DOUBLING VM DENSITY WITH HP 3PAR STORAGE

MARCH 2012

Server and client virtualization have revolutionized datacenters, bringing unprecedented levels of flexibility, scalability, and efficiency. But as many users have learned the hard way, a virtual infrastructure is only as strong as its weakest link. More often than not, that weakest link is storage. Legacy storage systems, and even some modern arrays, are not architecturally equipped to deliver the storage performance, capacity and management efficiencies required to match those on the compute side of the house, and to make a virtual infrastructure shine.

And when these storage efficiencies are not present, virtual machine (VM) density is one of the first things to suffer. Increasing VM density is a good thing, since it enables customers to non-disruptively run more VMs on a given server or cluster, and in the process, achieve greater resource utilization, simpler management, and lower costs.

Most storage solutions simply cannot overcome the I/O bottlenecks, long backup windows, and provisioning and management inefficiencies that plague traditional arrays in a virtual infrastructure. But fortunately, HP 3PAR Utility Storage can. As we'll see, the HP 3PAR platform is purpose-built for server virtualization environments, delivering the levels of storage scalability, flexibility and efficiency required for optimal functioning of a virtual infrastructure.

Together, HP 3PAR Storage Systems and virtualization software from VMware take this value proposition to the next level, providing customers with high levels of VM density over a range of different workloads. This joint compute-and-storage solution delivers high levels of virtual infrastructure performance, resource utilization, and manageability, which in turn leads to a lower cost of ownership.

In this paper, we'll begin by looking more closely at VM density and the impacts it has in the broader virtual infrastructure. Next, we'll turn our attention to the storage challenges that typically hamper the proper functioning of a virtual infrastructure, compromising VM density. We'll conclude with a detailed examination of HP 3PAR Utility Storage offerings, and describe how the platform overcomes typical virtual infrastructure storage issues, and increases VM density as a result.

FOCUS ON VM DENSITY

VM density has become a key metric for determining the cost and efficiency of a virtual infrastructure. What do we mean by "VM density"? Taneja Group defines VM density as the number of virtual machines that can run simultaneously on a single physical server. VM density has become one of the accepted measures of a hypervisor platform's architectural prowess, along with sustainable performance under load, and a short list of other factors.

Taneja Group has run a number of validation exercises to test and measure VM density, and to determine its impact on a virtual infrastructure. In our pioneering studies for VMware in 2009, for example, we compared the VM density of each of the leading virtualization platforms, including Microsoft Hyper-V, Citrix XenServer, and VMware vSphere. Our tests were designed primarily to evaluate hypervi-

87 Elm Street, Suite 900 · Hopkinton, MA 01748 · T: 508.435.2556 · F: 508.435.2557

sor performance under different types of workloads, and secondarily, to determine the impact that VM density levels have on key characteristics of a virtual infrastructure. The results of our testing received considerable industry attention, and helped put VM density on the "short list" of technical evaluation factors that customers now pay attention to.

In our testing, we validated that increasing levels of VM density generally lead to—or at least correlate with—other important metrics of a virtual infrastructure's effectiveness for running particular workloads, which in turn translate to concrete customer benefits. In particular, a higher VM density level can lead to benefits such as:

- Increased utilization of CPU and memory resources (i.e. greater efficiency on the host server)
- Lower CAPEX acquisition costs (since new server purchases can be deferred)
- OPEX savings, in the form of power, cooling and datacenter floor space reductions
- Reduced software licensing costs for Microsoft Windows OS and applications (e.g. MS SQL) that are licensed by physical CPU.

Together, these elements can create significant advantages in the Total Cost of Ownership (TCO) for a given virtual infrastructure.

STORAGE CHALLENGES TO VM DENSITY

Storage often presents significant obstacles to increasing VM density in a VMware infrastructure. In our experience, these obstacles are inevitable when a customer brings a legacy storage array into a modern virtualization environment. But many of the issues can also arise in conjunction with newer storage systems, which may not have been designed with server virtualization in mind.

