Creating expansive capacity pools to accommodate huge data sets. However, this sort of "capacity scaling" is of decreasing importance today, given current storage array capacities in the hundreds of terabyte and multiple petabyte ranges (and growing). Also, care must be taken when creating a pool of this sort, lest the pooling mechanism throttle back the performance of high-end arrays to match them with lower-performing units—virtualizing tends to standardize the behavior of pooled components.
- Continuous data replication among sites for business continuity, backup, or other purposes.

Server-based storage virtualization is implemented by software running in a server. The software may reside in the operating system, or in a driver, or as a layered software application. Server-based virtualization can be applied to any storage visible to the server, regardless of whether it physically resides on direct-attached (DAS) or network-attached devices (SAN, LAN). The flexibility of being able to create a pool from disparate storage and attachment technologies is one of server-based virtualization's strengths. Another strength is the level of integration possible with the operating system, which allows OS-specific optimizations and functionality to be easily and transparently implemented. Interesting applications for server-based virtualization include:

- Mirroring data among multiple storage systems such as external arrays and DAS RAID.
- Creating filesystems that can be provided to networked clients. This is the original implementation of file serving that has evolved today to "NAS."
- Creating RAID sets from embedded disk drives when a RAID controller is not available.
- Creating large volumes from a number of smaller individual disk drives.
- Host path failover among multiple SAN connections (HBAs).
- Performance load balancing among multiple SAN connections.

Storage virtualization tradeoffs

From the preceding discussion, it is clear that there may be many ways to achieve a particular result through virtualization. For example, one could aggregate several different arrays into a single pool using either network-based or server-based virtualization techniques. Likewise, RAID could be implemented just as effectively using server-based RAID software or subsystem-resident software. In choosing a "where to do it" implementation, one should be aware of a number of tradeoffs. These are summarized at a high level in Figure 4. Keep in mind that one can "stack" virtualization implementations, thus combining the benefits of all layers while reducing the drawbacks of each layer.

Figure 4. Storage virtualization summary: all storage virtualization implementations have tradeoffs. This chart summarizes key tradeoffs involved with choosing an implementation location

Attribute	Storage-based Benefits	Storage-based Drawbacks	Network-based Benefits	Network-based Drawbacks	Server-based Benefits	Server-based Drawbacks
Simplified management	Easier provisioning, less administrative effort	Span is controller or single subsystem	Fewer objects to manage (ie, small number of switches vs larger number of subsystems)	Management apps from different switch vendors can cripple each other; generally limited to managing arrays	May use familiar server management tools	Detracts from server's ability to do application work
Pool storage from different vendors/ subsystems	Simplified management for all elements in the storage subsystem	Limited heterogeneous capability (eg. XP24000)	Unifies multiple subsystems; Span can be huge; High performance	Requires additional virtualization appliance, software (complexity); Unpredictable application performance	Unifies multiple subsystems; can pool any storage visible to the server (eg. embedded + SAN)	Reduces server resources available to applications; requires additional software
Non-disruptive data movement	Snapshot, other services are vastly simplified within a storage system	Limited to individual subsystem; Different implementations among systems	Enabled and consistent among multivendor storage and multiple systems	Implementations and management vary among vendors	Enabled among multivendor storage and multiple systems	Reduces server resources available to applications; Server OS-dependent
Improved asset utilization	Capacity and performance distribution uniformly across subsystem elements, thin provisioning, DCM	Limited to individual storage system assets (disks, tapes)	Utilization of a group of subsystems can be improved in unison	Distribution not uniformly achieved among subsystems; adding assets can be cumbersome	Can effect all storage seen by the server	Access by other servers must be through the primary server

Storage-based virtualization derives many of its benefits from the performance advantages inherent in modern array controller architectures (including design integration and low latency), the ability to connect easily to switches and / or host adapters, and neutrality in presenting virtualized capabilities to hosts. Also, locating virtualization and other advanced features within the storage system is often the most technically straightforward, easily managed, and host-friendly approach. Most limitations of storage-based virtualization are related to the span of resources it can control—typically a single physical storage subsystem. The emerging class of scale-out storage systems such as HP EFS Clustered Gateway and Integrated Archive Platform offer the promise of greater scalability coupled with the ability to selectively scale specific storage service attributes (performance, capacity, availability) more independently than is generally possible with conventional disk arrays. When a single array, NAS appliance, or tape library is not large enough, virtualization that can span multiple subsystems may be employed as an alternative to buying a new, larger array. This is the realm of network- and server-based virtualization.

Both network- and server-based virtualization can be applied over a greater scale than storage-based virtualization. With these approaches, multiple subsystems—homogeneous, heterogeneous, and multi-vendor—can be aggregated into a single managed pool. Storage from the pool can be presented uniformly to consumers, with individual subsystem capabilities being almost totally masked to the outside world. Data can be migrated among subsystems within the pool, availability can be improved by mirroring among subsystems, and other useful inter-subsystem operations can be implemented—all totally transparently to consumers.

Network-based virtualization was introduced in the late 1990s to solve two key problems. First, customer data sets were often much larger than the disk subsystems available at the time. Hence, the concept of "beyond-the-box storage"—pooling the capacities of multiple arrays to create a "mega-array"—was very interesting. In addition to being able to store vast data sets on a single managed storage system, one could also think about mirroring within that huge system, and even having it be physically distributed across multiple locations. The second major motivation was data migration: non-disruptively moving data from one array to another—when an array is being replaced or when data must be moved from one location to another, for example. Today's storage systems are significantly larger that they were then, and it is not hard to find frames that can hold hundreds of terabytes, and even scale to more than 200 PB. Hence, there is significantly less need today to spread data sets across multiple frames for capacity expansion reasons alone. Performance, availability, or other considerations still exist, however. This leaves data migration as a key use for network-based storage virtualization, and this remains one of its greatest strengths. Also, switch- or appliance-based virtualization can add latency to the I/O path. The latency added by an appliance depends on whether it is in-band or out-of-band, as well as the design of the software running in the appliance. Switches are always in-band, but their latency is generally negligible.

There are relatively few drawbacks to network-based virtualization, but they are important. Recall that network-based virtualization can be implemented either in a switch or an appliance. In either case, the virtualization mechanism (pooling mechanism, data placement, metadata, and so on) is unique to the implementation product. Since multiple storage systems are being aggregated, one must be cognizant of the implications of changing virtualization technologies or mechanisms when switches or appliances are changed (primary concern), and of compatibility of the underlying subsystems with the physical and virtualized infrastructure (secondary concern). Put simply: when you replace a switch and its virtualization technology with another switch or appliance and its virtualization technology, what happens to the data in the pool, host presentation, and so on? Whereas virtualization uniqueness within arrays is self-contained and usually innocuous in this regard, the situation changes dramatically when a higher-order form of virtualization is imposed.

There is also a potential side effect of network- or appliance-based virtualization. Some virtualization devices include their own remote replication and other software. The good news is that these capabilities act on all members of the storage pool. The bad news is that often, these superimposed capabilities supersede—or totally replace—similar capabilities that may have been built into the individual storage arrays. For example, EVA Business Copy or Continuous Access software may no longer work or may no longer be utilized in a network virtualization solution. Instead, less specific or finely tuned, higher-level mechanisms could be used. This has two implications. First, useful and valuable features that have already been paid for are no longer useful. Second, the network virtualization solution's technical capabilities need to be assessed in the context of the job that needs to be done, and then a decision must be made as to the practical results of selecting network-based versus storage-based virtualization technologies.

