



The Need for Storage Virtualization

Storage Administration in a SAN World

A white paper from Vicom Systems, Inc.



Table of Contents

Introduction / 3

SAN Management Problems / 4

24-HOUR DATA AVAILABILITY / 4

NETWORK REDUNDANCY / 4

NON-DISRUPTIVE BACKUP / 4

STORAGE EFFICIENCY / 4

CENTRALIZED ADMINISTRATION / 4

SECURE ACCESS CONTROL / 4

MIXED STORAGE POOL, MIXED SERVER FARM / 5

SCALABILITY / 5

Vicom Systems Storage Virtualization Engine / 6

HIGH AVAILABILITY / 7

NON-DISRUPTIVE, LAN-FREE BACKUP / 7

STORAGE EFFICIENCY / 8

SCALABLE, HIGH-PERFORMANCE INFRASTRUCTURE / 9

CENTRALIZED SAN ADMINISTRATION / 9

CENTRALIZED ACCESS CONTROL / 9

FULL HETEROGENEOUS INTEROPERABILITY / 9

Storage Service Provider: An Application Focus / 10

ACCESS CONTROL / 10

STORAGE EFFICIENCY / 10

HIGH AVAILABILITY / 11

MULTIPROTOCOL INTEROPERABILITY / 11

CENTRALIZED ADMINISTRATION OF MULTIPLE SANs / 11

Conclusion / 11

Introduction

The storage networking debate is over. Protocol discussions aside, nobody doubts that the future of storage is networks. The subject now consuming the thoughts of enterprise Information Technology (IT) managers and Service Provider executives is how to manage the storage. This discussion has to be resolved quickly. The Gartner Group, among other research organizations, is predicting that most enterprises can expect storage capacity to double each year for the next five years, and Service Providers can expect it to double every quarter. Most organizations are not yet ready to handle this growth.

The focus on storage interoperability has provided a blueprint that clearly enumerates the benefits of a dedicated network for storage. In reality, most enterprises will be managing a mixed bag of newer Storage Area Network (SAN) technology and existing SCSI and SSA connected storage for some time. Providing centralized storage administration for this diverse array of storage technologies is essential: Forrester Research recently estimated that for every dollar spent on storage hardware, an organization should expect to spend seven dollars on storage management. The productivity gains from centralized storage administration in a distributed computing environment are self-evident.

The key to managing a diverse pool of storage hardware is abstraction. Abstraction enables the physical characteristics of storage hardware to be masked, giving users a pool of virtual storage to exploit.

To some extent, storage vendors already provide this virtualization facility on high-end storage hardware. The move to a networked storage model now provides an ideal platform for virtualizing all storage, regardless of vendor. Virtualization enables a common suite of management services to be administered across all storage hardware, regardless of the built-in capabilities of the individual device. Mirroring, LAN-free backup, access control, and storage partitioning are all services that can be uniformly applied by an administrator to a pool of virtualized storage hardware.

The Vicom Storage Virtualization solution provides a storage management infrastructure based on abstraction. The network-based Storage Virtualization Engine (SVE) resides between UNIX/Windows application servers and storage hardware, providing a common logical view of all networked storage. Administrators have a centralized

platform from which they effectively manage all storage resources. A wide array of services including three-way mirroring, non-disruptive backup, secure access control, and LUN carving are available through a single user interface.

This paper looks at the problems facing storage administrators as they experience the tremendous growth in the volume and diversity of the storage they manage. Data availability, network redundancy, centralized administration and control, and efficient storage allocation are issues that now come up daily. A special focus is given to Storage Service Providers (SSPs). The problems faced by SSPs, with business models dependent on efficiently managing storage for their customers, are early warning signals for enterprise IT. The challenges SSPs face—managing huge amounts of storage today—will face enterprise IT soon.

SAN Management Problems

24-HOUR DATA AVAILABILITY

Any system failure that results in data loss can mean disaster. Downtime, lost business, lost productivity, and the potential for legal problems can cripple an organization. Data availability is serious business.

