

1

Introduction

1.1 Storage Virtualization Overview

The data storage industry is one of the most dynamic sectors in information technology today. Due largely to the introduction of high-performance networking between servers and storage assets, storage technology has undergone a rapid transformation as one innovation after another has pushed storage solutions forward. At the same time, the viability of new storage technologies is repeatedly affirmed by the rapid adoption of networked storage by virtually every large enterprise and institution. Businesses, governments, and institutions today depend on information, and information in its unrefined form as data ultimately resides somewhere on storage media. Applying new technologies to safeguard this essential data, facilitate its access, and simplify its management has readily understandable value.

Since the early 1990s, storage innovation has produced a steady stream of new technology solutions, including Fibre Channel, network-attached storage (NAS), server clustering, serverless backup, high-availability dual-pathing, point-in-time data copy (snapshots), shared tape access, storage over distance, iSCSI, CIM (common information model)-based management of storage assets and transports, and storage virtualization. Each of these successive waves of technical advance has been accompanied by disruption to previous practices, vendor contention, over-hyping of what the new solution could actually do, and confusion among customers. Ultimately, however, each step in technical development eventually settles on some useful application, and all the marketing dust finally settles back into place.

No storage networking innovation has caused more confusion in today's market, however, than storage virtualization. In brief, storage virtualization is the *logical abstraction of physical storage systems* and thus, when well implemented, hides the complexity of physical storage devices

and their specific requirements from management view. Storage virtualization has tremendous potential for simplifying storage administration and reducing costs for managing diverse storage assets.

Unlike previous new protocols or architectures, however, storage virtualization has no standard measure defined by a reputable organization such as INCITS (InterNational Committee for Information Technology Standards) or the IETF (Internet Engineering Task Force). The closest vendor-neutral attempt to make storage virtualization concepts comprehensible has been the work of the Storage Networking Industry Association (SNIA), which has produced useful tutorial content on the various flavors of virtualization technology. Still, storage virtualization continues to play the role of elephant to the long lines of vendors and customers who, having been blinded by exaggerated marketing claims, attempt to lay hands on it in total darkness. Everyone walks away with a different impression. It is often difficult, therefore, to say exactly what the technology is or should be expected to do.

As might be expected, some of the confusion over storage virtualization is vendor-induced. Storage virtualization products vary considerably, as do their implementation methods. Vendors of storage arrays may host virtualization directly on the storage controller, while software vendors may port virtualization applications to servers or SAN appliances. Fabric switch manufacturers may implement virtualization services within the fabric in the form of smart switch technology. Some vendors implement storage virtualization commands and data along the same path between server and storage, while others split the control path and data path apart. Advocates of one or the other virtualization method typically have sound reasons why their individual approach is best, while their competitors are ever willing to explain in even greater detail why it is not. The diversity of storage virtualization approaches alone forces customers into a much longer decision and acquisition cycle as they attempt to sort out the benefits and demerits of the various offerings and try to separate marketing hype from useful fact.

In addition, it is difficult to read data sheets or marketing collateral on virtualization products without encountering extended discussions about point-in-time data copying via snapshots, data replication, mirroring, remote extension over IP, and other utilities. Although storage virtualization facilitates these services, none are fundamentally dependent on storage virtualization technologies. The admixture of core storage virtualization con-

cepts such as storage pooling with ancillary concepts such as snapshots contributes to the confusion over what the technology really does.

Although storage virtualization technology has spawned new companies and products, virtualizing storage is not new. Even in open systems environments, atomic forms of virtual storage have been around for years. In 1987, for example, researchers Patterson, Gibson, and Katz at the University of California Berkeley published a document entitled “A Case for Redundant Arrays of Inexpensive Disks (RAID),” which described means to combine multiple disks and virtualize them to the operating system as a single large disk. Although RAID technology was intended to enhance storage performance and provide data recoverability against disk failure, it also streamlined storage management by reducing disk administration from many physical objects to a single virtual one. Today, storage virtualization technologies leverage lower-level virtualizing techniques such as RAID, but primarily focus on virtualizing higher-level storage systems and storage processes instead of discrete disk components.