Server-based storage virtualization capabilities can be integrated tightly with operating systems and are able to span any storage that is visible to the server. Integration opens the possibility of building advanced capabilities into the virtualization software. This could be taken to the point of integrating storage virtualization with server virtualization, which could make it much easier to synchronize and orchestrate the automated provisioning of these resources—making it easier to manage demand-based provisioning and overall resource utilization. Also, storage virtualized by the server can be managed as part of the server, as opposed to requiring separate management interfaces, which can simplify overall management. As another example, storage virtualization capabilities could be linked in meaningful ways to the ways operating systems manage data placement and storage device utilization—EVA Dynamic Capacity Management is an example.

Server-based virtualization is often useful when the goal is to manage storage in the context of a particular server (similar to dedicated storage) as opposed to a shared environment (typical of network storage). The key compromises that server-based storage virtualization makes include:

- If the storage needs to be accessed by another server or client, access must be through the primary server (the one that is doing the virtualization). This introduces latency, places additional workload on the server, may require additional network connections (NICs), and adds the server as a point-of-failure in the data path.
- Server resources are consumed to accomplish the virtualization. This can reduce the resources available for accommodating production application workloads.

As you can see, regardless of where storage virtualization is implemented, it can offer significant benefits. It is up to the storage or system architect to balance product features with their tradeoffs, in the context of the challenges that needed to be solved.

A more complete virtualization perspective

The discussion up to this point has been based on a simple, fairly traditional, infrastructure-centric view of storage. While general concepts are helpful in understanding what can be done by implementing storage virtualization in various places and in various ways, it is increasingly useful to understand virtualization from a business solution perspective. In other words, we can **think of virtualization as a vast collection of enabling technologies that are combined into products and** solutions that do useful things for IT. From this perspective, the focus needs to shift to thinking about how applications and IT administrators interact with the resources they consume and manage, and how these interactions can be made more efficient. Efficiency is measured in terms of improvements in productivity, efficiency, and agility to respond to changing conditions. Indeed, this is how StorageWorks views virtualization: it is a growing set of technologies that form foundations and provide mechanisms for delivering customer value, embodied in products and solutions offered by HP and its partners. Now focus on virtualization from a broader, application-centric viewpoint.

Production applications, and the business processes they serve, consume resources. These resources can be described in terms of capacity (such as compute cycles, gigabytes storage, network bandwidth), performance (storage IOPs, bandwidth, network and storage latency, for example), availability (uptime, resilience to failure), security, and so on. Based on their ability to access the required resources when needed, business results are delivered.

These resources are simply seen and consumed: their underlying physical components and the management path that presents them to applications are invisible to the applications themselves. Thus, the entire path over which information flows and is processed is really a concerted progression of virtualization steps layered upon core tangible resources. Put another way, resources can be delivered from virtualized physical pools by conduits in the form “virtualized services.” Virtualized services are managed pathways that provision (allocate, present, and meter), aggregate, or in other ways warrant that appropriate resources are made available to applications in the most efficient ways. They provide end-to-end, application-relevant linkages between physical resources and the processes that consume them. This is a key to understanding storage virtualization in the greater context of IT and the HP Adaptive Infrastructure.

A different perspective

From an IT perspective, one can view virtualization as having two dimensions: business value and the strategic importance of implementations. Technically, these correspond to the range of resources being pooled, and the ongoing impact to business results. This is useful because it helps to rationalize technology with the benefits it delivers.

Figure 5 shows a comprehensive view of IT virtualization from HP. At the most basic level, physical resources are pooled. This is called “Element Virtualization,” which is where “inside-the-box” storage-based virtualization fits. EVA, MSA, and other subsystem-inclusive products deliver element virtualization.

Figure 5. Virtualization value chain

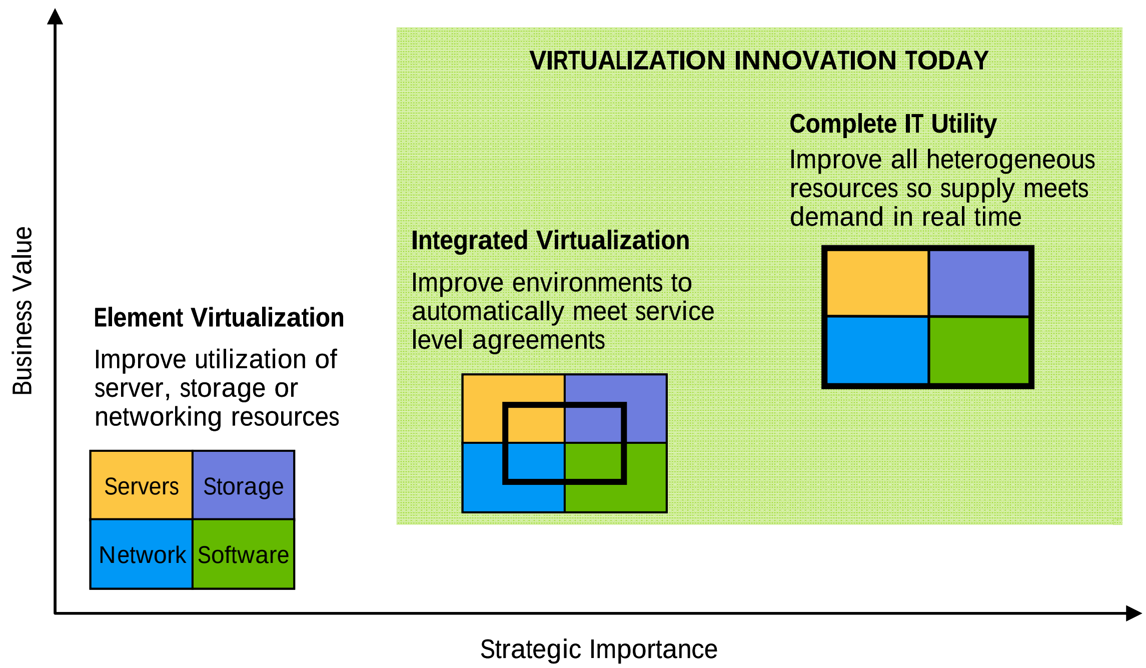

At a higher and more interesting level, heterogeneous resources can be unified. Pools of this sort encompass more diverse capabilities, with correspondingly greater business value, than more constrained element pools. Integrated virtualization creates pools that span multiple subsystems. These pools may combine diverse type of resources such as servers and storage, and different types of storage. The physical resources may even be geographically dispersed. Integrated virtualization necessarily involves resource management software that can provide capabilities beyond basic pooling. Data migration and remote mirroring are two examples.

Moving higher, additional automation, integration, and business application-aware software could be added to provide advanced, IT-wide capabilities. The Complete IT Utility delivers virtualization across the entire IT environment. In addition to integrating a broad and disparate range of physical resources, the Complete IT Utility uses virtualization technologies to provide automatic, dynamic, service-oriented resource delivery. More detail can be found in the references listed at the end of this paper.

Notice that each successive class of enterprise virtualization provides increasingly greater value to the enterprise. This is why, from an industry R&D perspective, there is more focus and innovation occurring at the higher levels. Indeed, HP StorageWorks is devoting more effort to technologies that deliver integrated virtualization, and is working with the rest of HP to deliver the Adaptive Infrastructure, which can provide Complete IT Utility virtualization.

Each successive level of management that is involved in providing resources to applications could be thought of as an element that provides an increased level of virtualization (that is, broader virtualization of the underlying resources).

