WHITE PAPER

Risk-Free Storage for Virtual Desktop Infrastructure (VDI)

VDI has unique performance characteristics that legacy storage architectures were never designed to support. This paper explains how Nimble Storage flash-enabled storage solutions provide a solid platform for VDI environments that eliminates extensive management of tiered legacy storage architectures to provide just-in-time performance.

Executive Summary

Desktop and notebook client computer prices have dropped steadily over time. However, management has remained the largest portion of computers' total cost of ownership. Providing break/fix support and updating operating systems, applications, and hardware make up the majority of these costs. Virtual Desktop Infrastructure (VDI) is a new client computing approach that takes over where thin computing leaves off, making desktop management a more sustainable model. Virtualizing a desktop computer takes a cue from mainframe computing by leaving the client hardware on the user's desk and moving the operating system and software to the data center.

Centralizing desktop computer resources allows IT to more effectively manage software updates and protect user data in a predictable way. VDI solutions leverage enterprise storage resources to provide advanced features such as high availability, snapshots, and replication to protect user systems. For example, when user problems arise, VDI implementations can solve them by quickly reverting to a recent snapshot of the system state that was prior to the issue, allowing users to remain productive.

The single greatest challenge for virtual desktop initiatives is providing enterprise-class storage performance, functionality, and management at a reasonable cost. Why is this a challenge? Because the management cost/benefit equation of VDI is well understood, but storage performance requirements are routinely underestimated. And it can be very expensive to meet the performance and capacity requirements of even a modest VDI implementation with traditional enterprise storage architecture.

Legacy SAN vendor architectures have traditionally relied on scaling performance with large arrays of high-RPM disk drives, but hard drive performance (as measured in IOPS) has hit its limits. In recent years, traditional storage vendors have attempted to meet VDI performance requirements by adding enterprise-class flash solid state drives (SSDs), but costs have been prohibitive. A new storage approach is needed to maximize VDI performance while reducing both capital and operational expense.

Nimble Storage has created the first converged storage solution that seamlessly merges flash drives and low-cost SATA disks to deliver extremely high performance at a reasonable cost. The performance of a single Nimble Storage array easily outpaces much more expensive tiered storage solutions. By combining primary storage and backup into the same storage architecture, Nimble eliminates backup windows, greatly improves recovery point objectives (RPOs), and dramatically reduces recovery times. Nimble has created a new storage approach that overcomes the limitations of legacy storage architectures with a much more affordable total cost of ownership that makes it a perfect fit for VDI.

VDI Storage Challenges

VDI Storage Performance Requirements

The daily use cycle of a virtual desktop has two primary phases: *boot* and *steady use*. The boot phase takes place when users power on and log in to their computers. This phase causes the most intense read activity that a desktop typically experiences. The steady use phase generates fewer total IOPS but is heavily weighted toward writes. Less frequent sustained storage I/O load occurs during update patching and virus scans, which also consist heavily of writes. A well-performing VDI storage solution must not only accelerate reads, but also must optimize write performance.