For storage hardware, availability invariably means mirroring. Mirroring creates two images of the operational data. If one image fails, the second is available to continue read/write operations. The mirror provides administrators with a window of time in which to repair or replace the failed copy without affecting business applications.

Simple two-way mirroring has been available on high-end storage hardware for many years. The challenge has been to provide three- and four-way mirroring in order to limit exposure to double- and triple-fault outages, in which multiple mirrors fail before administrators can repair them. Mirroring across miscellaneous storage hardware is becoming a necessity as deployment of low-cost Just a Bunch of Disks (JBODs)—which lack mirroring capabilities—increases.

NETWORK REDUNDANCY

Networked storage introduces a new range of infrastructure components that can fail. Mirroring can provide high availability at the storage hardware level, but redundancy covers availability in the network. Redundancy generally means providing multiple paths through the network that lead to the same place. In the event of a path failure, the application server can take an alternate route to the same storage location.

NON-DISRUPTIVE BACKUP

Shifting the network-congesting traffic associated with backup from the Local Area Network (LAN) onto the dedicated storage network is a huge benefit of SAN architectures. However, LAN-free backup can still impact business applications by creating a resource conflict between the backup server and the application server as both vie for the limited I/O capacity of a storage device.

Full-image data copy techniques offer a solution to the I/O conflict. Using the same low-impact copy technology from mirroring, a point-in-time image of the storage device can be accessed by the backup server, relieving contention on the operational data.

STORAGE EFFICIENCY

Storage efficiency has as much to do with the cost of administration as with the cost of hardware. The more storage there is to manage, the greater the administrative burden.

The move from captive storage—storage permanently attached to a server—to pooled network storage creates substantial improvements in the efficient sharing and allocation of storage resources. LUN carving, the partitioning of large devices into smaller units, can boost efficiency even further by allocating space to application servers as they need it, instead of over-allocating an entire multi-gigabyte drive.

CENTRALIZED ADMINISTRATION

Despite the trend toward consolidation of storage in a SAN, administrators face immense challenges as the volume of storage and variety of storage hardware increases. Centralized administration of the growing storage resource offers both productivity gains and a single platform for consistent management policies. Data protection, resource allocation, availability, performance, and security have a significant impact on business applications and require effective and uniform management.

SECURE ACCESS CONTROL

As storage resources move away from the isolation of direct server attachment and into the world of networks, security becomes an important issue. Availability of a storage device on the network does not mean it is open to access by anyone.

Trusting application servers to enforce access restrictions can work for some organizations, but for most it is a compromise. In general, end-user application access is managed at the application server, and lower level access to allocated networked storage is managed centrally by the storage administrator. Without centrally administered access control, inadvertent or malicious access from an unauthorized or misbehaving server can jeopardize security and integrity of the entire storage network.

Storage network access control mechanisms include: host-based zoning, which relies on agent software running on the application server; and switch-based zoning, which physically controls access at the channel or port on a network switch.

Placing enforcement and configuration of access restrictions in each application server can work for some organizations. But, for most, it introduces additional complexity of having to reconfigure each host with the appropriate agent software or Host Bus Adapter (HBA), increasing the risk of human error during configuration. And again, having the security layer at the application server makes it easier for a poorly configured or unauthorized server to gain access to the entire storage network. Only switch-based zoning meets the requirement for centralized configuration and non-server-based security. However, it lacks efficiency because large storage devices are physically mapped to a single port at the switch and cannot be divided into smaller multiple zones for granular allocation of storage capacity among servers.

Ideally, secure access zoning for storage networks must be centrally administered independently of the application server and must support efficient sharing of storage capacity. Neither switch-based nor server-based zoning meet all of these requirements.

MIXED STORAGE POOL, MIXED SERVER FARM

Most computing environments contain a wide variety of server and storage hardware. Getting this mix of platforms to talk together is a monumental challenge. SAN architectures can overcome this obstacle by providing a model for pooling resources. But the reality of SAN deployments with a single protocol and single operating system is that large numbers of servers and storage hardware can be left out of the pool.

