TANEJA GROUP
TECHNOLOGY ANALYSTS

S O L U T I O N   P R O F I L E

Igniting Storage Consolidation: 3PAR Proves Virtual Machines Are Not Just For Consolidating Servers

November 2007

THE NET NET

Over the past five years, server consolidation has become one of the top infrastructure spending priorities of IT departments. The catalyst for this shift in purchasing and deployment behavior can be directly traced back to the rise of the modern day hypervisor and server virtualization software. Server virtualization technologies, like VMWare and Xen, have allowed users to reap the true promise of a consolidated, centrally managed server infrastructure by eliminating management points and driving up CPU utilization. Despite its compelling ROI, widespread server consolidation only became feasible after the technology existed to satisfy three core consolidation requirements – isolation of application workload and administration, QoS for applications, and high utilization of shared physical server resources (e.g. CPU, memory, and I/O). The introduction of the hypervisor addressed these three technical requirements adequately, removing the last hurdle to server consolidation.

From our vantage point, we see the same dynamic unfolding in the storage consolidation market. Until recently, storage consolidation has lacked products that could satisfy this same technological trifecta. There was no analogous technology to hypervisors in storage systems that enabled true QoS and isolation for a given I/O workload, while still ensuring the optimal utilization of all physical resources within the storage system. With the recent introduction of virtual array technology, the storage consolidation market is poised to undergo a similar catalytic event to the one that occurred in the server market when hypervisors were released.

In the following profile, we will chart how storage consolidation is following a similar trajectory to how server consolidation evolved. To that end, we will examine the different technologies employed in server and storage consolidation and assess their ability to deliver the trifecta of isolation, QoS, and utilization. After that, we will spotlight 3PAR Virtual Domains, a new storage virtualization capability that creates virtual arrays on top of a single massively scalable physical storage system. We will discuss how Virtual Domains represents a transformative step forward in terms of removing the last remaining obstacles and objections to storage consolidation and ultimately ushers in a new era of self-service storage.

The Evolution of Open Systems Server Consolidation

The server infrastructure of most organizations has gone through multiple transitions over the past three decades. Centralized mainframe computing gave way to decentralized computing based on systems, arranged in client-server and, most recently, in web-based, scale out computing architectures. In the distributed computing wave, servers were typically dedicated to single applications. Because servers became so cheap during the client server and web-based computing build outs, IT deployed them in great numbers with little concern for the operating costs created by new management points. This led to server sprawl and created the conditions that drove server consolidation efforts.

From our perspective, the server consolidation market has gone through three distinct phases, characterized by different technological approaches to consolidation. The following phases are:

- Phase 1: Servers with No Partitioning
- Phase 2: Servers with Physical Server Partitions
- Phase 3: Server Virtualization and Hypervisors

In the following paragraphs, we will examine the technologies used in each of these phases with an eye for evaluating how they did or did not meet IT users' requirements for isolation, QoS, and resource utilization.

PHASE 1: SERVERS WITH NO PARTITIONING

The operating systems of the client-server and later web-based computing revolutions supported true multi-tasking, allowing multiple applications to run on a single operating system image. Unlike mainframes, open systems servers were not originally designed to run multiple copies or versions of the operating system concurrently on the same physical server hardware. Operating systems like Linux, Windows, and early versions of UNIX had no partitioning or hypervisor capabilities engineered into the operating system. As a result of their multi-tasking capabilities, these systems could in theory be consolidation platforms, sweeping applications running on multiple different servers and running them centrally on a single server with a single OS image. However, in reality, consolidating multiple different applications onto a single server with a single OS image proved too difficult. For example, applications running on Microsoft Windows, the most popular and widespread operating system, were particularly difficult to consolidate because they shared common DLLs and system resources. As those DLLs were upgraded as part of service pack upgrades, older applications that relied on the same component would break and cease to function. The end result is that to enable true server consolidation, a server needed to be able to isolate applications and their operating environments from one another.

In short, little server consolidation occurred because the technologies available in open systems operating systems did not provide effective isolation for applications running on a shared hardware platform. In the absence of partitioning functionality, application QoS became a matter of sizing servers upfront for the current and anticipated demands of a given application. The net result is that IT users chose to create isolation and QoS through the physical segregation of applications onto discrete server hardware. However, physically provisioning a server for each application resulted in exceptionally poor utilization rates. Thus, the practice of dedicating server hardware to a single application primed the economic pump for server consolidation once the technologies to accomplish isolation and QoS became available in the marketplace.

