

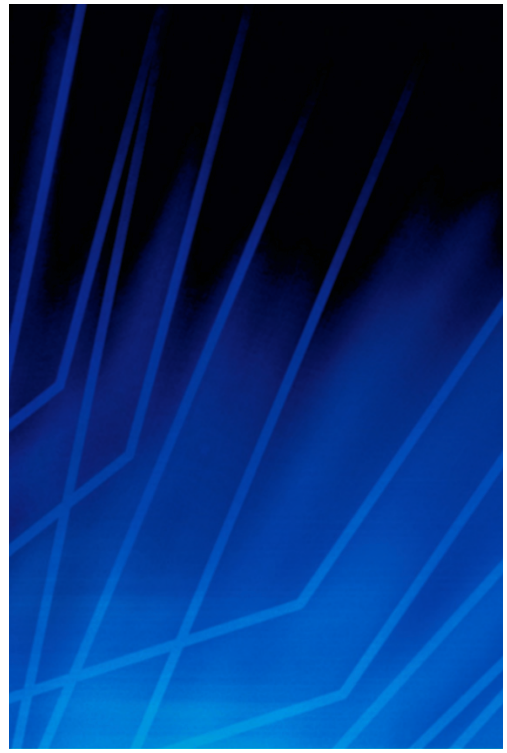

TECHNICAL PAPER
March 2009

Datacenter Virtualization with Windows Server 2008 Hyper-V Technology and 3PAR Utility Storage

Microsoft | Virtualization

3PAR
Serving Information

Table of Contents

1 INTRODUCTION

Hyper-V™ is an optionally-enabled, hypervisor-based server virtualization feature of Windows Server® 2008 that increases server infrastructure flexibility while saving time and reducing costs. 3PAR InServ® Storage Servers complement Hyper-V by providing a virtualized storage platform to support Hyper-V virtualized servers. The resulting combination is a powerful virtualized compute platform that helps enterprises and service providers reduce both capital and operational expenses while simultaneously increasing the overall flexibility of their datacenters.

This paper introduces 3PAR® Utility Storage and then details its benefits for a virtualized server environment with Hyper-V. Laboratory test results provided by 3PAR are used to illustrate and quantify these benefits; these results have not been independently verified.

1.1 Overview of 3PAR Utility Storage

3PAR Utility Storage is a highly-virtualized, tightly-clustered, dynamically-tiered storage array built for utility computing. 3PAR InServ Storage Servers use massive parallelization of all available controller nodes, physical disks, and other system resources to widely distribute volumes for optimal performance and resource utilization. The 3PAR InForm® Operating System, the software that manages 3PAR InServ arrays, automates the process of provisioning these massively-striped and scalable volumes, known as Virtual Volumes (VVs). The combination of ease-of-use and automated provisioning within 3PAR Utility Storage yields a highly-scalable, extremely adaptable, and easily manageable storage solution.

1.2 The Benefits of Virtualization

3PAR InServ Storage Servers and Hyper-V offer a unique alignment of benefits. Each product uses virtualization to sever the relationship between usable resources and the physical hardware behind those resources. By breaking these ties, previously stranded capacity and CPU cycles associated with individual applications are made available as a shared resource that can be used by additional applications. The resulting reduction in physical hardware and the associated management costs yields significant cost savings. The elimination of hardware also means less space and lower power consumption, leading to a greener datacenter and lower power, housing, and cooling costs. Additionally, virtualized environments enable greater flexibility and reduced complexity in provisioning new resources for applications. In short, virtualization enables IT organizations to reduce costs and to improve agility so that they can respond more quickly to new business demands.

2 STORAGE CHALLENGES FOR SCALING A HYPER-V ENVIRONMENT

While the advantages of server virtualization are well documented, the architecture of traditional storage poses numerous challenges to IT administrators trying to maximize the benefits of server virtualization. Combining multiple application servers onto one physical server results in higher-volume, less-predictable I/O. Traditional models of allocating storage by adding disks to each physical server result in extreme overprovisioning when spread over hundreds of virtualized application servers. Changing service-level requirements poses another challenge. While it is generally easy and non-disruptive to adjust some aspects of a virtual machine (VM), traditional storage arrays typically require a service disruption for data migration and LUN reconfiguration. This means that altering underlying volume characteristics such as RAID level or disk type (for example, moving data from Fibre Channel to SATA) are difficult and potentially disruptive. This section details these limitations of traditional storage and discusses how the limitations can prevent organizations from optimizing their server virtualization deployments.