By allowing servers and storage that support dissimilar communication protocols, such as SCSI or SSA protocols, into the storage pool, much greater efficiency can be achieved and corporations can protect their investment in non-Fibre Channel storage and server equipments. When all enterprise storage is in the pool, common access and data protection policies can be easily implemented and enforced.

SCALABILITY

Scalability is a significant issue for any storage administration platform aimed at enterprise IT and SSPs. The ability to accommodate growth of storage, servers, and network connections is essential. With storage capacity doubling every year, a solution that

fits the environment today can quickly be outgrown if tomorrow's needs are not considered.

Current SAN administration platforms vary in their ability to accommodate growing storage capacity. The alternative architectures include:

- » Host-based, in which software resides on each application server requiring access to the SAN
- » Storage-based, in which storage management software is provided by the vendor of the storage hardware and operates only on that platform
- » Network-based, with in-band (symmetric) and out-of-band (asymmetric) alternatives

Both host-based and storage-based management solutions offer limited scalability. For administrators coping with SAN access from a diverse array of application servers, maintaining storage management software on each host can be time-consuming and impractical. Storage-based solutions only tend to apply to the storage hardware they were supplied with. This limits the scalability of the software, tying the organization to storage from a single vendor.

In-band and out-of-band refer to the location of the network-based storage administration platform, whether it is in the path of the data request (in-band) or out of the data path (out-of-band). Out-of-band solutions generally require help from the host to identify the location of the management platform. This help can come from software running on the host or from the HBA. The reliance on the host server raises the same concerns as host-based storage management and the administrator's ability to effectively exercise control.

In-band solutions offer the most scalable and manageable alternative with the major architectural consideration being the use of a general-purpose computing platform—a PC—or a specialized, single-purpose device that is optimized for storage networking.

Vicom Storage Virtualization Engine

The Vicom Storage Virtualization Engine (SVE) combines a single, logical view of all networked storage with a simple, centrally administered suite of storage management services. The network-based SVE is an integrated solution that resides between heterogeneous application servers and storage hardware. Built on the Vicom Independent Distributed Routing (VIDR) architecture, Vicom SVE maps inbound storage requests to physical storage devices, providing logical, hardware-independent access to the storage pool from application servers.

With a single user interface for all networked storage hardware, a consistent suite of services—

including mirroring, partitioning, concatenation, access control, Instant Copy, and backup/restore—can be administered. Virtualization of the physical storage allows precise partitioning and access control, resulting in efficient and secure storage use.

The Vicom Independent Distributed Routing (VIDR) architecture provides an exceptionally scalable, redundant architecture for managing networked storage. Multiprotocol support enables Vicom SVE to incorporate a wide variety of storage hardware and server operating systems into the networked storage pool.