The first challenge revolves around the I/O performance bottlenecks that commonly occur, usually as a result of random VM I/O and storage resource contention. A collection of virtual machines running on even a moderately consolidated system (e.g. in the range of 6-10 VMs) tend to generate random I/O that can stress storage resources. Moreover, across many physical servers, the aggregate number of individual workloads rapidly becomes far greater than in the physical world, and at this scale the random I/O behavior often defeats the effectiveness of cache, and increases the chance of erratic spikes that may push storage performance beyond its limit. Throw some I/O-intensive applications into the mix, and you have a recipe for I/O gridlock. I/O nexus or aggregation points, such as controllers, ports and HBAs, are especially vulnerable to these pressures.

On top of this contention, attempting to execute any additional storage operation can make an entire storage system crumble. This in turn makes data protection enormously challenging. Disk performance can be entirely insufficient to handle yet another demand for disk access, much less the intense, fully sequential throughput request involved in backup. In turn, administrators turn to alternate protection mechanisms like snapshots, but yet again, the controller is already burdened, and we routinely see storage arrays allocate a third or more of their performance footprint to handling snapshot operations. Throwing such a request onto a burdened storage system can queue I/O and cause additional contention to the point that workloads time out and devastation follows.

Worsening the contention over precious storage I/O in the virtual infrastructure is a complete lack of control when it comes to storage performance, and the only recourse is reactive and density-defeating sprawl. When particular VM-based applications are not getting the I/O resources they need to meet quality of service (QoS) objectives, something's got to give, and one of the common remedies is to reduce the number of VMs running on that system.

Less well known, is that some of this contention in the virtual infrastructure is attributable to the overhead that the virtual server layer adds to all of these workloads. It isn't that the virtual infrastructure I/O is unusually inefficient, but the virtual infrastructure brings with it additional I/O, resulting

87 Elm Street, Suite 900 · Hopkinton, MA 01748 · T: 508.435.2556 · F: 508.435.2557

from operations such as LUN creation or re-sizing, VM snapshots, and Storage VMotions, which can put pressure on application performance, and prevent users from adding VMs to a given server or server cluster.

Finally, slow and inefficient storage provisioning can also be an obstacle to greater VM density. As each new VM is added to a server or cluster, storage must be provisioned, and when that involves steps such as cloning existing volumes or formatting new ones, both of which increase I/O overhead, then all virtual server applications running on that platform will potentially suffer. Adding more VMs will only worsen the problem, and at the same time, increase OPEX costs.

The standard practice to resolve the I/O contention issues in the virtual infrastructure is to reactively provision more storage. Such maneuvers exacerbate the operational management overhead that comes with storage, require unanticipated CAPEX, deteriorate hard dollar operational savings, and dilute the overall density of an infrastructure. At the end of the day, these tactical attempts to improve virtual server storage can completely negate the TCO savings enabled by virtualization.

HP 3PAR UTILITY STORAGE: PURPOSE-BUILT FOR VIRTUALIZATION

HP 3PAR Storage Systems have always been architected for demanding, highly consolidated, and contentious storage environments. It is little surprise then that HP announced in March 2012 a double density promise called the HP 3PAR Get Virtual Guarantee—in a nutshell, a guarantee that customers moving from legacy storage arrays to HP 3PAR Storage Systems will realize an increase in VM density by a minimum of 2X.

The HP 3PAR Architecture is built for flexibility and scalability, with well-controlled storage performance that enables 3PAR arrays to overcome the common storage issues that limit VM density. The platform's fine-grained virtualization divides physical disks into granular "chunklets", which can be dynamically assigned (and later re-assigned) to virtual volumes of different QoS levels. Mesh-Active clustering, along with a data layout methodology that includes a unique approach to wide-striping, dynamically distributes and load balances each workload across all storage system resources. This enables higher utilization and minimizes the opportunity for choke points. This approach eliminates issues with the narrow striping used by traditional arrays, as well as other wide-striping implementations that don't actually use all system resources. In both of these cases, storage controllers can become a choke point and the number of spindles dictates IOPS capability.