2.1 Server Density Strains Storage I/O

A dedicated application server tends to have a predictable storage I/O load. Some applications, for example a transactional database, generate a random I/O pattern that is confined to predictable parts of the storage. Other applications primarily generate sequential I/O, for example a database's journal files or the individual files of a Web server. These I/O loads tend to use very little of the available storage bandwidth because growth and surge capacity is built into the servers.

When combined in a virtualized environment, the modest and predictable loads of all of the child (guest) VMs blur into a high-volume, very random storage I/O load that can strain the abilities of traditional storage arrays. Large caches on the array may help with moderate I/O loads, but as the load climbs it can exceed the abilities of the cache and the array becomes spindle-bound, at least on selected "hot" disks. Streaming applications require raw disk performance equal to at least the streaming throughput, and they derive little benefit from caching. As a result, traditional storage arrays are ill-equipped to deliver cost-effective high levels of performance for modern virtual server technologies.

2.2 Storage Overallocation Leads to Overprovisioned Storage

Windows Server 2008 Hyper-V technology makes it easy to create new virtual servers; however, traditional storage arrays lack a similar agility to easily change storage allocations as requirements change. As a result, server administrators routinely request large amounts of storage upfront so that it is readily available when new VMs need to be provisioned. Storage administrators, seeking to avoid future pain and expense, often further pad the storage allocation. These practices lead to poor storage utilization rates that can average as low as 25% according to industry studies[1] and that yield significant cost inefficiencies.

Performance considerations may further exacerbate the problem of overprovisioning. Traditional arrays do not widely stripe volumes. The heavy, random I/O load produced by an aggregation of many virtual servers is usually solved by provisioning additional physical disks to the traditional storage array to increase performance, often resulting in more disks being purchased and allocated than are required to meet the capacity required.

2.3 Lack of Agility Limits Changes to Storage Service Levels

Service level requirements for an application change with the passage of time. Hot projects requiring RAID 1 to eke out every bit of storage performance may evolve and, after completion, may only have an ongoing need for inexpensive archival storage to service occasional requests. Alternately, a few pilot projects may blossom into critical, high-demand applications, while other pilot projects continue with minimal activity. Seasonal business may also drive periodic changes in performance needs. Server virtualization makes it easy to adapt application servers to these changing needs, but

1 Stephen Foskett: *Real-world storage utilization*, GlassHouse Technologies, in Storage, April 2003.

traditional storage solutions lack the same agility. For example, traditional arrays do not allow administrators to easily modify the RAID level on a volume. Once a RAID level and type of disk (for example: fast Fibre Channel or more cost-efficient SATA-based Nearline disks) have been chosen, any change requires significant effort and often a disruption in service.

3 SOLUTION: HYPER-V WITH 3PAR UTILITY STORAGE

3PAR Utility Storage uniquely addresses these challenges associated with traditional storage systems and greatly enhances Hyper-V environments. Massive parallelization of all system resources and wide striping of data across all disks ensures that the performance of all spindles is available for every volume. As a result, parent servers hosting many child VMs do not suffer from poor storage performance. Simple, flexible provisioning of storage combined with 3PAR Thin Provisioning software, which allocates fine-grained physical storage "chunklets" on an as-needed basis, solves the problem of overprovisioning. 3PAR Dynamic Optimization software provides the agility to adapt services levels simply and transparently.

3.1 Scalable Performance with 3PAR

The massively parallel 3PAR InSpire® architecture[2] distributes volumes over all disks in a system, allowing 3PAR to deliver consistently high I/O performance even as demand changes and grows. Traditional storage arrays require careful planning to avoid spindle-bound hot spots, which hinders the flexibility to adapt to changes at the server virtualization layer.

The scalability coupled with the high performance of 3PAR Utility Storage means more applications can be consolidated on each Hyper-V physical server without fear of storage I/O bottlenecks. High performance also allows larger volumes to be used, simplifying the job of Hyper-V administrators. By spreading I/O across all disks, customers can achieve a required performance level with fewer disks on a 3PAR Utility Storage array versus traditional arrays, resulting in lower storage costs for their virtualized computing environment. The lower costs allow IT investment to be directed towards other areas.