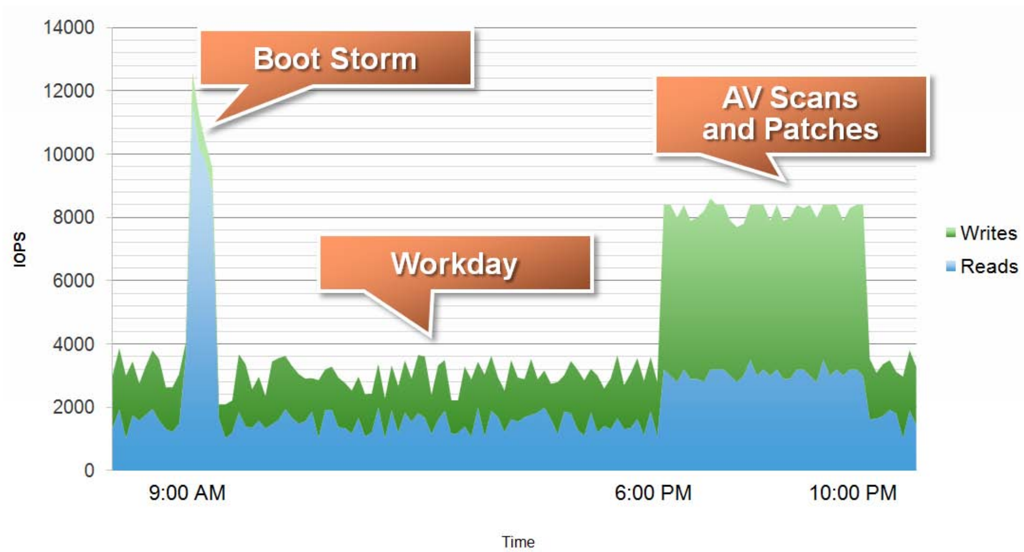

The boot phase is traditionally isolated on each user's physical desktop computer with dedicated CPU, memory, and storage resources. During the boot process, the operating system and third-party system-level tools such as antivirus are read from disk and loaded into memory. The VDI model shifts this isolated heavy read activity to the shared storage infrastructure and forces desktops to compete for storage resources. For example, Windows 7 boots to the logon screen, which requires reading about 350 Megabytes, and after processing files, an average desktop typically boots in about 30 - 60 seconds. Maintaining reasonable boot times requires much higher I/O performance in shared storage VDI environments when dozens, hundreds, and even thousands of virtual desktops boot from the same shared storage resource. This time span of large read activity is often referred to as a *boot storm*.

Once a desktop completes its boot cycle, read and write IOPS decrease dramatically to a steady state. Typical steady-state estimates for average desktops are 6-10 IOPS, while power users might require 25 IOPS. Heavy users (such as developers compiling code) might use as many as 50 IOPS. However, a boot storm far exceeds the steady use IOPS requirements, which means that sizing for boot storms will meet VDI requirements under most circumstances.

User Type	Avg. IOPS
Average User	6 - 10
Power User	20
Heavy User	50

VDI storage performance demands are also heavy during upgrades and anti-virus activity. These types of heavy IOPS activities tend to peak near boot storm levels. However, the load lasts much longer, with heavier write activity which puts additional strain on VDI storage. By itself, adding cache to meet heavy read IOPS boot storms isn't enough to meet heavy write IOPS demand. Well-designed VDI storage needs intelligent caching for read bursting as well as performance enhancements for write bursting.

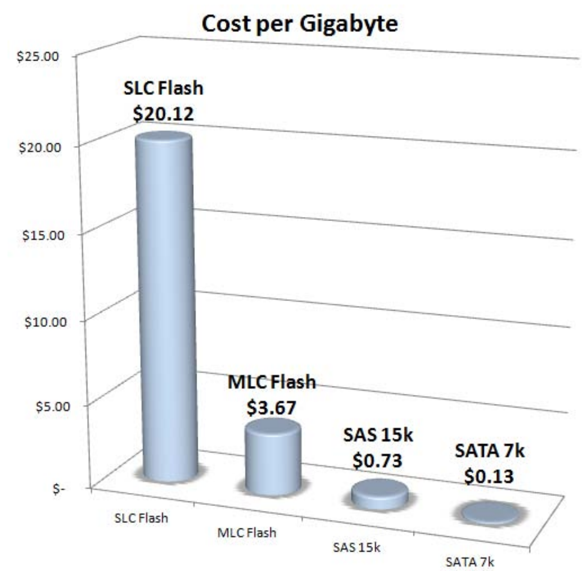

Meeting the performance demand for IOPS is the primary concern of VDI storage sizing. However, overprovisioning for storage performance dramatically increases costs. Traditional physical desktop implementations for 1,000 users would allocate 1,000 hard drives in each desktop computer. However, this level of overprovisioning would be prohibitively expensive to rack and power in a centralized datacenter. Storage vendors use I/O optimization and caching techniques to provide performance that is greater than the sum of the individual hard drives. In the past, SAN architectures attacked performance challenges by adding more power-hungry, high-speed disks as IOPS requirements grow. This solution simply isn't practical in the VDI world, because the much higher price point of replacing desktop SATA drives greatly increases up-front capital costs.

In recent years, the introduction of flash drives has dramatically increased storage performance, greatly benefitting VDI IOPS read bursts during boot storms. While flash provides high performance, it typically comes with a high price tag and behaves differently than spinning hard drives. This is because most file systems write data to flash the same way they write to a spinning disks, instead of recognizing the unique characteristics of flash media. When managed the same as spinning disks, flash is well-suited for reads but suffers from shorter useable lifespans: it wears each time it is written to. Storage vendors compensate by using extremely high-priced flash variants such as SLC (single level cell)-based enterprise flash drives rather than MLC (multi-level cell) drives which are much more cost-effective. Additionally, tiered flash solutions require RAID parity overhead that further increases cost by cutting usable capacity in half. Virtual Desktop Infrastructures require a new approach to storage that maximizes read performance without suffering the write penalties that limit lifespan—all at a reasonable per-desktop price point.