Moreover, intrinsic to the HP 3PAR Architecture, and pervasive beneath each of its storage features -- such as snapshots and replication-- is a "Thin-Built-In" ASIC that not only optimizes the capacity of each 3PAR array, but ramps the performance of every feature by streamlining storage transactions that exercise storage system cache and disks. Similarly, tight integration with VMware vSphere optimizes transactions that pass between the storage system and the virtual infrastructure (see VMware's perspective in the sidebar below). Together, these architectural features deliver the scalability, flexibility and density in the storage layer that are needed to support a highly scalable and consolidated virtual server infrastructure, and to maximize VM density on the physical hosts (see Figure 1).

While the performance claims of HP 3PAR Utility Storage are well known, HP's double density "Get Virtual" guarantee is a tremendous gauntlet to throw down in the marketplace. Let's now examine how 3PAR Storage Systems tackle the most common performance problems we see plaguing traditional storage infrastructures, in order to assess whether their guarantee will hold up in the real world.

87 Elm Street, Suite 900 · Hopkinton, MA 01748 · T: 508.435.2556 · F: 508.435.2557

VMWARE ON VM DENSITY

By Narayan Venkat, VP Storage Product Management, VMware

Prior to server virtualization, organizations typically ran one application per physical server, so comparing application costs was as simple as a cost-per-server calculation. VMware vSphere changed all of this, giving organizations a new model that places many applications on each server by using virtual machines rather than discrete hardware to run each application. The cost savings were dramatic, and cost-per-server comparisons no longer made sense. As more and more organizations deployed vSphere, they saw first-hand how important virtual machine density—the ability to support more VMs on each host—was to reducing the cost of business operations by supporting more and more applications on less and less hardware.

As we have continued to develop technologies to increase VM density on the hosts—including Memory Oversubscription, Direct Driver Model, High Performance "Gang" Scheduler, Support for Large Memory Pages and Nested Page Tables—we have seen our clients experience greater cost savings and return on their server virtualization investments. However, we also started seeing more and more of our clients struggling to realize the full benefit of vSphere because their storage simply couldn't keep up. We were seeing that legacy storage was presenting bottlenecks that restricted the number of applications that could be supported on host servers—not due to any limitations of vSphere or the servers themselves, but because, when they combined the workloads of multiple applications (each with its own unrelated I/O stream), they ended up with very random I/O patterns that their storage wasn't designed to handle. This means that application performance would suffer from increasing VM density, or they would have to trade performance for storage resiliency and availability—which, consequently, became all the more critical as they increased VM density.

Then came a new type of storage that solved this conundrum. From the outset of our partnership, the 3PAR platform has served as a natural companion to vSphere. Throughout this relationship, we have seen each generation of 3PAR array designed from the ground up to deliver higher application performance and support greater VM density on hosts by removing the barriers we had observed with traditional storage. Our customers were finding that the platform's advanced storage virtualization capabilities—such as wide striping, mixed workload support, and a unique, clustered Mesh-Active architecture—enabled them to maximize vSphere benefits by supporting more application workloads at higher performance levels and greater availability. During this time, we've also worked closely together on vSphere Storage technology integration (for example with VAAI, to allow ESX hosts to offload virtual operations to the storage array), host overhead and serving more applications with fewer and fewer physical resources.

87 Elm Street, Suite 900 · Hopkinton, MA 01748 · T: 508.435.2556 · F: 508.435.2557

How Does HP 3PAR Utility Storage Eliminate I/O Performance Bottlenecks?

HP 3PAR Storage overcomes I/O performance issues through a number of innovations. First, as previously stated, the platform's Mesh-Active controller clustering architecture reduces pressure on storage performance by dynamically load balancing and distributing each workload across all storage system resources (drives, nodes, loops, cache, and ports). Similarly, wide striping of volumes across all drives significantly reduces resource contention to take application performance to a new level. Since a higher number of VMs can now use shared storage resources without running into performance limits, VM density can be increased dramatically.

Second, the HP 3PAR Architecture is built specifically to handle mixed workloads, allowing transaction- and throughput-intensive workloads to coexist on the same array without contention. Data movement is handled by the 3PAR ASIC and Data Cache, while metadata is processed by Intel CPUs with support of the Control Cache. Because these activities are segregated, workloads of polar extremes –such as heavy random I/O behind virtual machines, combined with a streaming sequential backup– can happen in parallel without disrupting each other.