Each management component ("value-add service") that contributes toward ensuring delivery of data/content/information to applications can be thought of as contributing toward virtualization.

This perspective implies that:

- Automated provisioning can be thought of as automatically providing virtualized physical resources to applications. Thus, one could use provisioning software to present and manage the continued presentation of storage array resources to applications.
- QoS can be thought of as "virtualized resources that deliver [data/content/information] to business applications with assured service levels," and this is accomplished by layering provisioning management software that takes, for example, virtualized physical resources and adds QoS attributes (such as through the addition of metadata), and subsequently provisions the results to applications.
- Automated failover, business continuity, and other high-availability storage management software can be thought of as adding a "virtualizing layer" whose job is to mask specific kinds of storage failure from applications. In other words, these types of software are essentially providing a value-add layer of virtualization (virtualizing multiple arrays, perhaps across multiple sites) that results in applications not noticing when a specific physical resource dies.

HP StorageWorks virtualization today

Figure 6 shows a high-level mapping of virtualization delivered by the StorageWorks portfolio.

The figure focuses on the more advanced virtualization functionality offered in StorageWorks

systems themselves. HP Storage Essentials SRM software (which virtualizes the management of multiple systems), HP File Archiving software (which offers HSM functionality), and other HP storage software products are not shown on the chart. Also, while MSA offers advanced RAID functionality, HP realizes that the focus of today's virtualization efforts is on deeper levels that encapsulate more resources and capabilities, and provide greater benefits than older, more conventional approaches.

Grid-enabled virtualization products—types of scale-out storage systems—represent integrated virtualization. They offer a hybridization of the capabilities found in both storage- and network-based virtualization. For example:

- The HP StorageWorks EFS Clustered Gateway delivers file system virtualization from a scale-out, multi-node system that accesses SAN-based storage capacity. It includes tools that simplify provisioning, while improving client access to data.
- The HP StorageWorks 9000 Virtual Library System (VLS9000) and other virtual library systems ("virtual tape libraries") virtualize the physical resources (control nodes and the storage resources connected to them) into a pool. They then go farther, adding virtualized tape drive and tape library presentation to backup and other applications.

At the network level, intelligent switches provide virtual, changeable connections between servers and storage. Beyond this "virtualized hose" concept, intelligent switches can also run special applications that enable them to provide advanced management capabilities. For example, an application running within an intelligent switch could create a virtual pool of storage from a collection of arrays attached to the switch. In this example, virtualization might incorporate the ability to mirror LUNs across dissimilar arrays, or migrate data among arrays.

Figure 6. Virtualization is embedded in all of today's HP StorageWorks subsystems

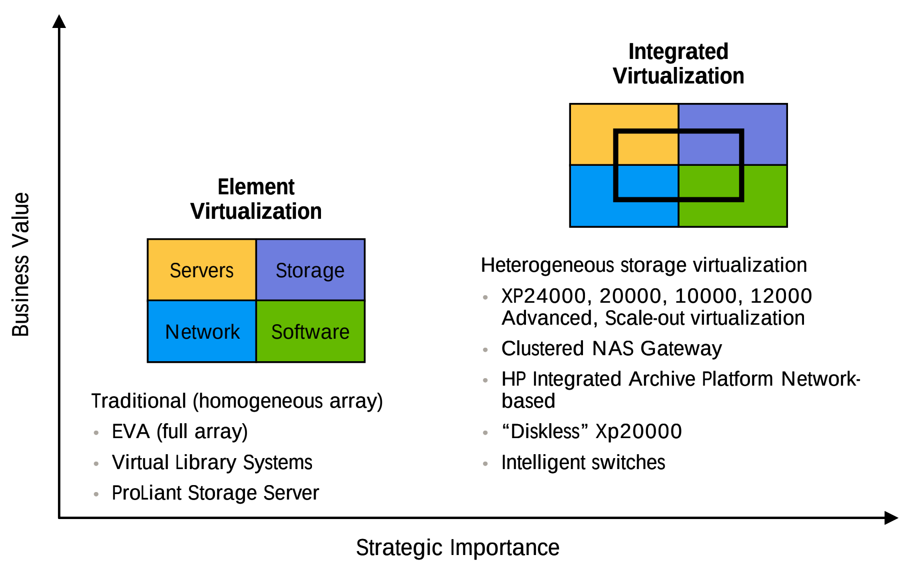

HP StorageWorks use cases

Nearly all HP StorageWorks hardware and software products incorporate and deliver virtualization today. Figure 7 shows some examples that are in wide use. This section describes them in high-level detail. For more information, refer to the QuickSpecs at http://h18006.www1.hp.com/storage/index.html.

Figure 7. HP StorageWorks storage virtualization is widely deployed today

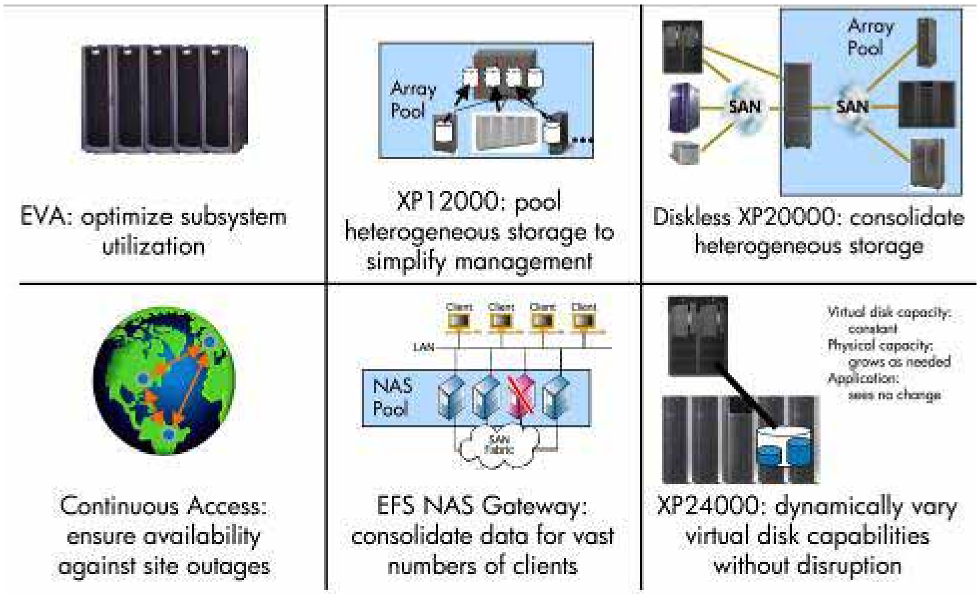

The EVA family features system-wide virtualization of all disk resources managed by the EVA controllers. Whereas other systems group disk drives into discrete physical RAID groups and thus have multiple physically distinct RAIDs, EVA uniquely binds RAID across groups of blocks and distributes all resulting RAID groups across all disk drives. EVA RAID uses the same algorithmic data placement patterns as traditional RAID, but applies them more intelligently with the addition of even more abstraction, so that a single disk drive can participate in multiple RAID groups. The result: automatic workload (capacity and performance) distribution to all disks in the array, while delivering one of the easiest arrays to configure in the industry. A secondary result is that data availability in an EVA can be higher than that of the equivalent conventional RAID implementation. The reason is simple. Consider RAID 5. By definition, RAID 5 has redundancy that protects data against an individual disk loss. However, after a disk loss, all data on the RAID is at risk until the failed member is replaced and its data is regenerated from the other disks—if **any other disk** in the array fails before the RAID has been regenerated, all data on the set is lost. However, because of EVA's virtualization mechanics, it turns out that an EVA RAID 5 (called "vRAID 5") can sustain up to **two** disk failures, **depending on which two disks fail**. HP does not recommend that you depend on always being protected against two disk failures, but it is reassuring to know that virtualization makes this possible.