VDI Backup Storage Challenges

As long as computers store data on media that can fail, some form of backup will be needed. IT practices have sought to eliminate data from end-user desktops and move it to the data center. This promotes sharing, but more importantly it provides centralized backup. While centralized data does make backups easier, it doesn't foster mobility for users who aren't tethered to their desktops. Thus, data has been allowed to roam with mobile users and reside on desktops over time for varying reasons, both good and bad. VDI solutions change this situation for the better by shifting the computer resources into the data center and providing a virtualized desktop computer experience that travels with users—even on non-PC devices. However, VDI implementations still need to be protected against data loss or corruption. That means some form of data protection is required.

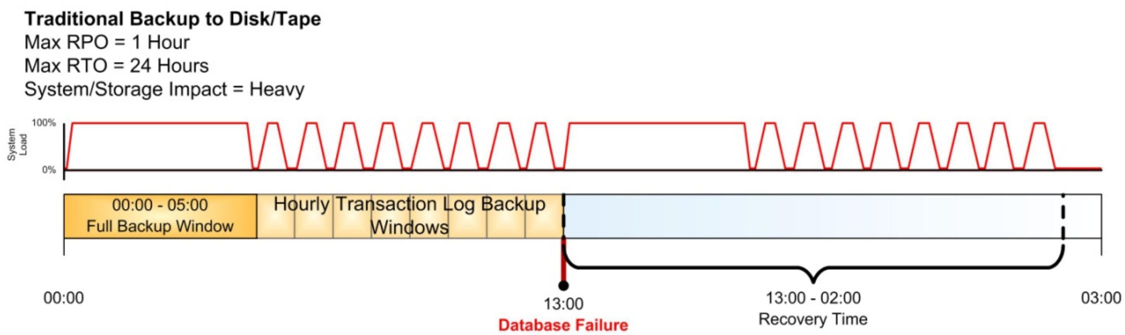

Backup solutions have evolved to provide better management capabilities since the introduction of virtualization technology. However, core backup functionality hasn't changed much over the years and still involves copying data from primary storage to backup storage, which takes time and has an adverse impact on production performance. The real pain of backup is truly experienced when performing data restoration: it is extremely cumbersome to find the right version of data that you need and then restore it back to primary storage. If that version of the data turns out to be wrong, the restoration process is repeated. For large data sets, that means waiting hours or even days before you know if the data is appropriate. All backup solutions share this painful cycle, regardless of media (tape or disk), or optimizations such as de-duplication. These solutions are really just a way to back up more data on the same disk media, and that data won't be useable until restored to primary storage.

IT is being forced to spend larger percentages of their budget on backup as data footprints get larger, recovery time expectations get shorter, and backup windows get smaller. Add virtual desktop protection and you have a potential management nightmare. Copy-based full backup technology provides a good solution for performing long-term data protection for the years of retention often mandated by regulatory requirements. However, the continuing sprawl of data has long since outgrown the ability of copy-based backup technology to provide near-term data restoration when you need it the most.

SAN- based snapshot technology has the potential to provide efficient data protection at a more frequent interval than once-a-day backup. Well-architected snapshot technologies are more lightweight than traditional backups because they don't have to scan the entire data set looking for modified data to copy. When applications write data to disk, they are telling the storage precisely what they are changing, which allows snapshots to create a frozen image of data without making a complete copy of data. However, snapshots have historically had high associated costs because they're stored on the same disk as high-performance primary storage. This made long-term snapshot retention too expensive. Some storage vendors have tried to alleviate the snapshot retention costs by allowing snapshots to reside on more cost-effective SATA media. However, the performance impact is very high for copy-on-write (COW) snapshots, because of the additional read and write of the existing data prior to allowing the new data write. This makes these COW snapshots undesirable in all but the smallest IT environments. While backup isn't exclusively a challenge to VDI, centralizing desktop computers brings with it a great deal of additional pressure to adequately protect them, and IT is growing desperately in need of a new backup approach.