The economic drivers for storage virtualization are very straightforward: reduce costs without sacrificing data integrity or performance. Computer systems in general are highly complex, too complex, in fact, to be administered at a discrete physical level. As computer technology has evolved, a higher proportion of CPU cycle time has been dedicated to abstracting the underlying hardware, memory management, input/output, and processor requirements from the user interface. Today, a computer user does not have to be conversant in assembly language programming to make a change in a spreadsheet. The interface and management of the underlying technology has been heavily virtualized.

Storage administration, by contrast, is still tedious, manual-intensive, and seemingly never-ending. The introduction of storage networking has centralized storage administrative tasks by consolidating dispersed direct-attached storage assets into larger, shared resources on a SAN. Fewer administrators can now manage more disk capacity and support more servers, but capacity for each server must still be monitored, logical units manually created and assigned, zones established and exported, and new storage assets manually brought online to service new application requirements. In addition, although shared storage represents a major technological advance over direct-attached storage, it has introduced its own complexity in terms of implementation and support. Complexity equates to cost. Finding ways to hide complexity, automate tedious tasks, streamline administration, and

still satisfy the requirements of high performance and data availability saves money, and that is always the bottom line. That is the promise of storage virtualization, although many solutions today are still far short of this goal.

Another highly advertised objective for storage virtualization is to overcome vendor interoperability issues. Storage array manufacturers comply with the appropriate SCSI and Fibre Channel standards for basic connectivity to their products. Each, however, also implements proprietary value-added utilities and features to differentiate their offerings to the market and these, in turn, pose interoperability problems for customers with heterogeneous storage environments. Disk-to-disk data replication solutions, for example, are vendor-specific: EMC's version only works with EMC; IBM's only with IBM. By virtualizing vendor-specific storage into its vanilla flavor, storage virtualization products can be used to provide data replication across vendor lines. In addition, it becomes possible to replicate data from higher-end storage arrays with much cheaper disk assets such as JBODs (just a bunch of disks), thus addressing both interoperability and economic issues.

The concept of a system level storage virtualization strategy occurs repeatedly in vendor collateral. One of the early articles was Compaq's Enterprise Network Storage Architecture and its description of a *storage utility*. According to the ENSA document, this technology would transform storage “. . . into a utility service that is accessed reliably and transparently by users, and is professionally managed with tools and technology behind the scenes. This is achieved by incorporating physical disks into a large consolidated pool, and then virtualizing application disks from the pool.”

The operative words here are *reliably* and *transparently*. Technical remedies, like doctors, must first do no harm. Reliability implies that storage data is highly accessible, protected, and at expected performance of delivery. Transparency implies that the complexity of storage systems has been successfully masked from view and that tedious administrative tasks have been automated on the back end. The abstraction layer of storage virtualization therefore bears the heavy burden of preserving the performance and data integrity requirements of physical storage while reducing the intricate associations between physical systems to a simple utility outlet into which applications can be plugged. Part of the challenge is to get the abstraction apparition conjured into place; a greater challenge is to ensure that the mirage does not dissolve when unexpected events or failures occur in the physical world. Utilities, after all, are expected to provide continuous service

regardless of demand. You shouldn't have to phone the power company every time you wish to turn on a light.

The notion of utility applied to storage and compute resources conveys not only reliability and transparency, but also ubiquity. The simpler a technology becomes, the more widely it may be deployed. Storage networking is still an esoteric technology and requires expertise to design, implement and support. The substantial research, standards requirement definition, product development, testing, certification, and interoperability required to create operational SANs was in effect funded by large enterprise customers who had the most pressing need and budget to support new and complex storage solutions. Once a storage networking industry was established, however, shared storage expanded beyond the top tier enterprises into mainstream businesses. Leveraging storage virtualization to create a storage utility model will accelerate the market penetration of SANs and, in combination with other technologies such as iSCSI, spread shared storage solutions to small and medium businesses as well.