Figure 1. Sample Vicom SVE topology

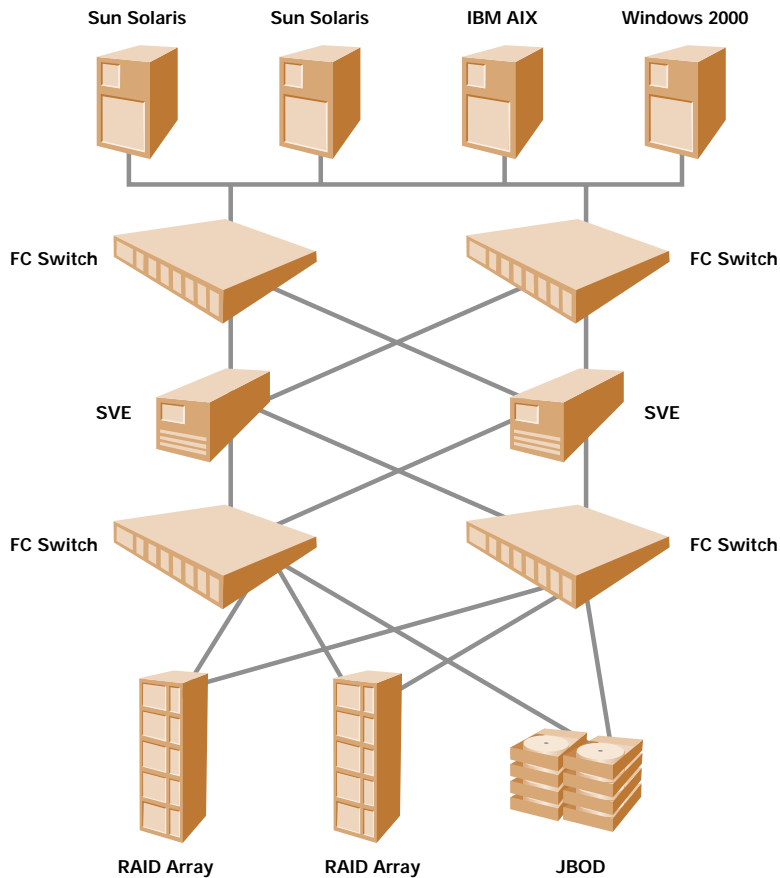


Figure 2. Redundant SVE topology

HIGH AVAILABILITY

Vicom SVE provides high-availability data access through advanced mirroring techniques and support for redundant network configuration.

Vicom SVE supports two-way and three-way mirroring, safeguarding business applications against catastrophic data loss. The unique virtualization offered by Vicom SVE enables mirroring across physically dissimilar devices. JBODs, high-end storage arrays, Fibre Channel, SSA, and SCSI hardware can all be sources and targets of mirror copies. When a failure occurs, a spare drive from the spare pool of storage is transparently allocated and automatically synchronized with the surviving mirror.

Additional mirrors can be configured using the uniquely addressable Instant Copy (IC) drive. The

IC drive can be synchronized with a mirrored set and then broken off and used for data migration, backup, or recovery without impacting access to the original data.

Vicom SVE hardware modules hold the logical to physical mapping data and provide a redundant network topology that guarantees a path to data in the event of an outage. Each SVE hardware module contains a full set of mapping data and automatically synchronizes changes with other SVE hardware modules.

NON-DISRUPTIVE, LAN-FREE BACKUP

Vicom SVE Instant Copy drive provides a separately addressable copy device for off-loading backup activities.

Vicom Storage Virtualization Engine

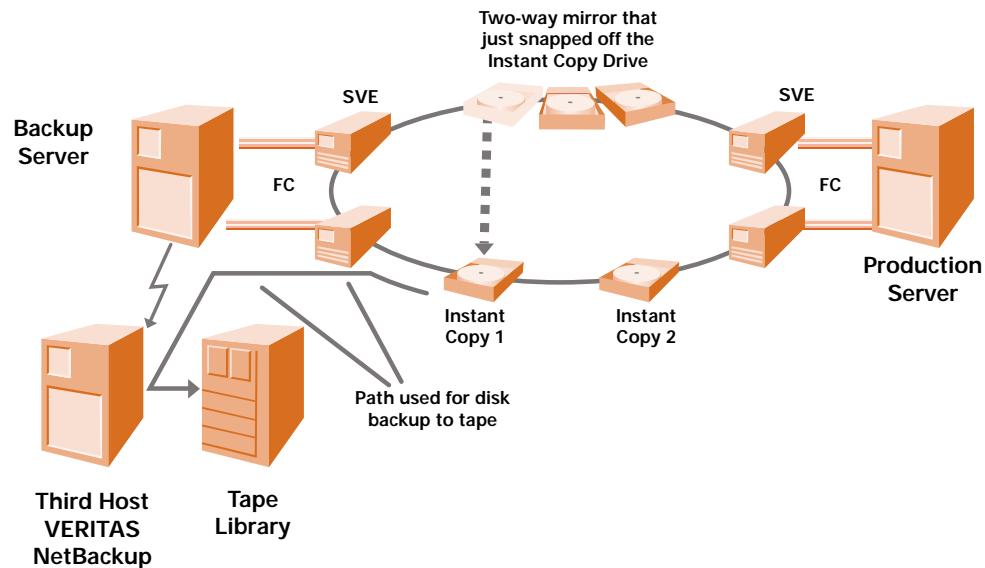


Figure 3. IC drive used for simultaneous backup and copy operations

Vicom SVE Instant Copy drive can be used with third-party backup and recovery software to provide a non-disruptive, LAN-free backup facility. The IC drive produces a point-in-time, logical device snapshot and is configured to use a separately addressable LUN. The drive can be mounted to a backup server alleviating I/O contention on the operational data. Multiple IC drives enable simultaneous copy and backup operations.

