TANEJA GROUP
TECHNOLOGY ANALYSTS

SOLUTION PROFILE

The Importance of Advanced Storage Virtualization in Virtualized Server Environments

June 2007

THE NET NET

Server virtualization has become a tried and true approach for enabling a new level of cost savings and operational efficiency in the data center. Enterprises everywhere are rapidly implementing virtualized servers in a quest for better resource utilization, simplified management, and a more flexible infrastructure. Over 85% of these virtualized servers use a SAN and a shared storage model. Unfortunately, server virtualization with shared storage exacerbates pre-existing storage management issues, and in many cases creates new challenges not present in a purely physical world.

We are finding that server virtualization planning and deployment are increasingly intertwined with the planning and scaling of a storage backend that can support the high levels of consolidation that many businesses demand today. As the scale of a virtual environment increases, the performance, availability, and cost-efficiency of the storage infrastructure becomes crucial to enabling the next level of scalability of a virtualized server environment.

In order to address these vexing storage infrastructure challenges, advanced storage virtualization features become a key infrastructure fulcrum for optimizing capacity and performance in a virtualized and consolidated server infrastructure. In this solution profile, we will review the key operational storage management challenges caused by server virtualization, posit the key criteria of advanced storage virtualization capabilities necessary to optimize and consolidate server infrastructure, and then spotlight 3PAR and how their advanced virtualization capabilities stack up against the criteria.

Key Storage Issues in Virtualized Server Infrastructure

There is a symbiotic relationship between planning and managing a virtualized server infrastructure and how to deploy, manage, and optimize the storage environment. From our conversations with end users, we frequently hear three major storage management and infrastructure challenges that are cited by users of large scale virtualization environments. Although these challenges are not necessarily unique to server virtualization, these challenges are exacerbated by the use of server virtualization and require careful planning and consideration before attempting a large scale server consolidation or virtualization project.

CHALLENGE #1: CAPACITY MANAGEMENT

Storage capacity planning and management are key disciplines and responsibilities of the

Copyright© The TANEJA Group, Inc. 2003. All Rights Reserved

storage administrator whether the environment is purely physical or virtualized. Ironically, while server virtualization is an excellent mechanism for driving server utilization, it can have the opposite effect on storage utilization rates.

Server virtualization, specifically VMware, favor large upfront allocation of capacity to a volume, even if little to none of that capacity is written or used at the outset. VMware's Virtual Machine sizing guidelines advise users to *double* application useable capacity estimates to account for additional suspend/resume space that VMware may require.

As a result, it is not uncommon that end users must allocate 1 to 2 TBs of physical storage capacity per volume per Virtual Machine (VM). In a highly consolidated environment, storage administrators must allocate 10s of TBs of capacity just to accommodate the deployment of a relatively small number of VMs. Therefore, over allocation of storage capacity results in low storage utilization and high CAPEX spend, mitigating the strong benefits that a virtualized server infrastructure brings.

CHALLENGE #2: APPLICATION PERFORMANCE

Application performance and responsiveness are always top of mind goals for IT. Users often use CPU utilization rates as a broad measure to determine how many workloads can be consolidated onto a given physical server. However, this oversimplifies consolidated workload performance management and optimization. From our experience, memory utilization and I/O latency must also be considered in optimizing application performance in a virtualized infrastructure.

Memory utilization at both the server and storage cache level, along with overall storage system responsiveness are key determinants of overall system level performance. For example, a typical server in a data center today will utilize over 60% of its memory, while its CPU will use between 5 to 15% of the available CPU cycles. When multiple physical servers with a similar memory and CPU utilization profile are consolidated onto a single server, memory becomes overcommitted and over utilized. Excessive memory utilization results in large number of page and cache misses. Each of these cache misses is expensive because VMware must make a request to the underlying disk system to recall the page back into memory. This in turn causes a highly randomized I/O pattern per LUN volume at the storage system, resulting in additional cache misses at the storage level too. More roundtrips to disk mean longer waits and lower application performance.

Consolidation best practices call for end users to provision servers with as much RAM as possible to counteract this excessive page swapping. However, modern day x86 servers were not designed to support the amount of RAM necessary for highly consolidated environments of greater than 6 VMs per server.