Third, technologies like Hardware Assisted Locking, supported by the platform's software Plug-In for VMware VAAI, integrates 3PAR storage operations more deeply into the virtual infrastructure—in this case, by performing more granular locking that prevents hosts from locking entire volumes to service the needs of a single VM—and allow more VMs to share the same storage pool.

Figure 1: HP 3PAR Utility Storage Enables High Virtual Machine Density

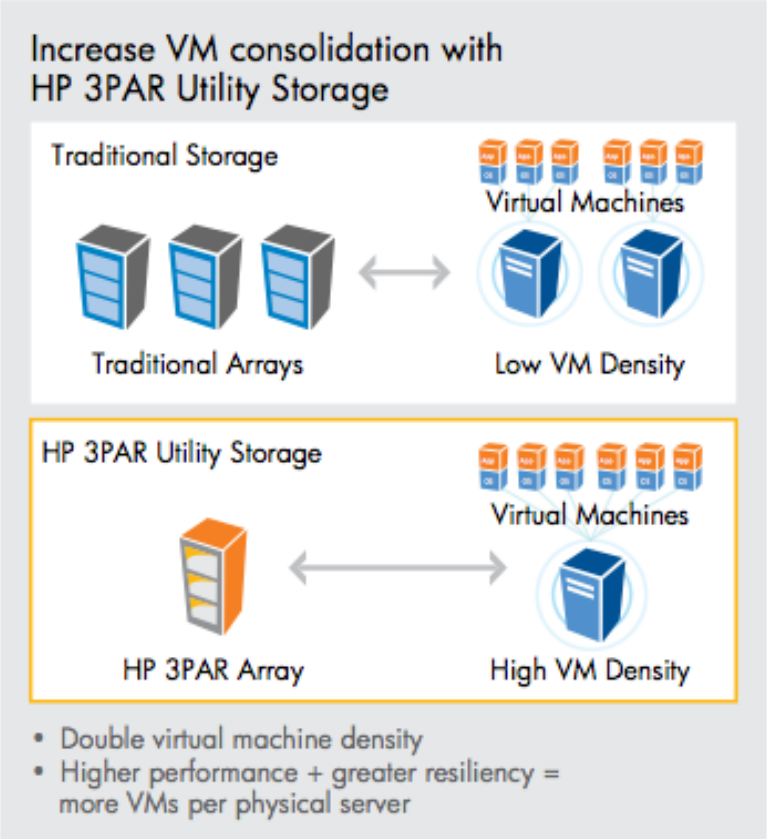

Source: HP 3PAR

87 Elm Street, Suite 900 · Hopkinton, MA 01748 · T: 508.435.2556 · F: 508.435.2557

How Does 3PAR Tackle the Performance Conundrum Behind Protection?

Highly efficient, thin copy technologies—combined with a deep level of VMware vSphere integration—enable HP 3PAR Utility Storage to significantly streamline and accelerate backup and recovery operations in a virtual infrastructure. HP 3PAR Virtual Copy Software allows users to instantly create hundreds of point-in-time snapshots per volume without impacting the performance of VM-based applications. A sister offering, HP 3PAR Remote Copy Software, facilitates efficient and non-disruptive storage replication for business continuity purposes. These solutions dramatically reduce backup windows and ensure plenty of headroom to non-disruptively add more VMs to each server and cluster.

HP 3PAR Recovery Manager Software for VMware vSphere, an HP 3PAR and VMware integrated offering, also plays a role in promoting greater VM density. This data protection solution enables users to non-disruptively create hundreds of space-efficient, point-in-time snapshots online in an HP 3PAR array for recovery purposes, without impacting the performance of other running VMs, thereby supporting greater VM density.

How Does 3PAR Provide Storage Control for the Virtual Infrastructure?

HP 3PAR Utility Storage provides several facilities to guarantee performance of critical applications, and this in turn allows more workloads to share a single system, without worrying about compromising the performance of key services. Several mechanisms are responsible for guaranteeing this storage control. First, HP 3PAR Adaptive Optimization (AO) Software optimizes service levels by intelligently moving and placing data at the sub-volume level on the storage tier that can meet its service-level requirements at the lowest possible cost. Data is moved at a granular level on an ongoing basis, to ensure that data is always in the appropriate tier at a given point in time. 3PAR AO software complements VMware Storage Distributed Resource Scheduler (SDRS). SDRS takes care of the initial placement of VMs and VM disk (VMDK) files, as well as load balancing of VMs based on space usage; while 3PAR AO dynamically load balances data across storage tiers to ensure VM workloads meet their quality of service objectives.