EVA virtualization also enables dynamic expansion of the disk pool. Even after the disks have been fully configured into RAID and presented to hosts, when disks are added to the EVA frame, they are automatically incorporated into the pool. The EVA goes further: after the new disks have been incorporated into the pool, the workload is automatically redistributed across all disks—in the background and almost totally invisibly to host applications. The ability to easily expand as needed—and automatically tune the results—makes it much easier for customers to acquire storage capacity as needed, as well as making for much more efficient use of storage.

Today, while some storage vendors are marketing their network-based virtualization approaches, HP StorageWorks offers similar capabilities in two well-integrated, proven storage systems. The XP10000 and XP12000 Disk Arrays can virtualize a number of storage systems—from HP and other vendors—into a single pool. When you connect any of a variety of arrays—from HP, EMC, HDS, or IBM—to an XP12000 Disk Array, they are combined into a single high-performance (up to 1.9 **million IOPS; 68 GB/sec**), high-capacity (up to 32 **PB**) virtual pool (the XP10000 Disk Array works similarly, but on a somewhat smaller scale). The administrator then manages one XP12000 Disk Array instead of a collection of individual arrays, and business applications receive their data by accessing the XP array. Thus, HP StorageWorks XP family virtualization has attributes of both storage- and network-based virtualization. This can be extremely useful for migrating data from one array to another: just pool the arrays behind an XP array, have the XP mirror (copy) the data from the old array to the new one, and the XP can do the rest—including "retiring" the old array when data transfer is complete. You can use variants of this approach to create, for example, a single pool of storage that contains multiple tiers—a very straightforward way to deploy tiered storage without the overhead of multiple management interfaces and other complexities. Interestingly, arrays that have been part of an XP pool can later be removed from the pool and used conventionally—and the data remains intact and usable. This contrasts with typical network-based virtualization approaches that deposit data on constituent arrays using unique and proprietary schemes, rendering the data on individual elements unusable directly if the element is removed from the pool.

Like the EVA, XP arrays incorporate a number of other virtualization technologies. For example, the huge cache in an XP array can be partitioned, and part of it used as a solid-state disk. Yet another example of useful virtualization is the ability to partition the XP array into multiple independent virtual arrays—a capability not found on competitive products.

The need to migrate data among storage systems—transparently to business applications—is a key reason for virtualizing multiple systems into a single pool. Traditionally, when a storage system was about to be retired or needed to be replaced by a larger system, data was manually moved from the old system to the new. This was usually a tedious, time-consuming, and application-disruptive process that required significant planning to lower business disruptions. While the XP array family can accomplish this, sometimes it can be lower cost to have a network-based approach. HP created the SVS200 to solve this and similar problems. (Editor's note: The SVS200 has been replaced by a "diskless" XP20000 that provides the same capabilities as the SVS200 did). Employing key technologies from the XP family, the SVS200 is a network-based storage virtualization appliance. SVS200 offers capabilities not found in some competing network-based products. For example, the SVS200 can mirror and migrate data among heterogeneous arrays **both locally and remotely**. SVS200 can use HP StorageWorks Continuous Access capabilities to mirror data from an EVA at one site to an EVA or even another vendor's array at a remote site. Taking this example a step farther, the remote system could then be used as a backup source to a co-located library—effective, "zero backup window," remote, lights-out backup. This is a very powerful and useful application of HP virtualization technology.

This leads into to another important set of virtualization solutions from HP StorageWorks. All IT organizations strive to keep the company's data available whenever and wherever it is needed. Virtualization from StorageWorks helps in a variety of ways. Storage-based solutions protect against loss of access due to disk failure (RAID) and controller failures (redundancy and automatic failover). At a higher level, SVS200 and XP inter-array mirroring protect against loss from entire subsystem failures. And while you do not often think of SANs as virtualized environments, in fact that is what they are—they virtualize the connections between servers and storage, helping to keep data flowing even if a switch or network connection (HBA) fails. But storage virtualization can be extended to verify that data is available even if a site goes down. HP StorageWorks delivers this with Continuous Access and HP StorageWorks Business Copy software for EVA and XP subsystems. These products create copies of data at one or more remote sites, either synchronously or asynchronously. Coupled with the extended clustering and hot-failover capabilities found in some operating systems and applications, HP virtualization solutions synergize to enable the availability of your production applications and their data if there is a problem at a site.

The foregoing use cases illustrated interesting and useful ways that HP applies traditional storage virtualization technologies to solve real-world problems. However, StorageWorks continues to evolve virtualization technology and its uses. Its long-term strategy includes a new, totally virtualized storage environment that delivers storage capabilities as a utility service. Fundamentally, this seeks to create a unified infrastructure upon which storage, server, and advanced management software virtualization technologies may be married to form a new, intelligent storage ecosystem. The result may be a service-oriented, scaled-out entity that can deliver the functionality of a number of today's disparate storage systems—arrays, NAS, and tape libraries, for example. These can ultimately be provisioned to respond to business application demands. StorageWorks has already begun the journey with the release of a number of **grid-enabled storage systems**, and evolving the management functionality of Storage Essentials. In addition, combining storage and server resource management through Systems Insight Manager provides the beginnings of unified management of very disparate resource pools—a step toward an IT utility. Going forward, these management capabilities may continue to evolve and ultimately deliver a service-oriented resource utility—layered over the types of infrastructures that are common in today's IT environments.

Grid-enabled storage systems warrant special attention, in part because they may portend an interesting direction for storage in a more general sense. As some industry observers have noted, servers and storage seem to have an opportunity for some convergence in the future—the increasing appeal of server blades in one catalyst. Among other things, this makes it possible to deploy useful and interesting new functionality into the storage domain. Following are two illustrative examples.

The HP StorageWorks EFS Clustered Gateway is an example of a grid-enabled product. It is composed of a number of processing nodes that run a storage application, storage capacity obtained from arrays in a SAN (remember: HP pioneered NAS/SAN convergence years ago—a useful application of virtualization), and a storage application that runs on each node. In this case the application provides a clustered filesystem that can be accessed by networked clients connecting to any node, and whose data is collectively stored in a SAN. Here, in addition to utilizing virtualized storage for capacity, the system virtualizes multiple access and file system delivery and processing points among the nodes. The results:

- Very high performance, measured in CIFS and NFS.
- Easy, reliable access: the aggregate of nodes, storage, applications, and filesystems are virtualized into a single unit from a client-access perspective.
- Simplified administration: the EFS Clustered Gateway presents a single system for administration tasks.

- Very high, **predictable** reliability: as long as any node and storage are available, users can access their files. The more nodes in the system, the greater the reliability. Also, if a node fails, the performance decreases fairly constantly at 1 divided by the number of nodes. SAN reliability is accomplished as for any other SAN installation.