Nimble Storage Solutions for VDI

VDI storage requires a new approach to cost-effectively meeting high performance bursts and long-term storage to maximize the storage density of large numbers of desktops. Using an all-Flash storage solution for VDI attacks the performance challenge directly, but at tremendous cost per desktop and for very little capacity, which impacts scalability. Adding flash to a legacy tiered storage architecture and treating it like a spinning disk doesn't provide an adequate solution due to the cost and resiliency issues previously described. Nimble Storage had the advantage of creating a storage architecture from scratch using the most recent advances in storage technology without the baggage inherent in legacy tiered storage architectures created over the past 20 years. Nimble Storage created the CASL (Cache Accelerated Sequential Layout) architecture that makes the best possible use of new technologies like flash SSD, high-density disk drives, and multi-core CPUs to create a new converged storage platform for the next 20 years.

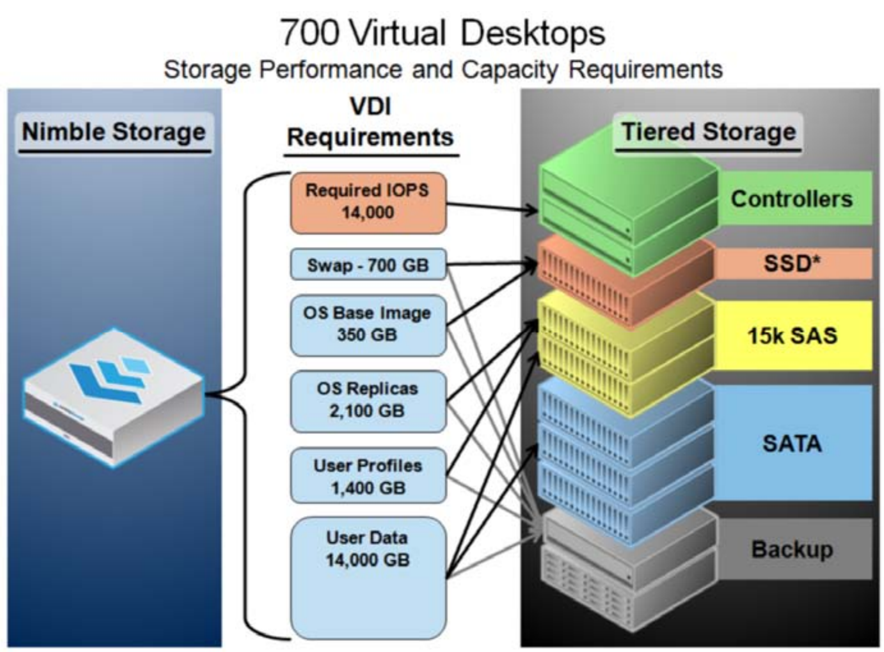

Nimble Storage Maximizes VDI Read and Write Performance

Nimble Storage's CASL architecture leverages multi-core CPUs and its variable-length block architecture to perform inline data compression on writes, maximizing hard drive and flash storage capacity. This approach generally results in a 50% increase in usable storage. CASL also optimizes write activity, which is critical for scaling VDI implementations and reducing overall cost per desktop. It's well known that hard drives have very good performance working with sequential I/O but suffer under random I/O loads because of the physical head movement (disk seeking) required by modern file systems. Nimble's CASL architecture uses a new approach that physically coalesces random writes before committing data to storage media, which greatly reduces hard drive head movement. In addition, all data changes are written to hard drives in full array stripes, eliminating the searching overhead associated with legacy hole-filling file systems that impacts performance and reduces storage density. CASL actively monitors I/O patterns and caches random I/O to flash, optimizing read performance of hot data in real time. Caching random I/O provides an instant performance boost to incoming random reads that are served out of high-performance flash, alleviating hard drive reads. CASL also writes to flash using full erase blocks, which improves the durability of flash storage and permits the use of lower-cost MLC drives without sacrificing longevity. An added benefit to treating flash SSD as a cache rather than tiered primary storage is that you no longer require RAID, saving cost and minimizing wear on the media because only hot data is cached. By optimizing the use of flash, Nimble Storage arrays dramatically outperform legacy tiered storage.