Overall application performance becomes tightly correlated with the overall performance and response time of the

Copyright© The TANEJA Group, Inc. 2003. All Rights Reserved

underlying storage system. High performance storage where I/O can be striped over large numbers of disks to deliver low latency response times becomes a crucial element enabling high scale virtualized server deployment.

CHALLENGE #3: DATA PROTECTION OF VIRTUAL MACHINES

Server virtualization environments require highly available shared storage since many virtual machines (VMs) share a common server. This magnifies the business downtime impacts of a failure since "all the eggs are in a single basket." VM data structures must be not only protected, but also easily shared so that VMs can be rapidly failed over or redeployed to another server platform to maximize uptime.

Server virtualization also complicates the data backup practices. Off-hours backups can no longer run simultaneously because the performance demand on a single server supporting multiple virtual guests is simply too high. In response, storage managers look to array-based snapshot backups, specialized clients, mirror splits, or other approaches, but those possibilities are mired in management complexity, negative performance impacts, and possible up-time complications. Assuring that virtual guests are easily backed up with complete data integrity becomes nearly impossible. Meanwhile, mapping constantly changing virtual configurations to typically scripted and fixed back-up tasks, introduces a new challenge.

Key Elements of Advanced Virtualization

Forward looking enterprises have found that pairing advanced storage virtualization capabilities with server virtualization offers a solution to these storage management challenges. In the following section, we will discuss the key capabilities that any storage virtualization (whether it is network-based or array-based) must have in order to deliver the business benefit of "truly" optimizing virtualized server infrastructure and enabling large scale server consolidation.

Requirement: Thin Provisioning

Thin provisioning is a crucial advanced storage virtualization capability that ensures that over allocation and low utilization of the corresponding storage infrastructure does not occur. As previously mentioned, VMware requires large volumes to be pre-allocated for each virtual machine. Thin provisioning counteracts this over allocation problem by ensuring that only storage that is used or written is allocated. In short, thin provisioning presents a logical volume of any size to VMware, but only allocates physical storage capacity when data is written or used on the volume. The physical storage is allocated on demand from a shared pool. As a result, storage capacity is used much more efficiency and only when needed. As a result, an organization can achieve not only optimal server utilization rates, but also high storage utilization rates by using a combination of server virtualization and thin provisioning.

Requirement: Performance

The ability of the storage system to process random I/O with low latency is a key

Copyright© The TANEJA Group, Inc. 2003. All Rights Reserved

determinant in enabling large scale consolidation. With server consolidation, IT is effectively collapsing many types of workloads and I/O access patterns onto a single physical server, creating an amalgamated server I/O access profile. This workload mix is compounded by frequent page swapping to disk due to high memory utilization that comes with a highly consolidated, virtualized server infrastructure. As a result, there is a strong dependency on fast, highly responsive storage.

Ideally, users need to deploy storage systems that can efficiently process small block, random I/O with low latency. From our experience, we have found that advanced storage virtualization functionality like clustered controller-based architectures and wide striping – the ability to balance data across as many available spindles as possible to get the best responsiveness and I/O throughput – are crucial to supporting the service levels required in a virtualized storage infrastructure.

Requirement: Simplicity of Operation

Simplicity of operation is a term that Taneja Group uses to describe the overall ease of management of a given storage virtualization solution. Simplicity of operation refers to not only the ease of initial installation and configuration, but how easily and efficiently an administrator can perform routine maintenance and key storage management tasks, such as storage provisioning.

In a virtualized server environment, end users are constantly provisioning new virtual machines, patching existing virtual machine operating environments, and upgrading the underlying physical server infrastructure. As a result, rapid provisioning, patch management, and online, non-disruptive data migration become critical operations that the storage administrator must perform to support the virtualized server environment. Specifically, the administrator must be able to optimize the performance of storage to meet the requirements of virtual servers, scale the storage environment at the same rate as the virtual server environment is scaled, and conduct routine storage maintenance activities – all without disrupting the virtual server environment. In total, these storage capabilities are crucial to contribution to a more agile, flexible IT infrastructure. We strongly urge users to consider how they will provision new virtual operating environments, patch them, and non-disruptively migrate data as part of their server virtualization implementation.