Vicom SVE provides a scripting interface on most of the major UNIX and Windows operating systems for support of multivendor backup and recovery software products, and database command-line interfaces.

STORAGE EFFICIENCY

Vicom SVE architecture supports efficient use of storage through pooling, composite drives, LUN carving, and LUN masking, and support for incremental growth of the storage pool.

Vicom SVE provides support for an expanded pool of storage hardware that encompasses multiple industry-standard protocols, including SCSI, SSA, and Fibre Channel. Vicom SVE increases the participation of all storage hardware in the storage pool and allows efficient sharing of all storage resources.

The abstraction of physical hardware characteristics to a logical layer enables Vicom SVE to create composite drives—logical drives that look and feel like physical devices to an application server. A composite drive is a concatenation of physical devices on the SAN allowing, for example, several small, high-performance disk drives to appear as a single, large drive.

Virtualization also enables the efficient allocation of storage through support for LUN carving. LUN carving allows a physical device to be partitioned, creating multiple virtual drives that can be as small as 500 MB. This allows administrators much greater flexibility when allotting resources and

permits efficient, incremental allocation of space to application servers.

LUN masking, on the other hand, is a mechanism for restricting application server access to a defined set of storage devices. The technology is also effective in providing storage-use policies. Vicom SVE supports primary and secondary zoning settings that allow application servers to be dynamically switched between the two settings.

An example of multiple storage access policies might be a business application that requires access to a database during normal business hours, but at night requires additional workspace for performing complex batch operations. Selectively providing access to the additional space, using Vicom SVE LUN masking policies, can result in enhanced sharing of the pool of storage, with other users able to access the workspace during business hours.

SCALABLE, HIGH-PERFORMANCE INFRASTRUCTURE

Vicom SVE offers infinite scalability with each SVE hardware module adding to the total throughput and bandwidth of the network.

The distributed in-band architecture of Vicom SVE uses a special-purpose operating system optimized for high-performance storage networking. Vicom SVE offers almost infinite scalability, unmatched by other virtualization implementations. Each SVE hardware module provides 8,000 I/O per second at a bandwidth of 100 MB per second. Adding more modules provides additional processing capacity. Scalability is linear, since additional SVE hardware modules can provide more bandwidth and throughput.

CENTRALIZED SAN ADMINISTRATION

Vicom SVE offers a scalable approach to storage management, enabling centralized administration of the entire pool of networked storage, including multiple SANs.

With a single graphical user interface or command-line interface, Vicom SVE provides a centralized view of all networked storage resources, including multiple SANs. Vicom SVE is independent of the server operating system and supports multi-vendor storage hardware.

Centralized storage administration provides the basis for a consistent suite of data management stan-

dards and policies for all storage resources. Vicom SVE enables storage administrators to exercise much greater control over a wider array of networked storage than would otherwise be possible.

CENTRALIZED ACCESS CONTROL

Vicom SVE provides secure, independent, and efficient/granular access control based on LUN masking and LUN carving.

The combination of LUN masking and LUN carving gives storage administrators precise control over access to storage devices in the SAN, providing the granularity needed to maintain efficient storage allocations. SVE LUN masking and LUN carving policies are enforced at the SVE hardware modules that are centrally configured through the SVE software suite. This eliminates the problems associated with host-based and out-of-band solutions, which require controls to be configured in software or the HBA at each application server.