A lab test was conducted to demonstrate the value of spreading volumes over many physical disks, a practice commonly referred to as wide striping. This test consisted of a random-write workload running in a Hyper-V child that ran against a volume spread over 32 disks. The number of disks was then increased to 80 (2.5 times) and the test was run a second time. The 2.5 times increase in disks resulted in a 2.4 times improvement in performance. The details of this test are located in Section 3.3.

3.2 3PAR Thin Provisioning

As previously mentioned regarding traditional storage environments, server administrators often request that storage be overprovisioned to accommodate anticipated future growth of an application's storage needs and to avoid requesting additional storage too frequently from the storage administrator. Additionally, having storage readily available allows server administrators to rapidly respond to new application requests. On the other hand, storage administrators are measured on minimizing the overallocation of storage. Because overallocated storage may remain

2 Refer to the *3PAR InSpire Architecture Technical White Paper* for further details.

unused for years[3], the storage team typically prefers not to allocate storage until it is actually needed. As a result, server and storage administrators are often at odds with one another regarding how to optimize the storage that is available.

3PAR Thin Provisioning solves this organizational conflict by allowing virtual volume capacity to be allocated without actually committing physical storage. Physical storage capacity is allocated on a fine-grained basis only when data is actually written, as illustrated in Figure 1. Instead of the traditional approach of tying storage costs to allocation, 3PAR Thin Provisioning aligns storage costs with the amount of data that is actually being written. As a result, server and storage administrators are both happy. The storage administrator can allocate a thin volume to the server team. The server team can then readily deploy new virtual servers and only consume the storage as data is written.

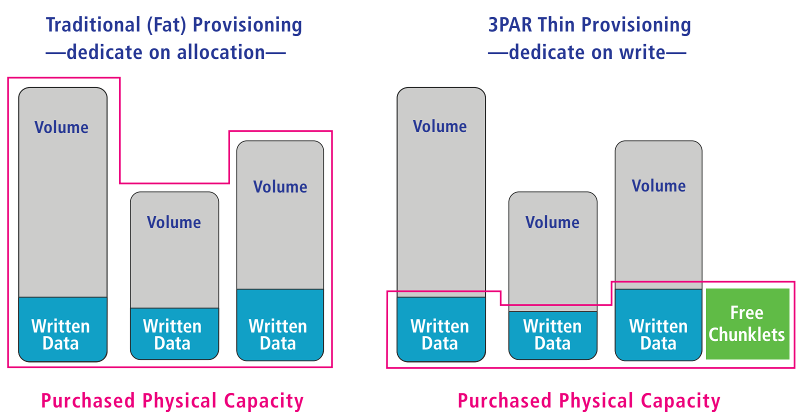

Fig. 01 Traditional Provisioning Versus 3PAR Thin Provisioning

3PAR Thin Provisioning also reduces IT costs by eliminating the need to carefully plan storage needs in order to minimize overprovisioning. A generous amount of storage may be provisioned and the application can grow into that space at its own pace.

Hyper-V greatly simplifies the process of bringing up new servers (and tearing them down) as needs change. However, allocating physical storage for the dozens of virtualized servers that might exist within a single Hyper-V server is expensive and wasteful. 3PAR Thin Provisioning addresses this paradox by deferring the allocation of physical storage until data is actually written (Figure 2), which restores agility to the entire virtualized server environment.

NTFS and VHDs (virtual hard drives) are Microsoft® technologies that are "thin-friendly", meaning they reuse "garbage" blocks instead of requesting additional storage from allocated but previously unused space. To evaluate how "thin-friendly" a Hyper-V environment is, Windows Server 2008 (64 bit) was installed on a server with a Fibre Channel HBA and then the Hyper-V role was added to it. A 40-GB Thin Provisioned Virtual Volume (TPVV) was provisioned on a 3PAR InServ Storage Server with an associated LUN exported to the Windows® server. An NTFS file system