Currently, all major storage providers have some sort of storage virtualization strategy in place, with varying degrees of implementation in products. Upon acquiring Compaq, Hewlett-Packard (HP) inherited the ENSA (and ENSA-2) storage utility white paper and has supplemented it with its Storage Grid and other initiatives. IBM has TotalStorage with SAN Volume Controller. EMC's Information Lifecycle Management (ILM) extends storage virtualization's reach throughout the creation and eventual demise of data. Hitachi Data Systems supports array-based storage virtualization on its 9000 series systems. Even Sun Microsystems has a component for pooling of storage resources within its N1 system virtualization architecture. These vendor-driven storage virtualization initiatives reflect both proactive and reactive responses to the customers' desire for simplified storage management and are being executed through both in-house development and acquisition of innovative startups.

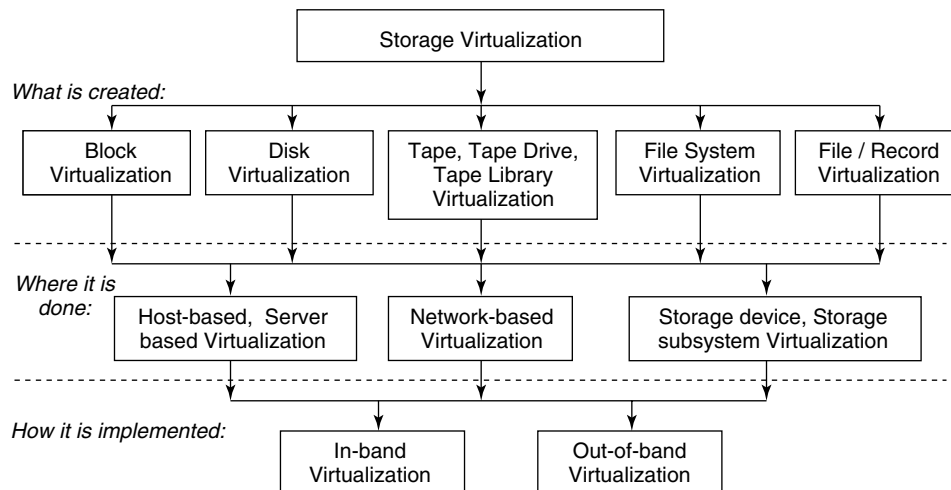
In addition, multilateral partnerships are being forged between vendors of virtualization software, storage providers, SAN switch manufacturers, and even nonstorage vendors such as Microsoft to bring new storage virtualization solutions to market. Despite the high confusion factor (and often contributing to it), storage virtualization development has considerable momentum and continues to spawn a diversity of product offerings. This is typical of an evolutionary process, with initial variation of attributes, cross-pollination, inheritance of successful features, and ultimately a natural se-

lection for the most viable within specific environments. Because storage virtualization is still evolving, it is premature to say which method will ultimately prevail. It is likely that storage virtualization will continue to adapt to a diversity of customer environments and appear in a number of different forms in the storage ecosystem.

1.2 Core Concepts

The SNIA taxonomy for storage virtualization is divided into three basic categories: what is being virtualized, where the virtualization occurs, and how it is implemented. As illustrated in Figure 1.1, virtualization can be applied to a diversity of storage categories.

What is being virtualized may include disks (cylinder, head, and sector virtualized into logical block addresses), blocks (logical blocks from disparate storage systems may be pooled into a common asset), tape systems (tape drives and tape systems may be virtualized into a single tape entity, or subdivided into multiple virtual entities), file systems (entire file systems may be virtualized into shared file systems), and file or record virtualization (files or records may be virtualized on different volumes). Where virtualization occurs may be on the host, in storage arrays, or in the network via in-



Copyright © 2003, Storage Networking Industry Association

Figure 1.1 The SNIA storage virtualization taxonomy separates the objects of virtualization from location and means of execution.

telligent fabric switches or SAN-attached appliances. How the virtualization occurs may be via in-band or out-of-band separation of control and data paths. While the taxonomy reflects the complexity of the subject matter, the common denominator of the various whats, wheres, and hows is that storage virtualization provides the means to build higher-level storage services that mask the complexity of all underlying components and enable automation of data storage operations.