Thin provisioning is one of HP's latest storage virtualization offerings. Available as an option for the XP24000 and XP20000 arrays, thin provisioning allows for higher disk utilization and pay-as-you-grow expansion. As described earlier, thin provisioning provides interesting capabilities to virtual disks created and managed by the XP24000 or XP20000. Virtual disks are created as usual during the initial configuration process. Upon creation, though, the administrator has the option to specify the fraction of virtual disk capacity that is backed by physical storage capacity—unlike traditional virtualization that requires physical disk space be equal to the virtual volume presented to applications. At the same time, a threshold is set so that additional physical capacity can be automatically added to the virtual disk if capacity utilization exceeds the predetermined value. The host sees, trusts, and works with the virtual disk capacity presented to it. The XP Disk Array monitors the volume's capacity utilization and, transparently to the accessing host, assigns additional physical capacity to the volume as needed. This saves administrative headaches by requiring less array management and manual disk usage monitoring. It also warrants that applications can run nearly continuously without worrying about running out of capacity (i.e, they are presented with larger storage volumes than they may ever need). And, since only the amount of actual physical capacity needed at a given time is provisioned (consumed), thin provisioning helps keep the total cost of storage as low as possible. Lastly, in these days of energy consumption sensitivity, thin provisioning is truly a "green" storage feature.

This summarizes some of the interesting ways that StorageWorks delivers the benefits of virtualization in interesting and unique ways. Virtualization technologies power all StorageWorks systems—tape libraries, virtual library systems, NAS appliances, and more. The focus of HP is to solve customers' business problems, leveraging appropriate technologies and extending boundaries wherever possible. Many StorageWorks capabilities are based on practical applications of virtualization technologies.

Storage virtualization—more than just hardware

Administrators need tools that aggregate complex environments into simpler, more understandable entities. They also need to be able to do their part to ensure that data is available only to authorized users, and that it can be easily accessed once it is made available. Also, given financial and regulatory realities, it would be useful if data could automatically be stored on the most cost-appropriate assets, and that the storage environment is sufficiently secure to satisfy laws and company auditors.

Delivering intelligent hardware and infrastructure virtualization is a good start, but HP offers much more complete capabilities. To address these requirements and more, HP offers a broad and evolving range of storage resource management, information lifecycle management, and data protection software. Figure 8 offers a representative view of the breadth of the storage management capabilities from HP. To the left, Storage Essentials provides resource management for all SAN-attached storage assets. To the right, HP Information Management solutions manage and use the SAN and other storage hardware for data- and application-focused purposes. This type of functionality includes intelligent data archiving—transparently managing data placement in the storage pool, with the goal to improve placement and retrieval for business application needs. HP solutions can also migrate data among tiers of storage and protect data according to application-relevant criteria.

Figure 8. StorageWorks software: virtualization applied to hardware and data management

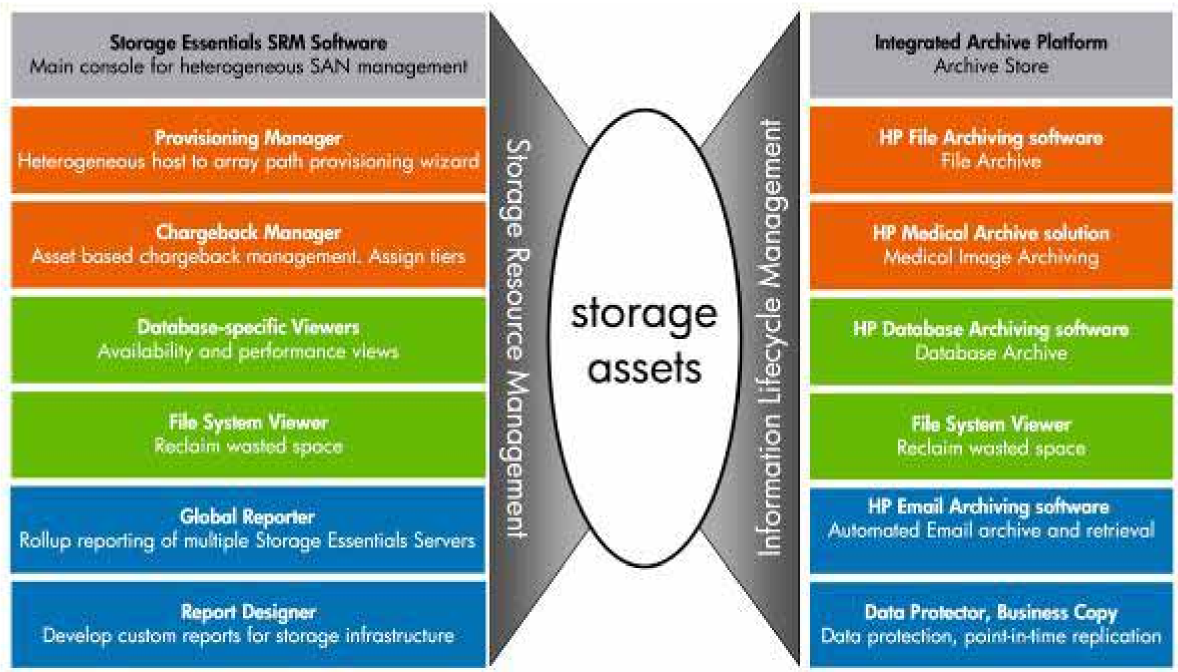

It is useful to understand how these relate to virtualization. Storage Essentials Storage Resource Management (SRM) software is a single point of contact for provisioning and monitoring the storage environment, presenting a unified portal to network storage devices. The console can be customized to enable the administrators to view and access only the resources they are allowed to view: they see a filtered abstraction of the entire environment. While Storage Essentials delivers little virtualization functionality today, the situation may change over time as provisioning tasks become increasingly automated. As its automated provisioning, monitoring, and other management functions reach more deeply into the physical hardware, Storage Essentials can provide managers with a highly virtualized interface that encompasses many features in the SAN's assets—administrators can increasingly deal with provisioning abstractions (called "policies") and work with virtualized resources that can be allocated as storage services as business applications see them. In other words, Storage Essentials is evolving along a path of allowing administrators to manage storage assets (hardware and software) in the context that applications consume them—virtualized management to deliver virtualized resources.

The storage-related aspects of Information Lifecycle Management (ILM) are fundamentally application-focused virtualization. These products and solutions incorporate application- or data-type-specific intelligence and meld them with user-defined policies to govern the placement of data throughout the data's stored life. A generalized schema includes:

- Administrative entry of policies or other criteria that describe data types, users, and lifecycle access requirements into a management application. These attributes may include availability, performance, protection level, security factors, and so on.
- A pool of storage devices that constitute a centralized data repository. This may include several storage tiers unified through a SAN or other network. The devices may include disk arrays and tape libraries. They may also include systems with embedded advanced intelligence, like the HP Integrated Archive Platform. In addition, data protection software (backup, replication, and related functionality) may be added.

- Management applications' control of data placement within the managed pool of devices. Since access patterns, authorized users, and other factors may change throughout the lifecycle, ILM includes the ability to transparently migrate data among devices in the pool, and to dynamically change storage attributes associated with the data, as dictated by the ILM policies.
- Business application- or data-type front-ends that provide the system with application contexts by which information can be properly identified and stored. Thus application integration software provides an intelligent lexicon between business applications and ILM policies, enabling appropriate classification and placement of data. Thus, a database front-end might identify database records; a file front-ed might identify Microsoft® Office documents; a graphics front-end might identify JPEG and TIFF images; a medical front-end might identify DICOM images, and so on.
- Optionally, incoming data may also be transformed so as to offer usability into the distant future. For example, files bearing proprietary formats may be converted to more standard formats like ASCII or JPEG, as they are transmitted to the storage repository.

Clearly the entire process is geared toward totally virtualizing data placement within a virtualized storage pool, with the added intention of moving data through the pool according to descriptive, user-set policies. Intelligent front-ends aid the storage abstraction process and provide a basis for properly tracking (and perhaps rapidly retrieving) data regardless of where it is physically stored. Thus, ILM represents a very complex multi-layered hardware and software virtualization solution.