PHASE 2: SERVERS WITH PHYSICAL SERVER PARTITIONS

In the second phase of server consolidation, high end UNIX vendors introduced the concept of physical server partitions and promoted them as a means to achieving server consolidation. Physical server partitioning is a mechanism that allows an administrator to carve up a single server into multiple domains or partitions, dedicating specific physical resources, such as CPU and memory to a given partition or domain. Each partition appears to clients as an independent server and is administered independently. Major UNIX vendors, such as Sun (with E10000 Domains), HP (with Superdome nPARs), and IBM (with P-Series LPARs), all introduced physical server partitioning schemes.

The physical server partitioning approach had two strong benefits. First, an application running within a physical partition was completely secure and isolated from other applications and partitions. In short, a single badly behaved application could not affect the performance or operation of another application operating in another partition. Second, within the limits of the underlying physical server hardware, administrators could allocate additional physical CPU and memory resources to a partition in order to scale performance of the operating environment. As a result, physical server partitions improved overall QoS for applications running within a partition.

Nonetheless, the major drawback of the physical partitioning scheme was that it detracted from the overall utilization of the server hardware. In short, by dedicating physical CPU and memory to a particular partition, CPU cycles that might have been consumed by other partitions were trapped and went unused. Nonetheless, physical server partitions represented a distinct improvement in terms of flexibly allocating physical CPU resources to applications with varying SMP requirements. Average utilization rates were above the 5 to 15% range because liberal sizing was not required upfront.

PHASE 3: SERVER VIRTUALIZATION AND HYPERVISORS

With the advent of the hypervisor, widespread server consolidation became feasible to achieve and the key technical barriers of isolation, QoS, and utilization were removed. Hypervisors allows an x86 server to run multiple virtual operating environments concurrently on the same physical hardware. Each of these virtual operating environments or virtual machines encapsulates a fully functioning, secure and isolated operating system and patch set, along with the application. The hypervisor allows multiple virtual machines to run different operating systems and patch sets without effecting the operation of other virtual machines. To that end, administrative privacy and control is enforced on a virtual machine by virtual machine basis.

Unlike physical server partitions, all the physical resources (CPU, memory, and I/O) of the server are shared across all the virtual machines. As a result, in a virtualized server, there are no dedicated resources. Therefore, hypervisors allow users to dramatically improve the CPU utilization of their server infrastructure, allowing IT to get more "bang for its buck." Moreover, virtual machines provide natural isolation, encapsulation, and security for applications and their operating environments. Each virtual machine acts and is managed as a separate physical server. As a result, each application and its corresponding operating system and patch set level is isolated from one another, while still sharing the same underlying physical hardware.

Lastly, given the fact that the hypervisor can schedule virtual machines across all available physical resources, a virtual machine can achieve good performance and QoS assuming the actual physical server hardware is not redlining. For example, VMWare ESX allows a virtual machine to be scaled across up to four CPUs or cores with its Virtual SMP option. Hence, hypervisor QoS has been sufficient for most mainstream applications. Only the most demanding, compute intensive applications require more CPU power than currently available today with hypervisor technology. These demanding applications are the very ones that have defied consolidation to date. Nonetheless, the level of performance and scalability has been sufficient for most legacy and line of business applications that have been the targets of consolidation activities to date.

Table 1. Summary of Server Consolidation Technology Evolution

	Isolation	**CPU Utilization**	**QoS**
Servers with No Partitioning	**(-)No isolation** If consolidating multiple applications on a single server and OS image	**(-)Poor utilization** Without isolation, limited consolidation of multiple applications occurs	**(-)Poor QoS** Little to no throttling and prioritization of resources
Servers with Physical Server Partitioning	**(+)Strong isolation**	**(-)Poor utilization** Multiple applications are consolidated, but little to no improvement to CPU utilization	**(o)Good QoS** Some additional CPU and memory can be dedicated to a partition for applications that may need it
Server Virtualization / Hypervisor	**(+)Strong isolation**	**(+)Highly utilized**	**(o)Good QoS** Applications can be scheduled across up to four CPUs or cores, but scalability above 4 CPUs is capped