As a second example, VMware Storage I/O Control (SIOC), also supported by HP 3PAR Storage Systems, provides QoS controls to ensure that the highest-priority VMs are always given a specified amount of I/O resources. This is particularly important in a highly consolidated environment in which multiple VMs, potentially extending across multiple instances of vSphere, are competing for the same I/O resources at the same time. The workloads need not be homogeneous, since VMware SIOC enables a large number of VMs with a mixed set of workloads to co-exist on the same server while delivering required QoS levels. These offerings jointly allow for greater VM density in a VMware and HP 3PAR combined infrastructure, even for a heterogeneous set of workloads.

Third, to facilitate both QoS and secure access for specific applications and user groups, HP 3PAR Virtual Domain Software allows users to create virtual partitions of storage resources in a physical array. Administrators can logically segregate and allocate storage resources to particular users and applications, without compromising the performance and availability benefits of 3PAR's wide striping of resources. This enables users to support the storage needs of an increased number of VMs, while still delivering the required storage service levels for critical applications.

How Does 3PAR Offload Virtual Operations to Reduce Overhead?

To complete the picture, HP 3PAR Utility Storage also makes virtual server storage management in VMware environments more efficient, enabling users to optimize resource utilization without incurring virtual infrastructure overhead that would bog down running virtual machines. Several HP 3PAR and VMware technologies play a role in making this happen.

First, a series of 3PAR thin technologies allow virtual machine capacity management to happen effi-

87 Elm Street, Suite 900 · Hopkinton, MA 01748 · T: 508.435.2556 · F: 508.435.2557 www.tanejagroup.com

ciently in the array (see Figure 2). At the beginning of the process is HP 3PAR Thin Provisioning Software, which is designed to help with the autonomic provisioning of highly-efficient storage volumes. Capacity is allocated only as needed and configured intelligently without active management. Next, HP 3PAR Thin Persistence Software takes over, automatically and non-disruptively ensuring that storage volumes remain compact and efficient over time by leveraging the HP 3PAR ASIC to reclaim unused space associated with deleted data. As a third 3PAR thin technology, HP 3PAR Thin Conversion Software ensures that storage volumes migrated from legacy arrays—which happens commonly with new installations—also utilize capacity efficiently, without requiring hypervisor resources. Collectively, these thin technologies offload capacity bloat from the virtual infrastructure, and use array power to non-disruptively keep storage thin over time.

Several VMware-specific integrations further offload virtual operations to the array, where they can be handled more efficiently. 3PAR Block Zero for VMware, supported in 3PAR's Plug-In for VMware VAAI, accelerates large, block-level write operations, making tasks such as provisioning VMs from templates much quicker and easier. 3PAR Full Copy for VMware environments, a second VAAI capability, offloads large data movements to 3PAR array hardware, allowing operations such as VM cloning and Storage VMotion-enabled workload migrations to be accomplished in about half the time they take on competitive arrays that lack the benefit of this integration. 3PAR Hardware Assisted Locking, a third VMware VAAI enhancement, offloads storage locking operations from the hypervisor to give IT managers the flexibility to perform VMware administrative functions—such as creating new VMs, powering VMs on and off, creating VM snapshots, and migrating VMs via VMotion—during peak hours, without significantly impacting application performance.