So far, the ILM discussion has focused on storage. For completeness, it should be noted that ILM solutions must also incorporate ways to input data to the ILM system. These could be direct input by way of keyboard, camera, x-ray machine, or other device. It could also be through a digitizing device like a scanner coupled with appropriate conversion software. ILM also necessarily includes human aspects to determine lifecycle storage requirements and other elements.

It should now be clear that virtualization plays fundamental roles in delivering most HP storage functionality. As virtualization has become more deeply embedded throughout the product line, HP has focused product descriptions on the business benefits delivered by its solutions, rather than on virtualization foundations and technologies. With this background, it is now time to revisit the virtualization model presented earlier.

HP storage virtualization: The big picture

Today, HP clearly has storage subsystems that offer a broad range of storage-based virtualization capabilities. It delivers network-based virtualization as well. At the server level, host-based mirroring and other capabilities exist. From the foregoing discussion, though, it is clear that HP has a more expansive view of virtualization, and it believes that this can benefit customers by providing more useful solutions into the future. To understand this, the HP view of the "virtualization stack," which is really much more comprehensive than what was described in Figure 3, must be introduced. StorageWorks is rationally extending storage virtualization in powerful ways. The HP perspective enriches the virtualization stack to encompass dynamic provisioning, QoS-oriented functionality (as might be needed by ILM and other business-derived needs), and so on, consistent with the discussion in the previous section. This affords HP a more complete approach and opportunity to deliver complete, integrated solutions than most storage vendors are able to realize, and it forms the basis for the HP virtualization strategy.

HP StorageWorks Extensible Virtualization Stack

First, remember that virtualization is the abstraction of storage that separates the host view from the storage system implementation (source: SNIA storage virtualization tutorial). It makes physical paths, device characteristics, physical data location, and other underlying aspects invisible to entities that access and manage it. Furthermore, it is dynamic.

HP, the Storage Networking Industry Association (SNIA), and others recognize that virtualization involves more than directly virtualizing storage elements. It also includes significant device management, presentation, and application integration elements. With this in mind, Figure 9 presents the HP StorageWorks Extensible Virtualization Stack—an inclusive and flexible model that integrates virtualization across all storage and management elements in the environment.

Figure 9. HP StorageWorks Extensible Virtualization Stack. In addition to traditional hardware-focused storage virtualization, the StorageWorks Grid strategy can deliver an interesting environment that encompasses the capabilities of both storage- and network-based virtualization, as well as embedding advanced management virtualization features

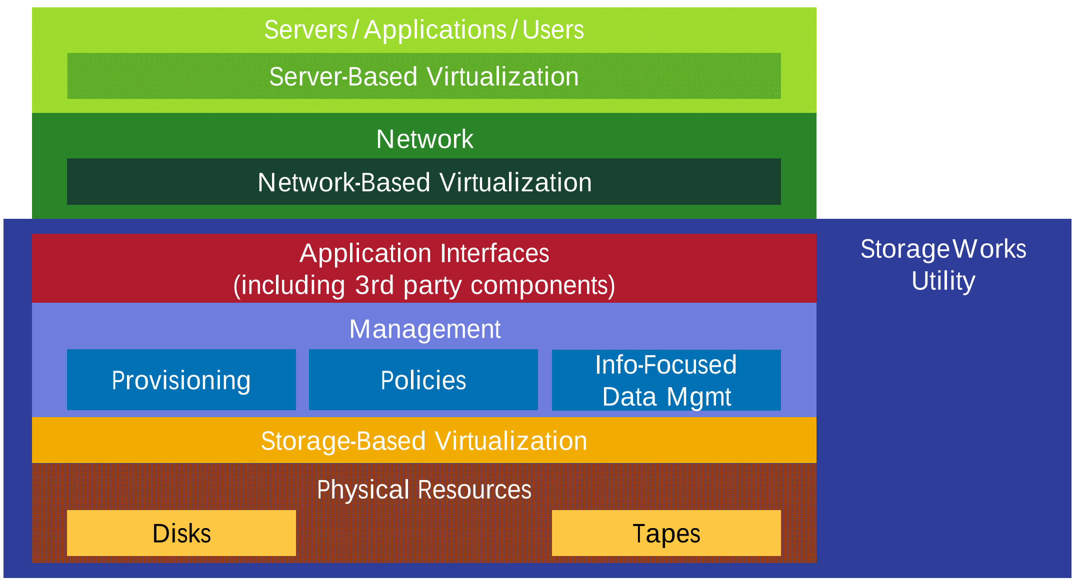

The stack recognizes that many hardware and software elements can contain virtualization technology and / or deliver virtualization capabilities. Furthermore, these elements can be combined in logical ways, thought of as layers contributing to efficient data delivery and asset utilization, that contribute to and amplify the power of a total virtualization solution. For example, network-based virtualization is complementary to storage-based virtualization—each has unique applications, and there are times when it is appropriate to use them together. Note that the layers are logical, and are not necessarily in the data path. This is consistent with the perspective in the previous section.

Being extensible, the model allows for future technology developments. For example, while specific categories of management are shown, the model allows for additional categories such as security, model-based automation, and other types of management can be added when appropriate. Indeed, this is part of the StorageWorks strategy. Also, new storage- or network-based categories can be added. The stack provides a useful way to visualize how different virtualization technologies might be complementary or redundant, for example.

Unified storage—delivered as a service

Importantly, this model includes a new element: the HP StorageWorks Utility. This is our strategic vision, intended to deliver an entirely new type of highly scalable, automated, integrated storage system. The effort embraces several concepts:

- Unified infrastructure. Our goal is to significantly simplify storage system decisions by creating a broadly applicable foundational infrastructure that can grow and be replicated as needed. This is expected to significantly reduce the numbers and types of disparate subsystems needed to provide complete storage solutions in many environments.
- Scale-out architecture. Just as our EFS Clustered NAS Gateway products can scale vastly and easily as needed, our goal is to broadly enable storage systems to physically scale over a very broad range—again, to reduce the number of architectural elements needed to satisfy dynamically changing and growing environments.
- Extensible storage service offerings. Historically, storage has been provisioned in granular, infrastructure-centric ways ("bottoms-up"), rather than in terms of the Quality of Service or other application-relevant perspectives (responsive, "tops-down"). It is increasingly becoming important to efficiently provision storage to support application Service Level Objectives. To do this, it would be very useful to manage storage according to the services it delivers—capacity, presentation, performance, availability, recoverability and so on—rather than simply as virtual disks. And, since application service requirements may change—and new storage technologies become available—our goal is to create virtualized storage systems whose service-oriented capabilities can be easily changed. Thus, our vision is to create storage systems that can efficiently and flexibly deliver many application-relevant services, as opposed to requiring a number of different types of storage systems to do the same job.
- Integrated at many levels. We view network-based virtualization as a strong complement to storage-based virtualization. This leads our vision to another major level of virtualization beyond subsystems. The StorageWorks Utility might ultimately be composed of networked storage assets and pervasive management software. Like today's networked storage, the utility may be able to scale by adding whatever is needed to the network. Software may be deployed more efficiently—within storage frames switches, and even on special management nodes. All management functionality may be orchestrated through common software interfaces, many of which might be based on industry standards. Software may provide intelligent, managed storage capabilities to applications. The resulting ecosystem may be an effective virtualized hybrid of traditional storage systems and a storage network. At the management layer, it may go beyond traditional "manage from outside the box" capabilities by embedding advanced capabilities by way of storage applications. Embedding provides better architectural integration between hardware and software. Among other things, this allows more interesting uses of virtualization.