Parallel to Open Systems Storage Consolidation

We now turn our attention to how the storage consolidation market has evolved over the past decade and analyze the technologies instrumental in propelling storage consolidation in IT. As with server consolidation technologies, we will examine each phase of market evolution through the lens of isolation, utilization, and QoS. We believe that the storage consolidation market evolution can be divided into three distinct phases:

- Phase 1: Traditional SAN Attached Storage System
- Phase 2: High End Storage Array with Physical Storage Partitions
- Phase 3: Virtual Machines for Storage

PHASE 1: TRADITIONAL SAN ATTACHED STORAGE SYSTEM

Storage consolidation came to fore as a concept and economic driver concomitant with the widespread adoption of Storage Area Networks (SANs). Ultimately, storage consolidation allowed IT to sweep up pools of underutilized DAS storage and consolidate it onto a centrally managed SAN. The predominant driver for this storage infrastructure transformation was not just the capital savings of a more highly utilized storage infrastructure, but the benefits of faster, more reliable backup and centralized Disaster Recovery (DR) practices.

However, traditional SAN attached storage systems did not address the twin problems of workload isolation and QoS. For example, traditional SAN attached storage systems were not well suited for consolidating multiple I/O workloads onto a single storage system and providing QoS SLAs for each workload. As a result, many IT users opted to dedicate modular arrays to specific applications that demanded high performance. By dedicating hardware, IT could isolate I/O workloads and administration, and provide known, if not high, levels of QoS for a particular application or workload. However, this practice negatively impacted storage utilization rates on the SAN because it potentially created separate pools of underutilized capacity and performance. The end result: SANs ushered in an improvement in storage utilization over DAS, but it did not solve the vexing problems of isolation and QoS that were prevalent when multiple I/O workloads were consolidated onto a single physical storage system.

PHASE 2: PHYSICAL STORAGE PARTITIONS

Similarly to the evolution of server consolidation, the storage vendor community introduced physical storage partitions on high end SAN attached storage as a means to address the issues with isolation and QoS. Physical storage partitions allow an administrator to divide an array into multiple partitions that appear and are administered as a single discrete storage array. Administrators dedicate specific cores, controllers, cache, and disks to a particular partition. In short, physical storage partitions allowed administrators to create arrays within an array. Two examples of physical storage partitions are HDS Virtual Partition Manager and IBM DC8300 LPARs.

By creating separate partitions and dedicating specific physical resources within the array, physical storage partitions isolate I/O workloads from one another and provide a certain level of QoS and performance. Administrators can add additional physical resources (e.g. cores, cache, and disks) to improve or scale performance as necessary within the overall limits of the physical storage array.

However, dedicating physical resources to a particular partition mitigates the utilization benefits of a centrally, consolidated storage infrastructure. Physical resources dedicated to a particular partition cannot be used or shared by other partitions. Therefore, storage capacity or cache can be trapped and underutilized once dedicated to a particular partition.

PHASE 3: VIRTUAL MACHINES FOR STORAGE - "VIRTUAL ARRAYS"

Until very recently, physical storage partitioning represented the apex of the storage consolidation technology evolution. However, similar to how server consolidation technologies evolved, we are now witnessing the birth of a soft partitioning scheme akin to hypervisors. If we are correct, this new virtual array scheme should drive the next stage of growth in storage consolidation solutions in much the same way that it fueled the rise of server virtualization technologies.

This next wave of storage consolidation is premised on creating virtual arrays within a single physical storage system. As a result, the storage system provides logical segregation and partitioning, but still ensures that all physical resources are shared across all virtual arrays. The promise of virtual arrays within a single physical storage system is that multiple types of workloads can be consolidated onto a single physical storage system without affecting the performance and operation of the other workloads. Each of these workloads can enjoy isolation and be assigned different levels of QoS, while ensuring that all physical resources (controller, cache, disk types, etc) are fully utilized. In short, this next evolution is analogous to a virtual machine layer for storage.