FULL HETEROGENEOUS INTEROPERABILITY

Vicom SVE provides full interoperability, ensuring a seamless transition to SAN and investment protection for existing storage and server equipment.

Vicom SVE provides enhanced SAN interoperability by enabling non-Fibre Channel storage hardware and non-Fibre Channel-ready servers to be part of the SAN. Supporting SSA and SCSI protocols through protocol-converting SVE hardware modules, Vicom SVE simplifies the integration of existing resources into the Fibre Channel SAN.



Storage Service Provider: An Application Focus

SSPs highlight the storage management problems administrators are likely to face as storage capacity escalates. Management overhead directly affects the profitability of the SSP, and efficient control of the hundreds of terabytes of storage they manage for customers is essential. Of all businesses, SSPs are probably the most keenly aware of the promise of storage virtualization. Using the SSP as a case study offers a valuable insight into the future of enterprise IT.

ACCESS CONTROL

The corporate business data housed at the SSP data center is likely to be one of the most valuable assets the SSP customer owns. Securing that asset from unauthorized access is a critical issue for the SSP.

The stringent access control requirements facing the SSP provide a litmus test for competing SAN management solutions. Host-based access control, with its delegation of responsibility to the host server, is unlikely to satisfy such strict security requirements. Host servers are not under the control of the SSP. Switch-based zoning can be effective at limiting access, but is a blunt instrument resulting in inefficient and costly over-allocation of storage to customers.

The precise access control offered by Vicom SVE gives SSPs absolute control over security without compromising storage efficiency. With control over access residing in the network rather than at the host, SSP clients can be assured of bulletproof security.

