

Efficient and Adaptive Flash-Optimized Storage for Virtual Desktop Infrastructure

With desktop and application virtualization gaining momentum, IT needs to carefully examine their infrastructure's ability to meet the unique storage characteristics for virtual desktop infrastructure (VDI). This paper explains how Nimble Storage flash-optimized hybrid storage arrays and comprehensive SmartStack solutions eliminate the complexity of traditional storage architectures by delivering efficient and adaptive performance and capacity for VDI environments.

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. 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 datacenter.

Centralizing desktop computer resources allows IT to more effectively manage software updates and protect user data in a predictable way. VDI solutions leverage enterprise network 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 efficiently providing enterprise-class storage performance, functionality, and management. 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.

Traditional SAN architectures have 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 flash-optimized hybrid storage solution that seamlessly combines flash-based solid state drives as a read cache and cost-effective high-capacity disks with a write-optimized data layout to efficiently deliver adaptive performance. The performance of a single Nimble Storage array easily outpaces much more expensive traditional storage solutions by as much as 10x the performance and with sub-millisecond latencies.

By integrating data protection the same storage architecture, Nimble Storage eliminates the need for separate solutions for user and VM/application data protection and backup windows, greatly improving recovery point objectives (RPOs) and recovery time objectives (RTOs). This new storage architecture facilitates a much more affordable total cost of ownership (TCO) for VDI deployments.

VDI Storage Challenges: 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. Additionally, sustained storage I/O load occurs during OS/Application updates or patching as well as virus scans, both of which also heavily consist of writes. A well-performing VDI storage solution must not only accelerate reads, but also must optimize write performance.

The boot phase on traditional desktops 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 hundreds of megabytes, and after processing files, an average desktop typically boots in about 30 to 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 storage performance challenge is directly proportional to the number of desktops being booted, and the time span of large read activity is often referred to as a “boot storm.”

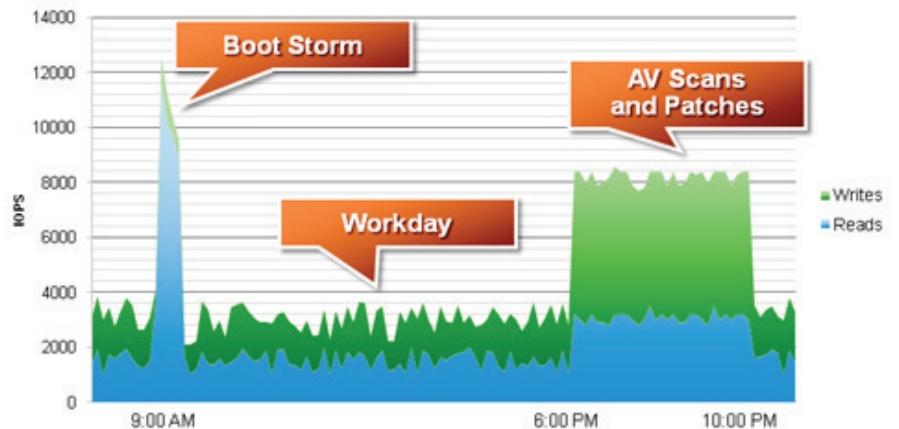
Once a desktop completes its boot cycle, read and write IOPS decrease dramatically to a steady state. Typical steady-state estimates for average desktops are 10 to 15 IOPS, while power users might require 25 to 35 IOPS. Heavy users such as developers compiling software might use as many as 50 to 75 IOPS. 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.

Popular VDI brokers such as Citrix® XenDesktop™ and VMware® Horizon View™ try and address some of the performance needs for boot storms by leveraging server resources to cache commonly read data, but this may not work well in all cases. For example in the scenario where end-users are given persistent desktops and high degree of customization in terms of type of operating system and boot options, it might not be practical to leverage server resources for caching purposes.

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-intensive operations such as boot storm as well as performance enhancements for write-intensive operations such as patching and updates.

Meeting the performance demand for IOPS is the primary concern of VDI storage sizing. However, over-provisioning 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 over-provisioning 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, storage 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 low cost desktop hard drives greatly increases up-front capital costs.

In recent years, the introduction of flash drives has dramatically increased storage performance, greatly benefitting VDI performance 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 storage 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 life spans: it wears each time it is written to. Storage vendors compensate by using extremely high-priced flash variants such as SLC (single level cell) or eMLC (enterprise multi-level cell) flash drives rather than MLC (multi-level cell) drives, which are much more cost-effective.



Storage in VDI environments needs to handle high-impact events such as boot storms and AV scans.

User Type	Average IOPS
Task Worker	10-15
Knowledge User	25-35
Power User	35-50

Storage IOPS vary greatly based on the type of end-user.

Additionally, tiered flash solutions require RAID parity overhead that further increases cost by cutting usable capacity in half. Virtual Desktop Infrastructures require a new, efficient approach to storage that maximizes read performance without suffering the write penalties that limit life span—all at a reasonable per-desktop price point.

VDI Storage Challenges: Data Protection


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 a 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 datacenter 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—either in the same system using copy-on-write snapshots, or more frequently from primary storage to a separate backup storage. Both methods take time and have an adverse impact on storage performance. The real pain of backup is truly experienced when performing data restoration, especially with a separate storage system for backup. It 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 rapidly grows, recovery time expectations get shorter, and backup windows get smaller. Add virtual desktop and end-user data 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.