3 Case Study: *Evaluating the virtualization and thin provisioning benefits achieved by 3PAR customers*, Wikibon, 2008.

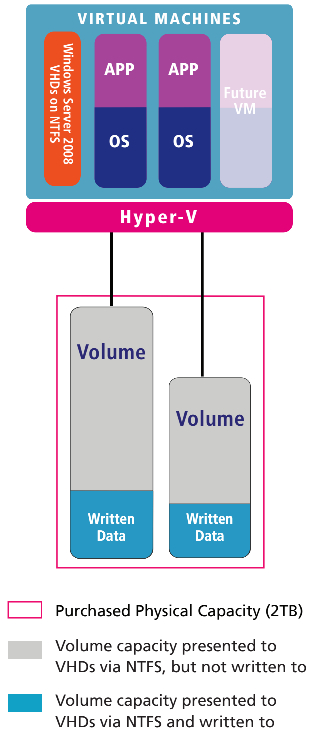

Fig. 02a Traditional Provisioning

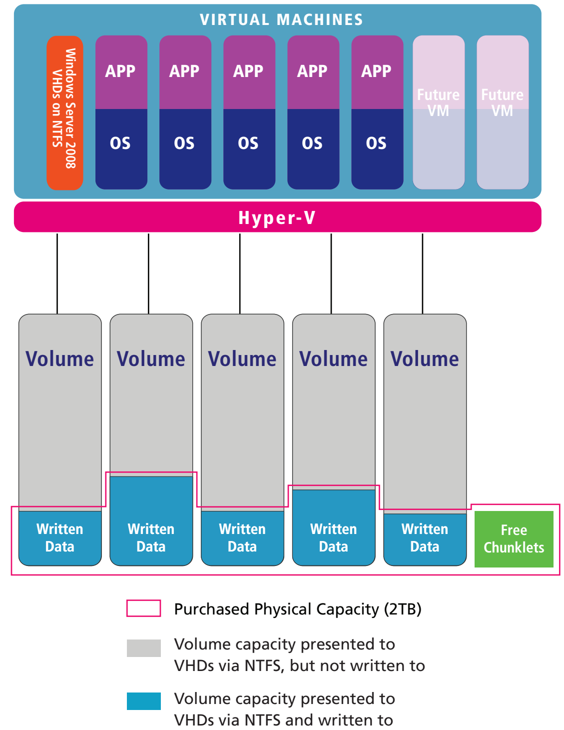

Fig. 02b Thin Provisioning

was then installed on this LUN using a quick format. Next, a Hyper-V child running Windows Server 2008 was installed using a VHD in the new file system. Storage usage was checked and then several copies of a large (646-MB) file were copied to the desktop of the child server. These files were deleted and copied again over several iterations. The results are illustrated in Figure 3 and summarized in Table 1.

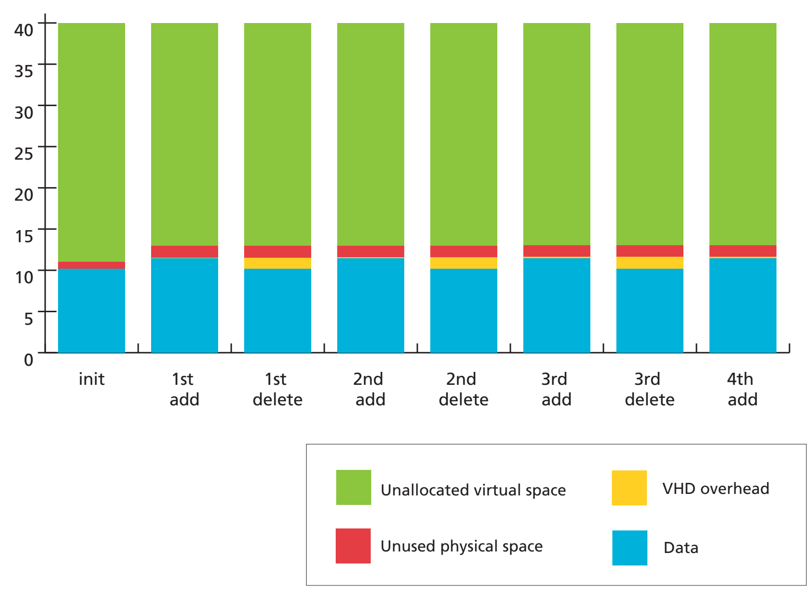

Fig. 03 Storage Usage as Files Are Added and Deleted

Iteration	TPVV Virtual size (GB)	TPVV Physical size (GB)	% Allocated	VHD Size (KB)
(initial state)	40	11.000	28%	10,171,990
1	40	13.000	33%	11,483,990
2	40	13.000	33%	11,550,615
3	40	13.000	33%	11,601,865
4	40	13.000	33%	11,601,865