Nimble Zero-Copy Cloning

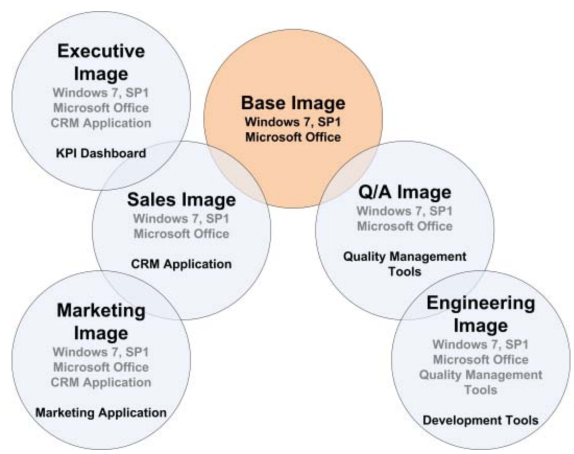

Nimble Storage provides advanced zero-copy cloning technology that quickly provisions new volumes that reference base images in seconds, regardless of their size. This allows fast provisioning of new VDI desktop boot images while taking advantage of block sharing and inline compression to optimize valuable storage resources. For example, you can create clones of a base image and add role-specific applications and configurations. Nimble also allows you to create clones of clones for additional reuse of pre-existing configurations. While each clone has applications and configuration associated with specific roles, the clone doesn't need to consume the common storage blocks used by parent images and clones. Nimble zero-copy clones greatly reduce the total storage footprint required for VDI across the organization.

Better VDI Backup

Nimble Storage converges primary storage and backup as an integrated function of the storage architecture, which provides a 24/7 backup window without production impact. Converged backup allows you to quickly protect data stored on hundreds and thousands of virtual desktops within seconds, rather than using legacy storage disk-to-disk backup that requires lengthy network copies.

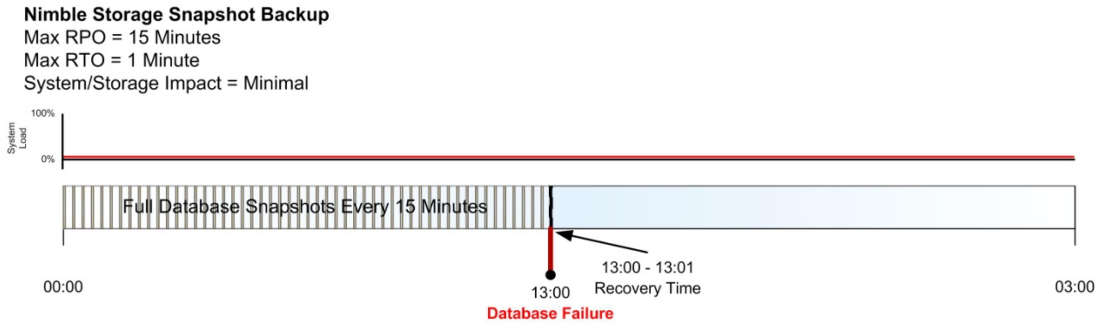

Nimble arrays are typically sized to retain 60-90 days of data change, which meets the vast majority of restoration requirements for VDI environments. Nimble's CASL architecture uses redirect-on-write (ROW) snapshots with compressed single block-level granularity. These snapshots are dramatically superior to traditional copy-on-write snapshots in terms of performance and space savings. Additionally, CASL natively uses compressed block sharing technology to further reduce snapshot backup storage requirements by avoiding duplication of data. Nimble's snapshot backup technology allows more frequent snapshots which greatly improve recovery point objectives while maximizing long-term retention.

Summary

VDI has unique performance characteristics that legacy storage architectures were never designed to support. However, Nimble Storage's modern storage architecture provides a solid platform for virtual desktop infrastructure environments. Nimble provides performance auto-tuning that eliminates extensive management of tiered legacy storage architectures to provide just-in-time performance for heavy-read boot storms. By optimizing storage writes, Nimble more cost-effectively meets IOPS demands during update storm and anti-virus scans while maintaining high read/write IOPS during steady use. Nimble’s zero-copy cloning optimizes the storage footprint, using compressed, shared blocks that reduce deployment and management costs. At the same time, converged primary and backup storage architecture was designed to provide instant snapshot backup of VDI images, eliminating traditional backup overhead.

Nimble Storage provides the most complete storage solution for VDI implementations today and scales performance seamlessly to grow with your infrastructure in the future. To learn more, visit www.nimblestorage.com.

2740 Zanker Road, San Jose, CA 95134
Phone: 408-432-9600; 877-364-6253
Email: community@nimblestorage.com
www.nimblestorage.com