Requirement: Data Protection & Fault Tolerance

By consolidating, enterprises are effectively putting "more eggs in one basket." As a result, fault tolerance and data protection become paramount issues to ensure IT achieves their uptime SLAs for their virtualized environment. First, the fault tolerance and overall availability of the underlying physical storage is crucial for a supporting a virtualized environment. An advanced storage virtualization solution must have no single point of failure and be capable of surviving cascading failure scenarios, while ensuring that data is fully protected from physical disk failures.

Copyright© The TANEJA Group, Inc. 2003. All Rights Reserved

Second, an advanced storage virtualization solution should provide advanced snapshot and data copy capabilities that allow enterprises to create duplicate data sets for backup and recovery. This facilitates backup, diminishes the performance impact of I/O intensive backup activity on primary storage, and provides the ability to quickly restore the data set to a previously consistent point in time copy. Furthermore, an advanced virtualization solution should automate the creation, movement, and retention of snapshots and copies, reducing the administration for these tasks.

Requirement: Quality of Service (QoS)

With server virtualization and advanced capabilities, such as VMotion and Distributed Resource Scheduler (DRS), enterprises are beginning to gain greater control and flexibility in terms of setting and maintaining service levels for their virtual machines and applications. Given that storage IO performance is a key determinant of application performance in a virtualized server environment, QoS at the storage level becomes a key requirement for guaranteeing low latency, high throughput I/O to the application and virtual machines that need it most.

Traditionally, upgrading data service levels involved manual migrations of data among arrays or storage tiers or lengthy RAID set rebuilds. These migrations and volume layout changes impacted performance and could result in significant downtime, while the migration or rebuild occurred. However, in a virtualized context, data migration and data service levels must be able to respond as quickly as the server virtualized environment to changing demands and workloads. As a result, the ability to dynamically change service levels on the fly and non-disruptively becomes an important complement to the flexibility of a virtualized server infrastructure. Specifically, advanced virtualization solutions that allow administrators to change RAID settings, promote and demote data among different tiers or classes of storage, add additional disk resources to a given volume, and adjust the placement or stripe-width of data on disk for optimum performance are critical capabilities for delivering dynamic QoS at the storage layer.

Spotlight on 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

Copyright© The TANEJA Group, Inc. 2003. All Rights Reserved

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.

3PAR Compared Against the Criteria of Advanced Virtualization

In addition to a highly scalable, available clustered architecture, 3PAR Utility Storage boasts a wide range of advanced storage virtualization capabilities within the storage system. Unlike other storage systems where advanced storage virtualization were bolt-ons or retrofits to a legacy design, 3PAR had the luxury of designing key capabilities like thin provisioning, clustered volume management, and snapshots at the core of its architecture. The result: an elegant execution of advanced internal virtualization capabilities married with a massively scalable clustered architecture. In the following section, we will examine in greater depth the key storage virtualization capabilities of the 3PAR product and how they match up to the needs of an enterprise considering a large scale virtual server consolidation project.

3PAR: Thin Provisioning

3PAR recognized early on the problems caused by over allocating storage capacity to volumes and responded with a virtualized storage system capable of making more efficient use of physical storage capacity than traditional fat provisioning approaches. Since 2003, the company has offered thin provisioning as an option on its 3PAR Utility Storage. In contrast, it is only recently that large tier 1 vendors have added or announced plans to add bolt-on thin provisioning approaches to their arrays.

The 3PAR system presents to hosts virtualized volumes of any size and service level. Meanwhile, physical storage capacity is intelligently and autonomically dedicated only as the application writes data. 3PAR Thin Provisioning is "Dedicate-on-Write," as opposed to "Dedicate–on-Allocation." Physical capacity can be added to the InServ's single unused capacity pool non-disruptively at any time. Given the fact that over allocation of storage is such a rampant problem in VMware environments, 3PAR's thin provisioning implementation offers an ideal antidote to this costly problem.