The ultimate goal of storage virtualization should be to simplify storage administration. This can be achieved by a layered approach, binding multiple levels of technologies on a foundation of *logical abstraction*. Concealing the complexity of physical storage assets by only revealing a simplified logical view of storage is only a first step towards streamlining storage management. Treating multiple physical disks or arrays as a single logical entity segregates the user of storage capacity from the physical characteristics of disk assets, including physical location and unique requirements of the physical devices. Storage capacity for individual servers, however, must still be configured, assigned, and monitored by someone. Although one layer of complexity has been addressed, the logical abstraction of physical storage alone does not lift the burden of tedious manual administration from the shoulders of storage managers.

To fulfill its promise, storage virtualization requires *automation* of the routine soul-numbing tasks currently performed by storage administrators. Allocating additional storage capacity to a server, for example, or increasing total storage capacity by introducing a new array to the SAN are routine and recurring tasks begging for automation.

Ideally, storage automation should be *policy-based* to further reduce manual intervention. Virtualization intelligence should automatically determine whether a specific storage transaction warrants high availability storage or less expensive storage, requires immediate data replication off-site or simple backup to tape on a predetermined schedule, or becomes part of a lifecycle management mechanism and retired at the appropriate time. A tiered infrastructure leveraging *class of storage* provides policy engines with repositories that meet the requirements of different types of storage transactions.

Finally, storage virtualization should become *application-aware*, so that policy-based automation responds to specific data types and identifies the unique needs of each upper layer application. Digital video, for example, gains more consistent performance if it is written to the outer, longer tracks of physical disks. Likewise, financial transactions for banking or

e-commerce would benefit from frequent point-in-time copy policies for safeguarding most current transactions. An intelligent entity within the storage network that monitors and identifies applications and, based on preset policies, automates the handling of data for class of storage brings storage virtualization much closer to the concept of utility.

Application-aware storage virtualization provides the potential for dynamic communication between upper-layer applications and the storage services beneath them. As demonstrated by Microsoft's initiative to provide enhanced interfaces between the operating system and storage utilities such as snapshot, mirroring, and multi-pathing, it will become possible for upper layer applications to more fully leverage underlying storage services. *Storage virtualization-enabled applications* could, for example, seek out those services that more closely align to their current requirements for capacity or class of storage or, via APIs, inform the storage network of unique policies that should be enforced.

The viability of storage virtualization is enhanced by, but not dependent on, *interoperability* between storage assets. Although storage virtualization vendors highlight the benefits their products bring to heterogeneous data centers that may include HP, IBM, EMC, HDS, or other storage, some customers are quite happy with single vendor, homogeneous storage solutions. Logical abstraction of physical storage, automation of tedious tasks, policy-driven data handling, and application awareness have significant value for both single-vendor and multivendor storage networks. Interoperability, however, is a key component of the storage utility, since a utility should accommodate any type of application, operating system, computer platform, SAN infrastructure, storage array, or tape subsystem without manual intervention.

As shown in Figure 1.2, storage virtualization technology is a layered parfait of more sophisticated functionality that drives toward greater degrees of simplicity. Current products provide bits and pieces of a virtualized solution, from elementary storage pooling to limited automation and policy engines. Vendors and customers, however, are still struggling toward more comprehensive, utility-like storage virtualization strategies that fully leverage the potential of the technology.