Table 01 Storage Growth Over Time

After the first copied files were deleted, a second set of copies caused the VHD (and the TPVV in which it resided) to grow slightly (2%), but by the fourth iteration of the test, the VHD remained the same size. After reaching this state, a continued ebb and flow of similar amounts of data would be expected to result in no change to the VHD size, and no further allocation of physical storage.

The above example demonstrates that, at the end of the tests, only 33% of the 40 GB virtual volume has been consumed by the Hyper-V environment. In other words, only 33% of the possible 40 GB volume has been physically allocated. The remaining 67% of physical storage has not been allocated. In contrast, a traditional storage array would have allocated all 40 GB of physical storage, even though only 16 GB was actually required by these tests.

Note that the effectiveness of thin provisioning depends on the behavior of the application. A situation where a database pre-extends all of its files by writing zeroes to them or where a user immediately fills up all the expected space (for instance, through a backup to disk) will nullify the benefits of thin provisioning. Many applications are "thin-friendly" because their storage use grows gradually over time. As long as the application's file system is thin-friendly, the application can take advantage of thin provisioning. Since Hyper-V uses the Microsoft thin-friendly NTFS file system, Hyper-V and 3PAR Thin Provisioning complement each other nicely.

3.3 3PAR Dynamic Optimization

Hyper-V Live Migration allows server administrators to tailor application server performance by transparently migrating running VMs from one parent (host) server to another. 3PAR Dynamic Optimization provides a similar ability for storage, allowing the administrator to convert a volume to a different service level (Figure 4) with a single command. This conversion happens within the 3PAR InServ Storage Server, transparently and without interruption.

The agility of 3PAR Dynamic Optimization makes it easy to alter storage decisions. For example, a database for an active project might be allocated in RAID 1 (mirrored) storage on high-performance Fibre Channel disks. Later, when the project has transitioned to a maintenance mode, 3PAR Dynamic Optimization makes it simple to transparently convert the database volume to more cost-effective RAID 5 storage on more economical Nearline disks. Conversions could even be done to accommodate seasonal performance requirements, using RAID 1 for the busiest seasons and RAID 5 during off-peak times.

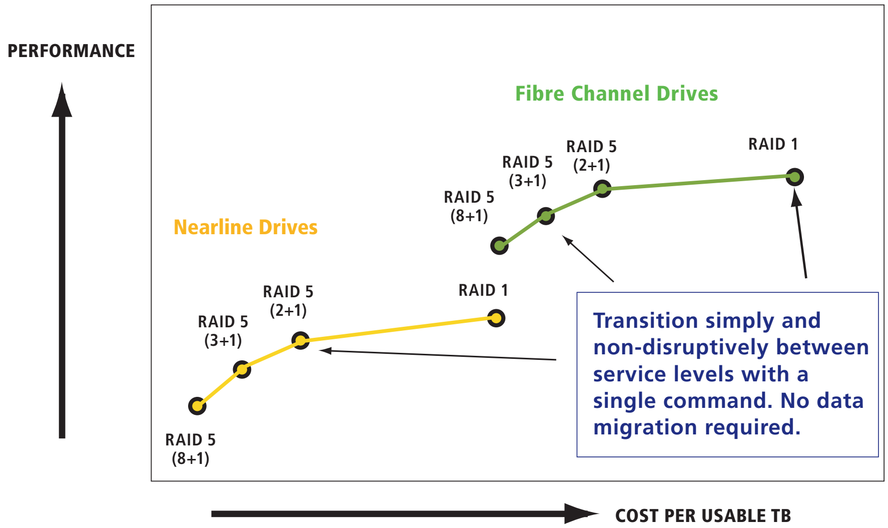

Fig. 04 Dynamic Optimization

Another use of 3PAR Dynamic Optimization is to redistribute volumes after adding disk drives to a 3PAR Utility Storage array. Using Dynamic Optimization, existing volumes are automatically striped across existing and new drives for optimal volume performance following capacity expansions. The increase in the total disks for the provisioned volume contributes to higher performance, as illustrated in the example given in Section 3.1, which describes how a 2.5 times capacity expansion yielded a performance increase of 2.4 times.