3PAR: Performance

3PAR's massively parallel InSpire architecture automatically spreads all data over all internal disk resources, delivering demonstrably higher and more predictable levels of performance (high IOPS, low latencies) with full capacity utilization. For example, an eight node 3PAR clustered system delivers 100,046 IOPS in the SPC-1

Copyright© The TANEJA Group, Inc. 2003. All Rights Reserved

benchmark at $14.81 per IOPS. Unlike many other SPC-1 results, 3PAR achieves their SPC-1 result with a nearly complete, utilization of installed capacity. The SPC-1 benchmark represents a good proxy for real world business critical applications, such as those being consolidated with VMware.

Due to page swapping and high levels of memory utilization, VMware's performance is highly dependent on fast, low latency storage performance. 3PAR has shown that its architecture is a strong fit with the I/O profile of a virtualized server infrastructure. We believe that 3PAR's advanced storage virtualization capabilities like wide striping and highly parallelized clustered design are critical capabilities in delivering low latency, high IOPS performance to virtualized applications. With 3PAR's ability to deliver low latency, high throughput IO to VMWare, end users have found that a 3PAR storage system can enable them to stack more VMs per server without impacting application performance. As a result of 3PAR's architecture and internal virtualization capabilities, enterprises can realize greater efficiency and higher levels of VMware consolidation than traditional storage systems with external network-based or appliance-based virtualization.

3PAR: Simplicity of Operation

3PAR provides a single point of management for administering the cluster system as a single physical resource. Both UI and CLI have been designed to minimize administrator complexity and training requirements.

Furthermore, 3PAR has invested heavily in streamlining the provisioning and patch management processes for their storage system. A 3PAR administrator can configure one or hundred volumes and export it to hosts in a matter of seconds. Moreover, provisioning can be automated to utilize templates that stipulate preset volume parameters for desired service levels. The net effect is that the provisioning process can be compressed from days or weeks in many organizations to seconds and minutes using 3PAR technology. Often times, provisioning is a gate to the rapid deployment and re-deployment of virtual machine. With 3PAR, administrators get an easy to administer storage backbone that can respond rapidly to the dynamic, flexible nature of a virtualized server environment.

3PAR: Fault Tolerance & Data Protection

3PAR provides a comprehensive approach to ensuring that data is accessible and protected at all times. 3PAR's clustered architecture provides a completely fault-tolerant system that can survive multiple individual component failures, while continuing to service requests for data. In addition, the 3PAR system features rapid rebuild times to safeguard the system from disk failures and non-disruptive hardware and software upgrades to minimize or eliminate downtime from outages.

In addition to 3PAR's redundant architecture, 3PAR boasts advanced data protection capabilities designed around 3PAR Virtual Copy – its highly efficient and flexible snapshot implementation built on thin copy technology. Virtual Copy allows an

Copyright© The TANEJA Group, Inc. 2003. All Rights Reserved

administrator to create hundreds of point in time copies of a single volume, while consuming minimal amounts of physical storage. More importantly, Virtual Copy's fine-grained auto-growth capability is drawn in small increments from a free space pool on an as needed basis, eliminating the traditional need to reserve conservative amounts of copy space on a per-volume basis. Furthermore, users can assign distinct quality-of-service levels to this copy space apart from those of the base volumes. In short, 3PAR's Virtual Copy capability provides an excellent data protection capability for supporting the backup and recovery of virtual machine images and their data without disruption to the application service quality.

3PAR: QoS

3PAR Dynamic Optimization technology provides the facility to online and non-disruptively change data service levels on the fly, allowing the 3PAR storage system to respond in concert with server virtualization software to changing application and business demand. With a single command, 3PAR's Dynamic Optimization engine can reliably convert a data volume from one service level to another with no disruption to the application. 3PAR's system acts as the perfect complement to VMotion and DRS to facilitate an agile, flexible virtualized infrastructure that can support varying degrees of service quality on demand.