Figure 10. The StorageWorks Utility is built upon a networked infrastructure; embedded storage and management applications deliver application-focused functionality. Virtualization permeates this unique storage system

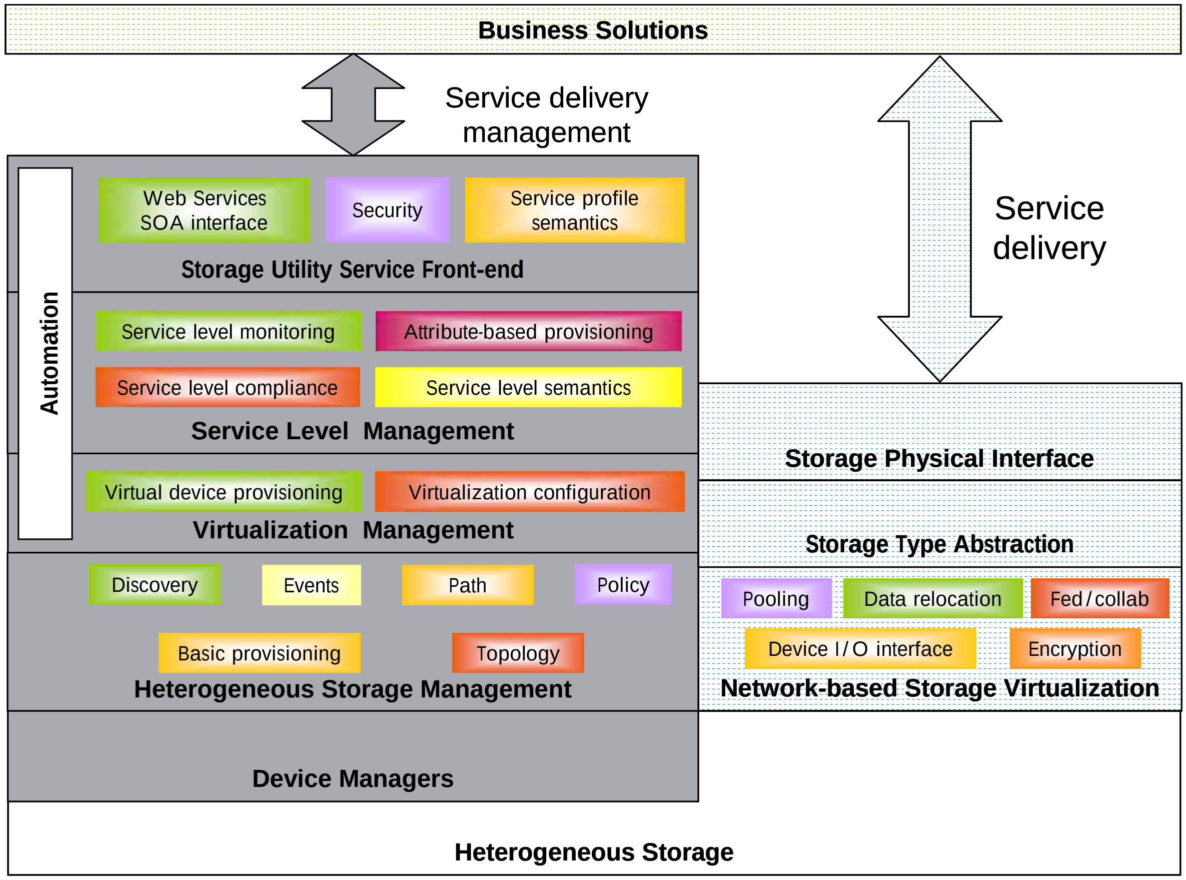

Managing storage to present its capabilities as services is illustrated in Figure 10. At the foundation is heterogeneous storage: disk, tape, and other systems from various vendors. The left side, shaded in gray, represents the management path. Notice that multiple layers of virtualization are combined in useful ways. In addition to storage-based virtualization that resides in arrays, the management stack treats the virtualized systems as a pool from which specific application-relevant capabilities could be carved. This is analogous to network-based virtualization in that it aggregates among network storage systems. Moving higher in the management stack, storage capabilities are translated into application-relevant services, which in turn are allocated to applications as needed. Automation will occur among the higher-level, more advanced virtualized layers.

The right side, shaded in blue, illustrates the data delivery path, which is deliberately separated from the management path so as to allow more flexibility in managing storage non-disruptively to applications. The result is that applications may continue to communicate with storage along the same sorts of data paths they do today (that is, no application modification is needed), and the data presentation is managed to enhance asset utilization while providing the service levels required by the business. Notice that:

- The range of capabilities incorporated in a StorageWorks Utility can be expanded as needed, by incorporating relevant storage applications into new, extensible subsystems.
- The utility can ultimately be managed to automatically provide storage capabilities to meet service level objectives.
- Virtualization functionality spans the traditional domains of storage resource management and Information Lifecycle Management—unifying them in useful ways. This can represent a significant management simplification.

These capabilities illustrate some of the key evolutionary advances that delivering storage as utility services may introduce. Notice how virtualization permeates the utility:

- Storage functionality is delivered by software and is scalable and changeable by adjusting the software or adding new applications. Regardless of the software content, applications that access the system see only the storage services they need and have no visibility to its underlying infrastructure or management mechanics.
- At a higher, longer-term view, networked infrastructure—with associations among components being controllable by software applications—behaves as a logical entity even though its physical underpinnings may change.
- The management environment presents a single entity to administrators, who may not need to focus their attention on specific physical or software components. Instead, administrators will interact with management applications that can control lower level provisioning and other functionality.

This is an important element of the HP storage virtualization strategy. As mentioned very early in this paper, constraining span of control to a single physical subsystem is an important limitation of storage-based virtualization. New scale-out subsystems largely escape traditional capacity and functionality limitations. At a larger scale, the StorageWorks Utility escapes this limitation through the management of its scale-out networked infrastructure. The level of management integration planned for the utility can allow it to deliver virtualization as seamlessly as today's storage systems do. In short, the StorageWorks Utility can combine the benefits of both storage- and network-based virtualization. This is the most comprehensive and highly virtualized storage environment that HP can create, and it draws broadly on technologies developed by HP, its partners, and even other vendors (through their inclusion in the grid). This is what the large blue box in Figure 9 illustrates.

Figure 11 summarizes the benefits the StorageWorks Utility brings to virtualization. The yellow regions highlight specific storage- and network-based drawbacks that are solved by this approach. The last columns present examples of business and administrative benefits the grid is being designed to deliver.