Figure 2: HP 3PAR Thin Technologies Ensure Space Efficiency Throughout the Storage Lifecycle

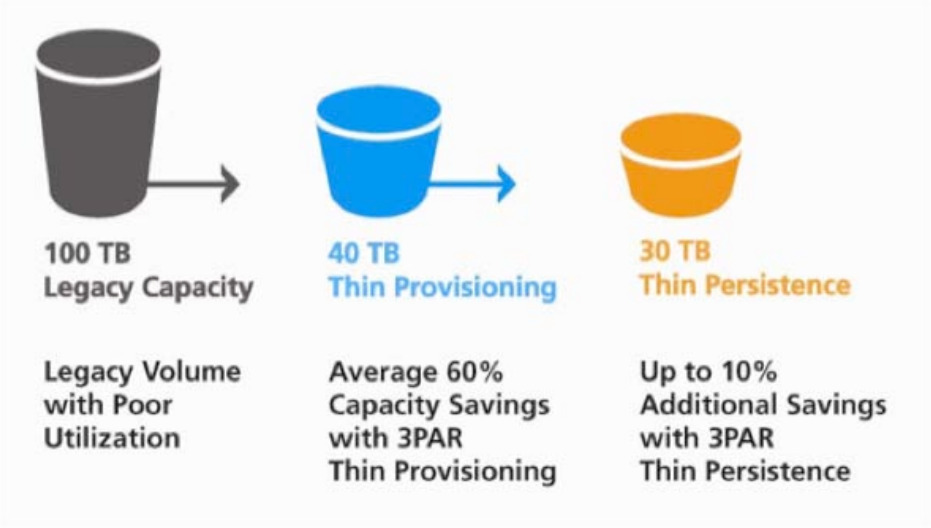

Source: HP 3PAR

How Does 3PAR Streamline Provisioning and Management?

HP 3PAR Utility Storage speeds storage provisioning and simplifies management tasks, enabling the storage for a large number of VMs to be consolidated in a single array. For example, HP 3PAR Autonomic Groups Software significantly streamlines the cumbersome, manual repetition of storage provisioning in both virtual server and virtual desktop infrastructures, by dramatically reducing the number of actions required to provision new volumes. HP 3PAR Recovery Manager Software for

87 Elm Street, Suite 900 · Hopkinton, MA 01748 · T: 508.435.2556 · F: 508.435.2557 www.tanejagroup.com

VMware vSphere enables VMware administrators to rapidly provision new virtual machines using VM copies, and at the same time, to automatically protect and recover VMs, VMDKs, or even individual files.

HP 3PAR Management Plug-In Software for VMware vCenter allows administrators to manage 3PAR storage directly from the vCenter Server, where they can view how storage resources are being exercised, including precise insight into how individual VMs are mapped to datastores and specific volumes.

While management overhead is seldom considered a component of VM density, it can make or break virtual environments. Density only goes so far if the data center cannot hold enough administrators to manage it. The HP 3PAR emphasis on zero-impact administration not only simplifies the lives of administrators, but also enables a higher number of virtual servers or desktops to run on any given server or cluster.

TANEJA GROUP OPINION

As we've learned in frequent discussions with end users, storage is an incredibly important part of a virtual infrastructure, and a key factor in determining whether such an infrastructure will meet customer expectations. To excel in a virtualized environment, a storage array must overcome some thorny challenges to deliver the levels of flexibility, scalability and efficiency demanded by virtual server users and applications.

Unfortunately, a majority of storage arrays currently operating in customer environments are not up to this task. This includes nearly all legacy storage systems, as well as a number of arrays that have been introduced to the market during the past five years.

HP recognizes the importance of taking a holistic approach to virtual infrastructures, and ensuring that all of the ingredients—servers, hypervisors, storage, networking, software and services—are in place to deliver customers an optimized solution. HP 3PAR Utility Storage is the crown jewel of HP's block-based storage offerings, and is perfectly designed to meet the stringent demands of a virtual server infrastructure.

VM density is a key measure of the overall effectiveness of a virtual infrastructure, and is a metric that a joint VMware and HP 3PAR solution excels in. In fact, HP is convinced that a joint VMware-3PAR solution can at least double the VM density that customers running legacy arrays are achieving in VMware environments today.

HP is so convinced, in fact, that it has launched a new program called the HP 3PAR Get Virtual Guarantee to stand behind its belief, with a "double your VM density" guarantee for HP 3PAR arrays running in VMware environments. Given 3PAR's impressive track record as a storage solution in virtual infrastructures—with the satisfied customers to back it up—we are willing to bet that HP can deliver on this claim, and think that many virtual infrastructure customers will too.