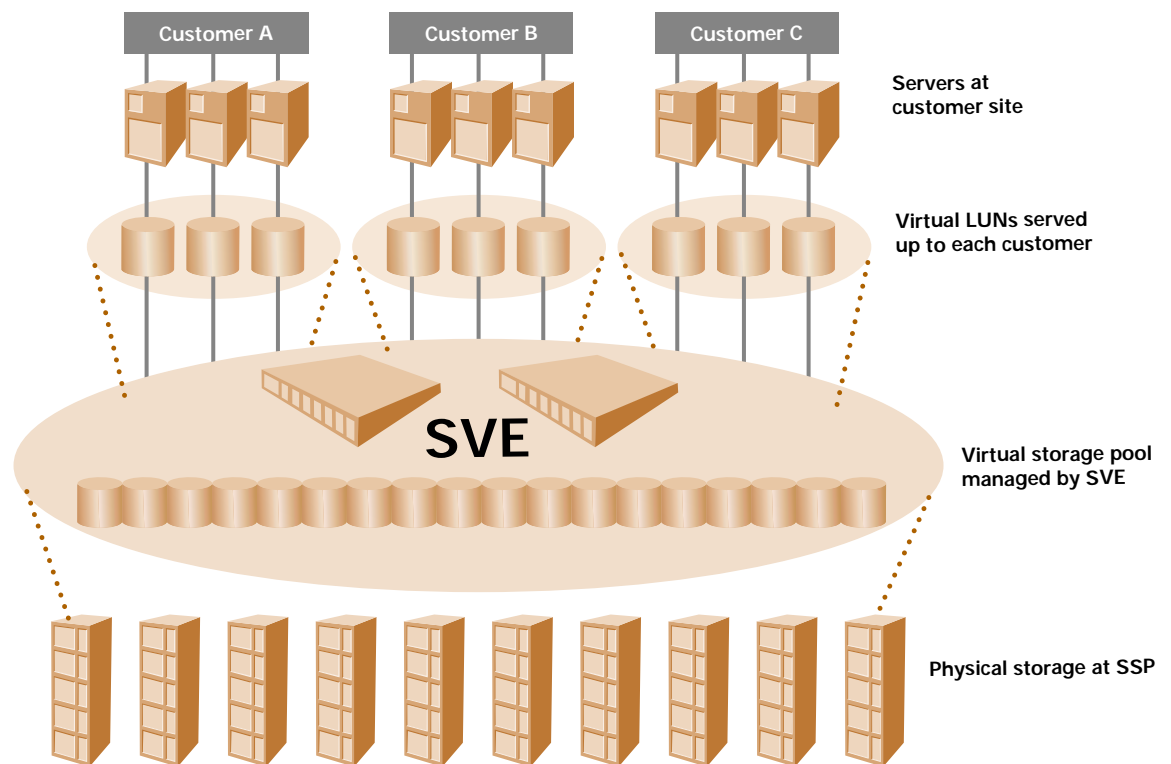


Figure 4. Sample SSP topology

STORAGE EFFICIENCY

With storage capacity measured in hundreds, if not thousands, of terabytes, SSPs are critically aware of the costs of inefficient storage allocation. Wasted space is expensive, both in hardware costs and in the extra burden it places on storage administrators. Vicom SVE provides the most effective method for maintaining streamlined storage allocation.

SVE virtual drives provide a mechanism for partitioning large storage devices into smaller units, the smallest unit available being 500 MB. Storage administrators gain the flexibility of allocating space to customers in the exact amount needed.

HIGH AVAILABILITY

SSPs have no illusions about the true cost of data availability. With detailed Service Level Agreements spelling out the costs of downtime, the SSP is likely to suffer financial penalties if a customer's data is unavailable.

The redundancy and enhanced mirroring features of Vicom SVE provide high-availability coverage for the SSP. Support for two- and three-way hardware-agnostic mirroring guarantees continuous availability, while mirroring across hardware from different storage vendors gives SSPs the flexibility to make cost-effective multivendor storage sourcing decisions. The distributed architecture of Vicom SVE also provides a robust, fully redundant network with no single point of failure.

MULTIPROTOCOL INTEROPERABILITY

SSPs have a vested interest in providing the widest possible connectivity to their storage. By limiting access from older servers that do not comply with Fibre Channel connectivity requirements, SSPs risk alienating a large group of potential customers.

Vicom SVE allows true heterogeneous access to the SAN, removing limitations to sharing storage. Using multiprotocol virtualization, Vicom SVE supports access to the SAN from older servers that are not Fibre Channel ready.

CENTRALIZED ADMINISTRATION OF MULTIPLE SANs

Of primary importance to SSPs is the ability to reduce the cost of managing storage. A centralized view of all storage resources is an essential compo-

nent of improving storage administration productivity. Without a centralized view of storage and the necessary management services, administrators are forced to accommodate multiple sets of operating policies and standards in managing the diverse range of storage under their control.

Vicom SVE offers a single user interface to all storage in the network. Storage management services can be administered uniformly to all storage resources at the SSP data center.

Conclusion

Recent research from Gartner, Forrester, and IDC all point to rocketing growth in IT storage capacity. The volume of storage managed by even modestly sized organizations is likely to double annually. Gartner points to a need for enterprise IT to scale storage capacity tenfold while retaining the same "people print" for managing the storage. The only way to achieve such a productivity boost is to provide an efficient administration platform for all networked storage resources.

For Storage Service Providers, administration problems are even more acute. With the potential for thousands of terabytes of storage under management, it is essential to find effective ways to centrally administer and control access. Host-based and storage-based administrative alternatives fail to meet the specific needs of the SSP for centrally administered, autonomous control. Host-based solutions fail to address essential security issues. Single-vendor storage-based software limits the SSP's ability to make multivendor, cost-effective storage hardware choices.

Vicom SVE provides the only effective solution to the problem of enterprise network storage administration. Virtualization holds the key to productivity gains that will allow enterprises to keep pace with the growing demand for storage. With many years of experience in managing business-critical data connectivity, Vicom Systems has the domain expertise and proven track record to provide enterprise IT and SSPs with the best networked storage administration model.



Vicom Systems, Inc.

47281 Bayside Parkway

Fremont, CA 94538

T 510.743.1130

F 510.743.1131

www.vicom.com

© 2001 Vicom Systems, Inc.

All rights reserved. Trademarks and service marks appearing herein are property of their respective owners.

05/01 310-606-098