Table 2. Summary of Storage Consolidation Technology Evolution

	Isolation	CPU Utilization	QoS
Traditional SAN Attached Storage	**(-)No isolation** if consolidating on single SAN attached array. Isolation achieved by dedicating modular arrays to a particular workload	**(-)Poor utilization** Arrays are too often dedicated to deliver workload isolation and QoS	**(-)Poor QoS** Somewhat mitigated if arrays are dedicated to a particular workload
High End Storage Array with Physical Storage Partitioning	**(+)Strong isolation**	**(-)Poor utilization** Dedicated resources to individual	**(o)Good QoS** Some additional controllers, cores,

		partitions leads to underutilization	cache, and disk types can be dedicated to a partition to improve performance
Virtual Array within an Array	**(+)Strong isolation**	**(+)High utilization** All resources can be shared by each and every application. Purchase on pooled need basis.	**(+)High QoS** All physical resources can serve each and every application workload

Enter 3PAR

Founded in 1999, 3PAR is a next generation storage system provider focused on improving resource utilization and simplifying management in departments and data centers. 3PAR Utility Storage is designed from the ground up to deliver a simple to use, efficient, massively scalable solution for the next generation data center and its mission critical applications. The product is ideally suited for storage consolidation, integrated data lifecycle management and performance intensive applications.

3PAR has eschewed the traditional dual controller designs of most traditional midrange and the costly layered architectures of high end storage systems in favor of a multi-node, massively scalable clustered design. 3PAR's Utility Storage is based on the innovative InSpire clustered architecture that provides a modular, fault tolerant storage platform that scales continuously from the very small to the very large. 3PAR InServ Storage Servers are composed of a cluster of up to 8 3PAR controller nodes interconnected together over a high speed, low latency meshed backplane. Together the 3PAR system forms a cache coherent, active-active cluster.

Each InServ Storage Server acts as a single system, so that hosts can access virtual volumes over any host-connected Fibre Channel or iSCSI port on any controller node, even if the particular portion of data to be accessed is directly managed by another controlled node. The modular, clustered architecture allows 3PAR to scale to meet the most demanding workloads, while delivering capital efficient "pay as you grow" scalability to meet changing business demands.

Spotlight on 3PAR Virtual Domains

In November 2007, 3PAR announced the general availability of a new storage virtualization capability dubbed 3PAR Virtual Domains. Virtual Domains allows administrators to define multiple discrete, secure virtual arrays within a single physical 3Par InServ system. In fact, 3PAR Virtual Domains supports up to 2,000 Virtual Domains in a single InServ Storage Server. From the user or host perspective, the Virtual Domain appears and is accessible as if it was a discrete, standalone physical storage system. A Virtual Domain may export volumes to exclusively defined WWN names, contains its own virtual pool of storage capacity and configuration policies, and its own defined administrators.

Unlike the use of multiple storage arrays or of physical storage partitioning schemes that rely on the physical segregation of resources, 3PAR Virtual Domains uses a logical implementation that preserves the benefits of distributing each application workload across all system resources (ports, processors, cache, loops, and disk drives). In short, a Virtual Domain is a fully isolated and secured storage system within a high performance, scalable storage platform.

Virtual Domains represents the next logical step in the storage consolidation evolution. It is analogous to a virtual machine for storage – a virtual array within a physical array.

As previously mentioned, the problem with storage consolidation approaches in the past is that they lacked one or more of the three key attributes (isolation, QoS, and higher utilization) that users need in a consolidation solution. These three attributes have been key inhibitors to large scale, widespread storage consolidation in the past. Virtual Domains is the first product of its kind to allow end users to consolidate onto a single physical storage system, while still providing the isolation, performance, and high levels of utilization that users demand from their storage infrastructure. In short, Virtual Domains eliminates one of the last technological and business reasons for maintaining physically diverse and dispersed storage devices.

BENEFIT #1: IMPROVED UTILIZATION WITHOUT COMPROMISE

With Virtual Domains, IT gains the utilization benefits of a consolidated storage environment, while getting the isolation and QoS benefits for individual I/O workloads. Virtual Domains does not make IT have to compromise or choose between two suboptimal choices – lower utilization or better QoS. IT can get highly utilized, shared storage infrastructure that delivers true isolation and QoS to each and every application.

BENEFIT #2: SECURITY, PRIVACY, & COMPLIANCE

A common impediment to storage consolidation has been that a particular department or business unit requires the control and privacy of its own dedicated storage array. With Virtual Domains, 3PAR allows departments, business units or external businesses the control and privacy they demand for their application data while enabling IT to centralize and consolidate storage onto a single scalable storage platform.