To demonstrate how 3PAR Dynamic Optimization can enhance performance after adding additional disks to a system, a 64 GB volume was created on the test system described in Section 3.2 using 32 disks, eight in each of four drive chassis, in a RAID 5 (3+1) configuration. An NTFS file system was created on the volume and a 60-GB VHD was allocated to the Hyper-V child. Within the child, Iometer was used to generate an I/O load consisting of 100% random writes of 512 bytes per request, an unusually challenging workload. An average throughput of 2,516 I/Os per second (IOPS) was reported during this initial test.

After establishing the baseline performance, 3PAR Dynamic Optimization was used to non-disruptively expand the volume from 32 disks to 80 disks in the same four drive cages. This expansion simulates starting off with partially-populated drive chassis and later adding additional disks as storage requirements grow. A 3PAR InForm OS command similar to the following was used to reconfigure the volume's storage:

```
tunealdvv usr_cpg dynopt_raid5 dynopt_test_vv
```

The conversion of the volume from 32 drives to 80 drives took slightly over three hours, with Iometer continuing to write data to the volume at an average rate of 2,035 IOPS. After completion, Iometer averaged 5,914 IOPS, approximately 2.4 times the rate with fewer disks (Figure 5).

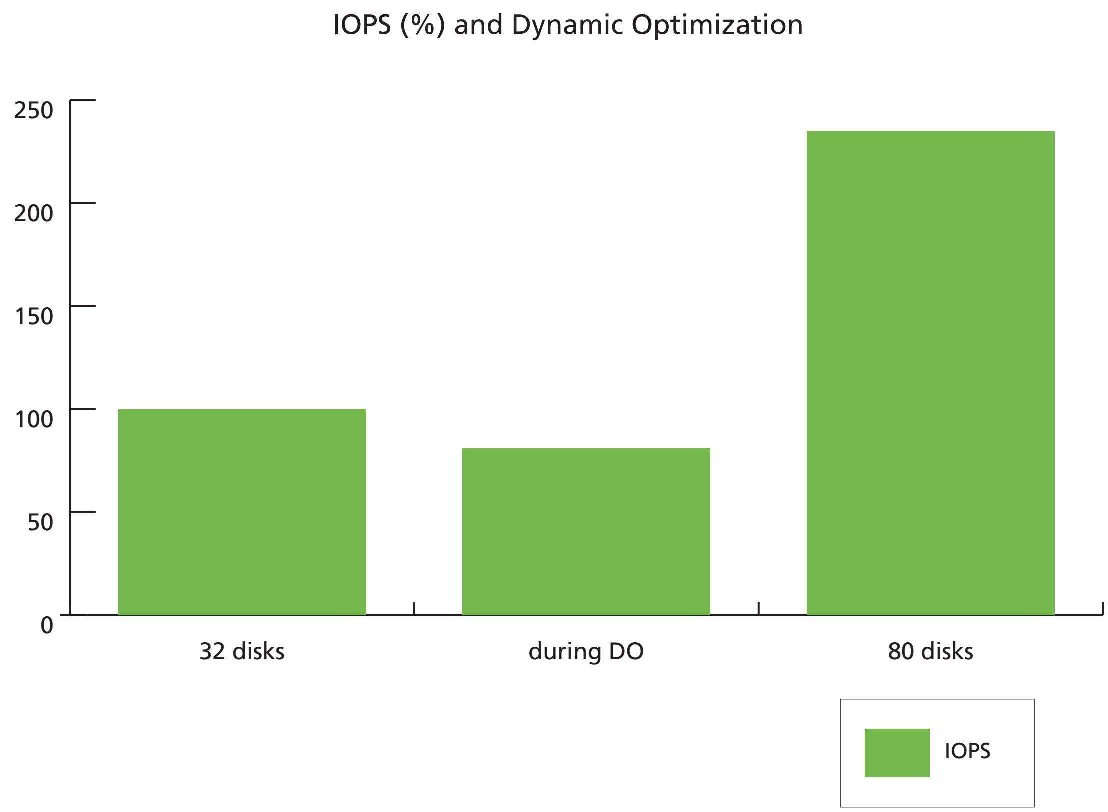

Fig. 05 Performance Before, During, and After Dynamic Optimization

4 SUMMARY

Virtualization of servers using Hyper-V increases server flexibility while reducing costs through simplified management and more effective use of hardware resources. Traditional storage arrays struggle to meet all of the performance and agility needs of a virtualized server environment. 3PAR Utility Storage solves the storage challenges of this environment using massive parallelization for exceptional performance, 3PAR Thin Provisioning to allow physical storage to be allocated when used rather than when planned, and 3PAR Dynamic Optimization to tailor storage performance without disruption. As a result of these features, 3PAR Utility Storage perfectly complements Windows Server 2008 Hyper-V technology for building an optimal virtualized server environment.

ABOUT MICROSOFT VIRTUALIZATION SOLUTIONS