Dynamic Optimization allows operators to change RAID type, restrict or dedicate the use of individual processors, drives, and or ports to a specific application or virtual machine, place data on the outer tracks for optimum performance and response time, and decide what data should be placed on FC or SATA class drives. To that end, 3PAR offers a simple, straight forward approach to delivering high service quality to a dynamic, rapidly changing virtualized server environment. Given the importance of storage system responsiveness in virtualized environments, 3PAR's QoS capabilities are critical for high performance virtualized server environments.

Copyright© The TANEJA Group, Inc. 2003. All Rights Reserved

Table 1. Summary of 3PAR Capabilities Compared to the Storage Virtualization Criteria Needed for a Highly Virtualized Server Infrastructure.

Storage Virtualization Requirements	3PAR Findings
1. Thin Provisioning	3PAR provides an elegant implementation of thin provisioning that an organization to achieve high server utilization rates without compromising storage utilization.
2. High IOPS, Low Latency Storage Performance	3PAR allows data to be striped across all available spindles and controllers in order to deliver low latency, high IOPS performance. As a result, more virtual machines can be consolidated on a single server, resulting in higher utilization and cost savings.
3. Simplicity of Operation	3PAR streamlines common storage maintenance operations, such as provisioning and patch management, allowing the storage system to be scaled in tandem with the virtualized server infrastructure.
4. Fault Tolerance & Data Protection	3PAR clustered architecture and advanced virtualization capabilities, like Virtual Copy, provide a comprehensive approach to ensuring that virtual machine data is accessible and protected at all times.
5. QoS	3PAR Dynamic Optimization technology ensures storage I/O responsiveness in virtualized environments by providing the capability to quickly, non-disruptively dial up or down data service quality by volume.

Taneja Group Opinion

Virtualized server environments continue to proliferate as businesses look to drive greater operating efficiencies across the data center. Many forward looking enterprises have goals to virtualize over 50% of their server infrastructure in the next couple of years. Clearly, server virtualization is reshaping how applications are deployed and managed in the data center. However, this sea change is not without ramifications for the storage infrastructure. In fact, we have come to believe that the planning and deployment of large consolidated, virtualized server environments is explicitly intertwined with the design and implementation of a flexible, high performance, resilient storage infrastructure.

Our advice to end users: think holistically about your infrastructure before embarking on a large scale server consolidation project. As we have discussed, server consolidation is about improving server utilization, but it

Copyright© The TANEJA Group, Inc. 2003. All Rights Reserved

makes no sense to drive up server utilization to the detriment of lowering storage utilization. In short, we believe that user must consider both sides of the infrastructure equation – server and storage optimization – in order to realize the highest ROI possible inn a server consolidation project.

To that end, selecting a storage backend with rich internal storage virtualization capabilities is critical to the successful rollout of a large-scale consolidated, virtualized server infrastructure. From our research and discussions with end users, selecting the right storage system with internal storage virtualization allows users to stack more virtual machines per server and hence drive better operating efficiencies across their environment. Furthermore, the right advanced virtualization solution will require a fraction of the storage capacity otherwise needed. Lastly, the right storage virtualization solution complements the flexibility and management benefits of a virtualized server environment by streamlining and automating storage provisioning on demand. In short, we believe that advanced internal storage virtualization should be an integral element of any large scale server consolidation implementation.

3PAR, with its innovative high scalability, clustered architecture and rich storage management capabilities, has proven itself to be a leading vendor of advanced storage virtualization technology. If you are considering a large scale server consolidation rollout, we strongly urge you to consider 3PAR for your storage requirements. We think that the combination of efficiency, simplicity, and scalability that 3PAR brings to bear are the perfect complement to a virtualized server environment.

NOTICE: The information and product recommendations made by the TANEJA GROUP are based upon public information and sources and may also include personal opinions both of the TANEJA GROUP and others, all of which we believe to be accurate and reliable. However, as market conditions change and not within our control, the information and recommendations are made without warranty of any kind. All product names used and mentioned herein are the trademarks of their respective owners. The TANEJA GROUP, Inc. assumes no responsibility or liability for any damages whatsoever (including incidental, consequential or otherwise), caused by your use of, or reliance upon, the information and recommendations presented herein, nor for any inadvertent errors which may appear in this document.

Copyright© The TANEJA Group, Inc. 2003. All Rights Reserved