Figure 11. The StorageWorks Utility combines the benefits of storage- and network-based virtualization, while solving many of the limitations of each

	Storage-based		Network-based		HP Storage Utility	
Attribute	Benefits	Drawbacks	Benefits	Drawbacks	Benefits	Drawbacks
Simplified management	Easier provisioning, less administrative effort	Span is controller or single subsystem	Fewer objects to manage (ie, small number of switches vs larger number of subsystems)	Management apps from different switch vendors can cripple each other; generally limited to managing arrays	Unifies multiple tiers of storage; Span is the entire storage utility Encompasses all tiers of storage, plus management	Some configuration limitations
Pool storage from different vendors/ subsystems	Simplified management for all elements in the storage subsystem	Limited heterogeneous capability (eg. XP24000)	Unifies multiple subsystems; Span can be huge; High performance	Requires additional virtualization appliance (complexity); Unpredictable application performance	Convergence: allows traditional arrays to be incorporated into the storage utility virtualized environment/pool	Limited heterogeneous capability (supports 'qualified' HP and third party storage systems)
Nondisruptive data movement	Snapshot, other services are vastly simplified within a storage system	Limited to individual subsystem; Different implementations among systems	Enabled and consistent among multivendor storage and multiple systems	Implementations and management vary among vendors	Enabled consistently and internally throughout the ecosystem, without special considerations	
Improved asset utilization	Capacity and performance distribution uniformly across subsystem elements, thin provisioning, DCM	Limited to individual storage system assets (disks, tapes)	Utilization of a group of subsystems can be improved in unison	Distribution not uniformly achieved among subsystems; adding assets can be cumbersome	Storage-wide optimizations for capacity utilization, performance delivery, billing	

Factors addressed by the StorageWorks Utility

HP StorageWorks virtualization strategy overview

Looking forward, the key elements of the HP storage virtualization strategy are:

- Continue to focus on storage-based virtualization. This includes embedded capabilities in arrays, tape libraries, and other storage systems. This is a key focus for HP product teams. HP has a heritage of delivering excellent storage-based virtualization implementations and it believes that this is the most effective way to continue to deliver the benefits of storage virtualization to its customers. In the future, HP intends to produce scale-out storage systems with broader, more extensible capabilities. At the same time, HP understands the importance of complementary approaches.
- Introduce network-based storage virtualization capabilities appropriately. HP will continue to pursue its current strategy of partnering with leading OEMs to deliver solutions, examining switch- and appliance-based approaches. HP believes that the right approach is to evolve tested, trusted network solutions rather than risk increasing SAN complexity by inventing radical new products. The goal of HP is to enable customers to employ storage virtualization with the least risk, and in the most flexible, effective ways to meet their needs.
- Work with operating system and other OEMs to incorporate server-based virtualization capabilities into server software. HP believes that this functionality should be tightly integrated with operating systems to deliver one of the best customer experiences. StorageWorks understands storage and storage virtualization, and it will strive to leverage its expertise with operating system and application vendors.
- Deliver products that enable service-oriented storage systems and networked storage environments to take virtualization—and other storage capabilities—to a much higher level. HP understands the needs to simplify operating environments, to make it easier to scale storage capabilities, and to enable storage to be more accurately and reliably delivered to applications.
- Integrate storage virtualization with broader virtualization and other strategies from HP. This will be implemented within the Adaptive Infrastructure. HP will also continue to extend the integration of storage and server management through HP Systems Insight Manager (HP SIM) and HP Storage Essentials SRM software, and Business Technology Optimization software. Storage is an important part of the HP portfolio and must be tightly integrated within it. The Adaptive Infrastructure is about total IT integration, automation, and virtualization. StorageWorks will provide the storage utility that services the Adaptive Infrastructure. HP has already started integrating storage with other resource management capabilities (by way of HP SIM); it plans to extend this significantly in the future.
- Overall, deliver the strategy so that customers can continue to deploy HP products with lowered risk. HP wants to enable customers to evolve their environments with little or no disruption to productivity.

HP has roadmaps and product teams in place to deliver this strategy. In addition, HP is actively engaging with a number of partners to deliver the portions of the strategy that call for partnerships. Recent enhancements to EVA and XP software, the ever-expanding scale-out storage product set from HP, initial service-oriented All-in-One storage systems, and the continuing evolution of Storage Essentials and consolidated server-storage management are recent examples of the strategy rollout.

Benefits of the strategy

The HP storage virtualization strategy is designed to be as painless for customers to implement as possible. From the start, customers do not have to replace anything to get StorageWorks virtualization benefits—they are embedded in all HP products. Through virtualization, StorageWorks enables storage consolidation, cost reduction, business continuity, and more—it is merely a case of creating a solution that best meets the business need. HP and its partners can help design and deploy the solution if needed. Furthermore, through continued development of HP SIM and Storage Essentials, updating existing products and/or adding storage systems as business needs demand can make it relatively painless to deploy the latest virtualization enhancements HP offers.

Lower costs by simplifying storage

Storage virtualization simplifies the environment by reducing the number of storage systems that must be managed (storage consolidation), by consolidating and standardizing management interfaces, and by making it relatively easy to scale storage systems and the storage network they are attached to. Host connections and driver software are also standardized to lessen infrastructure compatibility issues.

Accelerate growth with agile infrastructure

The flexible, scalable, extensible network storage environments created by StorageWorks products can be accessed by most of today's popular servers. StorageWorks SANs can accommodate most of today's business requirements, and can flex to accommodate changes. Growing a SAN is straightforward, as are adding file and data protection services. Our unified management infrastructure enables administrators to quickly analyze their storage environment and provides the basis for understanding how to change it when needed. StorageWorks disk subsystems are designed with the ability to easily accommodate a number of changes, with lowered or no application disruption. Indeed, virtualization technologies form a foundation that allows StorageWorks solutions to support change in many ways.

Mitigate risk with reliable storage

Virtualization technologies are at the heart of StorageWorks consolidated data infrastructures. The broad range of products from HP allows customers to select solutions at price points appropriate for a vast range of availability and recoverability requirements. StorageWorks products have a heritage of reliability, durability, backward-compatibility (new systems coexist in the same infrastructure as old systems), and longevity. The breadth of HP's storage & IT infrastructure portfolio addresses can help to protect data with disaster tolerant solutions and Implement a comprehensive backup and recovery strategy that is right for your business.

Summary

The term “storage virtualization” is in broad use, but it can mean many things. This paper has described some of the basic concepts, technologies, and implementations that are included under the virtualization umbrella. At its heart, virtualization encompasses a vast and increasing number of technologies that aggregate resources into useful pools, and may extend to provisioning and managing them to deliver interesting results. Virtualization drives business value through efficiency (higher resource utilization), consolidation, management simplification (fewer objects to manage), higher application productivity (greater availability, resilience), and more.

An Extensible Storage Virtualization Stack that models the HP view of storage virtualization has been described. The stack uses a top-down approach, viewing virtualization from the business application and end-user perspective, rather than the traditional storage-centric perspective. It is useful in understanding how storage virtualization contributes to deliver consumable storage capabilities to applications.

HP has been delivering a broad range of storage virtualization capabilities for many years, as shown by practical examples in this paper. HP has been an innovator in storage virtualization for several decades and has plans in place to continue as a leader. The ways in which HP implements virtualization and the capabilities it enables continue to advance. HP is evolving conventional storage- and network-based virtualization and ways in which virtualization can be applied to storage management. HP is simultaneously pushing virtualization to new limits by extending Storage Essentials and seeking to manage storage as a service. In addition to broad and increasingly integrated storage virtualization capabilities, StorageWorks is actively working with the rest of HP toward virtualizing and automating the entire IT environment as part of the HP Adaptive Infrastructure strategy.

For more information

For more information on HP StorageWorks virtualization, visit:
www.hp.com/go/StorageVirtualization

References

- ENSAextended technical white paper (2003)
- Storage and Adaptive Enterprise white paper (May, 2003)

Further reading

- Storage Virtualization, by Tom Clark (Addison-Wesley, 2005, ISBN 0-32-126251-4)
- SNIA Virtualization Tutorials: http://snia.org/education/tutorials/
- XP24000 Thin Provisioning White Paper
- EVA Dynamic Capacity Management Software

4AA0-7358ENW Rev. 1, February 2008