Microsoft provides a complete suite of technologies to enable an integrated, end-to-end, virtualized infrastructure. Using products that span the desktop to the data center, Microsoft technologies bring capacities online in real-time, as needed; streamline and provision applications, services, and data on-demand; accelerate backup and recovery, and enhance availability to help protect against system failure and service interruptions. Microsoft's extensive partner ecosystem complements and extends the Microsoft virtualization toolset with products for desktops, servers, applications, storage, and networks. Together with our partners, we deliver the most robust, complete solutions for the virtualized infrastructure. For more information about Microsoft Joint Virtualization Solutions, visit: http://www.microsoft.com/virtualization/partners.mspx.

ABOUT 3PAR

3PAR® (NYSE: PAR) is the leading global provider of utility storage, a category of highly virtualized, tightly clustered, and dynamically tiered storage arrays built for utility computing. Organizations use utility computing to build cost-effective virtualized IT infrastructures for flexible workload consolidation. 3PAR Utility Storage gives customers an alternative to traditional arrays by delivering resilient infrastructure with increased agility at a lower total cost to meet their rapidly changing business needs. As a pioneer of thin provisioning—a green technology developed to address storage underutilization and inefficiencies—3PAR offers products designed to minimize power consumption and promote environmental responsibility. With 3PAR, customers have reduced the costs of allocated storage capacity, administration, and SAN infrastructure while increasing adaptability and resiliency. 3PAR Utility Storage is built to meet the demands of open systems consolidation, integrated data lifecycle management, and performance-intensive applications. For more information, visit the 3PAR Website at: www.3PAR.com.

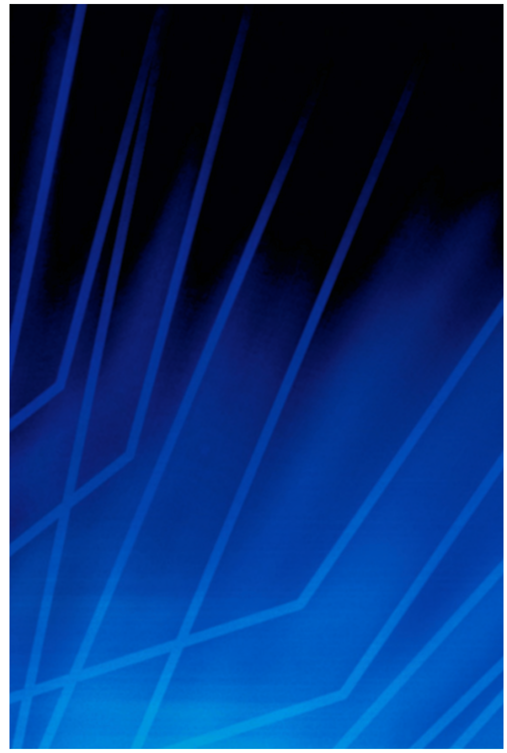

U.S. CORPORATE HEADQUARTERS

3PAR, Inc.
4209 Technology Drive
Fremont, CA 94538
Phone: 510-413-5999
Fax: 510-413-5699

Email: salesinfo@3PAR.com

Microsoft | Virtualization

3PAR
Serving Information

hyperv-overview-wp-09.0