BENEFIT #3: HARNESSING PERFORMANCE OF THE ENTIRE SYSTEM

With Virtual Domains, each domain spans potentially eight controller nodes and hundreds of disk drives in a 3PAR InServ storage system. In this way, each Virtual Domain is automatically load balanced against all physical resources within the system, including controllers, cache, and drives. The result is: higher and more predictable levels of performance, greater resource utilization, and no need to manually load-balance Virtual Domains over silos of physical resources.

Virtual Domain Usage

Virtual Domains challenges some of the widely held assumptions around storage infrastructure configuration and management. To that end, it is worth discussing how Virtual Domains complement or enable new usage and administration models for storage infrastructure.

USAGE: SELF SERVICE STORAGE

Virtual Domains has the power to free storage administration from the repetitive mundane storage provisioning and management tasks to become a more strategic partner of managing and delivering storage SLAs for the business units and their critical applications. Coupled with the notable ease-of-use of InServ Storage Servers, Virtual Domains allows a storage administrator to delegate the day to day administration tasks of the storage infrastructure to the lines of business. For example, a test/dev Virtual Domain can be created with a defined provisioning policy and fixed pool of capacity allowing all application developers to create volumes and take snapshots easily without the need for central IT to be involved. This frees storage teams up to focus on the higher value work of ensuring service quality and achieving SLAs for the business units. Furthermore, the lines of business get what they want – greater control and faster response to provision, allocate, and manage their storage within the predefined constraints of a domain and without the expense of dedicated storage.

USAGE: MITIGATES OPERATOR ERROR

Virtual Domains provides a secure and compartmentalized environment for an administrator to provision and manage storage. Operator error remains the number one cause of downtime and lost data today. Virtual Domains safeguards administrators from inadvertently issuing a command that destroys data or irreparably alters storage configurations across the entire system. Once logged onto a Virtual Domain, an administrator can only change volumes and storage configuration within the domain. As a result, administrators are protected from fat fingering or mistyped commands that could have led to hours of downtime or lost productivity.

USAGE: SECURITY, COMPLIANCE, & PRIVACY

One of the last objections to centrally managed and consolidated storage infrastructure has been the need to segregate sensitive application data from other forms of data for compliance or security reasons. Virtual Domains solve this problem because it allows the sensitive data to be logically segregated and contained within a secure Virtual Domain. However, IT still enjoys all the economies of scale that centrally managed, consolidated storage enables because Virtual Domains ensures that the data and access to the data is distributed over all available physical resources (controllers, cache, and disks). In short, Virtual Domains obviates the need for lines of business to invest in dedicated, underutilized storage

87 Elm Street, Suite 900 Hopkinton, MA 01748 Tel: 508-435-5040 Fax: 508-435-1530 www.tanejagroup.com

arrays to satisfy security, compliance, or privacy business reasons.

Taneja Group Opinion

The parallels between the evolution of the server consolidation and storage consolidation markets are striking. The introduction of the hypervisor ignited the current market tornado around server virtualization and server consolidation. Virtual Domains has the same potential to drive a new wave of storage consolidation, allowing IT to consolidate previously impossible-to-consolidate data sets and application workloads. Moreover, Virtual Domains has the power to eliminate the common objections that lines of business raise against centrally managed, consolidated storage infrastructure.

From our vantage point, Virtual Domains represents an industry first and a pivotal milestone in driving greater storage consolidation throughout organizations' storage infrastructure. Looking forward, we believe that virtual array capabilities will become another hallmark of advanced virtualization solutions and should become a non-optional capability of any high end storage system that aspires to be a bonafide storage consolidation platform.

Throughout its history, 3PAR has always been a technological pioneer of advanced storage virtualization functionality. They were an initial proponent of key concepts like thin provisioning, autonomic management and clustered controller designs that have driven greater efficiency, reduced administration and higher performance in high end Fibre Channel storage. With the introduction of Virtual Domains, we see that 3PAR is once again flexing its innovation muscle. This is great for IT users. Finally, IT has the ability to consolidate storage infrastructure without compromise. We think that introduction of Virtual Domains cannot be overlooked and represents a significant step forward in the quest to consolidate and centrally manage all storage infrastructure.