Where the intelligence to do all these virtual things resides is interesting from a technical standpoint, but of less interest to the ultimate consumers of storage resources. The transparency that storage virtualization provides for storage assets should eventually apply to the storage virtualization solution

Virtualization-Enabled Applications	Services
Application-Aware Storage Virtualization	Dynamic Capacity Allocation Tape Backup Processes Storage Consolidation Heterogeneous Storage Tiered Storage (ILM) Point-in-time Snapshots Replication / Mirroring Auditing / Service Billing
Policy-based Management	
Automation of Storage Processes	
Logical Abstraction Layer	
Physical Storage Systems	

Figure 1.2 Storage virtualization enables successive layers of advanced functionality to fully automate storage administration.

itself. The abstraction layer that masks physical from logical storage may reside on host systems such as servers, within the storage network in the form of a virtualization appliance, as an integral option within the SAN interconnect in the form of intelligent SAN switches, or on storage array or tape subsystem targets. In common usage, these alternatives are referred to as *host-based*, *network-based*, or *array-based* virtualization. Each approach has strengths and weaknesses that we will describe in subsequent chapters.

In addition to differences between where virtualization intelligence is located, vendors have different methods for implementing virtualized storage transport. The *in-band* method places the virtualization engine squarely in the data path, so that both block data and the control information that govern its virtual appearance transit the same link. The *out-of-band* method provides separate paths for data and control, presenting an image of virtual storage to the host by one link and allowing the host to directly retrieve data blocks from physical storage on another. In-band and out-of-band virtualization techniques are sometimes referred to as symmetrical and asymmetrical, respectively, but for the sake of accuracy and simplifying the virtualization vocabulary, this text uses the hyphenated band terminology.

Simplifying storage administration through virtualization technology has many aspects. Centralizing management, streamlining tape backup processes, consolidating storage assets, enhancing capacity utilization, facilitating data integrity via snapshots, and so on, are not really attributes of storage virtualization, but rather beneficiaries of it. Storage consolidation, for example, is enabled by networking storage assets in a SAN. If there is only one large disk array to manage, virtualization may not contribute sig-

nificantly to ease of use. If there are multiple disk arrays, and in particular, arrays from different vendors in the SAN, storage virtualization can help streamline management by aggregating the storage assets into a common pool. Current vendor literature is often punctuated with exclamations about the many benefits of storage virtualization, and then proceeds to focus on backup, snapshots, and so on. In some cases, customers may indeed benefit from these enhanced utilities, but may not need to virtualize anything to use them. As always, the starting point for assessing the potential benefit of a new technology is to understand your application requirements and your existing practices and measure potential benefit against real need.

In the following chapters, we will examine the basic problems that storage virtualization is trying to solve, the different approaches technologists are proposing, and the intersections between current capabilities and productive use by customers. In the process, hopefully, the confusion factor generated by this subject can be kept to a minimum as the virtual state of this technology is separated from its real one.

1.3 Chapter Summary

Storage Virtualization Overview

- Storage virtualization depends upon lower layer standards but is not defined by a standards body.
- Storage virtualization provides an abstraction layer over physical storage systems and enables multiple storage assets to be treated as a single logical entity.
- Granular forms of storage virtualization occur in technologies such as RAID.
- The value of storage virtualization is in its ability to automate tedious administrative tasks and reduce costs of storage management.
- The Compaq ENSA initiative leverages storage virtualization for creation of a storage utility.
- Storage virtualization must maintain reliable and transparent access to underlying storage assets.
- Virtualization technology is being driven by all major storage providers.

Core Concepts

- Virtualization provides a foundation of logical abstraction upon which more advanced functionality can be built.
- To streamline storage operations, logical abstraction should support automation of routine administrative tasks.
- Automated processes should be policy-based to further reduce manual intervention.
- Logical abstraction, automation, and policy-driven management enable storage virtualization to support application-aware storage management.
- The tight integration of operating systems and virtualization technology will enable applications to more fully leverage virtualized storage.
- Storage virtualization facilitates interoperability of storage systems.
- Virtualization technology can be supported on hosts, networks, or storage arrays and tape subsystems.
- In-band virtualization provides both control and data along the same connection path.
- Out-of-band virtualization separates control and data on different connection paths.
- Not all storage environments require or benefit from virtualization.