Efficient redirect-on-write (ROW) snapshot technology has the potential to provide efficient data protection at a more frequent interval than once-a-day backup. Well-architected ROW 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 ROW 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.



RPO and RTO delivered by traditional backup solutions.

Nimble Storage Solutions for VDI

VDI storage requires a new approach to efficiently meeting high performance needs and long-term storage to maximize the storage density of large numbers of desktops. Using a flash-only storage solution for VDI attacks the performance challenge directly, but it comes at tremendous cost per desktop and with very little capacity. Simply adding flash to a traditional tiered storage architecture doesn't provide an adequate solution due to the cost and resiliency issues previously described. Flash-optimized hybrid Nimble Storage had the advantage of creating a storage architecture from scratch without the restrictions inherent in legacy storage architectures. 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 flash-optimized hybrid storage platform.



Nimble Storage's CASL architecture delivers adaptive performance for VDI.

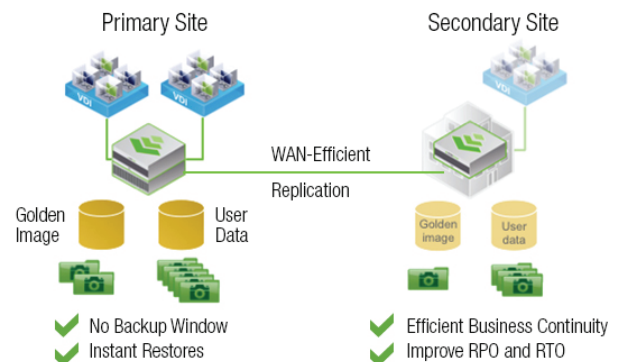
Nimble Storage's CASL architecture leverages multi-core CPUs and its variable-length block architecture to perform inline data compression on writes, maximizing hard drive and flash caching capacity. This approach generally results in a 30 to 75 percent increase in usable storage. CASL also optimizes write activity, which is critical for scaling VDI implementations and reducing overall cost per desktop.

CASL leverages the excellent sequential performance of hard drives and avoids any random writes to disk. This is done by coalescing or collecting random writes in system memory first before committing data to disk. 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.

The CASL architecture also actively monitors storage access and caches random data to flash, optimizing read performance in real time. Caching random I/O provides an instant performance boost to repeating incoming random read requests that are served out of high-performance flash, alleviating the reliance on hard drives. As flash is being used as a cache rather than as a tier, and the data integrity is ensured using a best-in-class checksum algorithm, there is no need for RAID, saving cost and maximizing utilization. Bottom line CASL-enabled Nimble Storage arrays efficiently and adaptively meet the needs of VDI better than traditional and flash-only storage architectures.

Nimble Storage 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 virtual 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 Storage 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 Storage zero-copy clones greatly reduce the total storage footprint required for VDI across the organization.



Keep VDI up and running with highly available storage and a combination of snapshots and efficient replication.

Better VDI Backup

Nimble Storage integrates data protection in the CASL storage architecture. This provides a 24/7 backup window without any production impact. Snapshot-based backup allows you to quickly protect data stored on hundreds and thousands of virtual desktops and user data within seconds, rather than using legacy storage disk-to-disk backup that requires lengthy network copies.

Nimble Storage arrays can store months' worth of data change, which meets the vast majority of restoration requirements for VDI environments. Nimble Storage CASL architecture uses redirect-on-write (ROW) snapshots with compressed single block-level granularity. These snapshots are 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 Storage snapshot backup and efficient replication technology allows more frequent snapshots and disaster recovery, which greatly improve recovery point objectives while maximizing long-term retention.

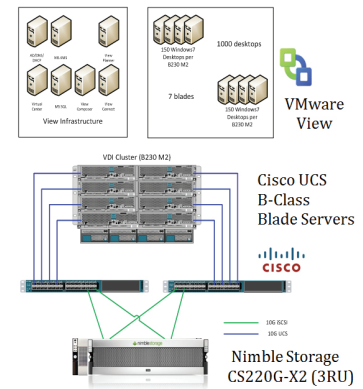
Nimble Storage SmartStack for VDI

Nimble Storage SmartStack for VDI is a pre-tested and validated end-to-end solution for desktop virtualization leveraging best-of-breed compute, networking, storage, and software. It eliminates much of the guesswork involved in upfront planning and sizing as well as simplifying common tasks such as provisioning and operations. SmartStack for VDI leverages the Cisco® Unified Computer System™ (UCS) to create a scalable foundation for both VMware Horizon View and Citrix XenDesktop.


By deploying a best-of-breed SmartStack for VDI, you take advantage of:

- Lower cost of acquisition, better ROI, and a clear recipe for successful implementation
- Simplified management and flexible scalability
- Better desktop performance to drive higher levels of productivity
- Better support for SLAs and compliance initiatives

Nimble Storage had published validated reference architectures for both VMware Horizon View and Citrix XenDesktop, supporting storage for thousands of VDI users on a single CS-Series array and occupying only three rack units (3U) of space.



Leverage validated SmartStack solutions for VDI to accelerate deployments.



Support thousands of users on a single Nimble Storage array.

Summary

VDI has unique storage characteristics that traditional storage architectures were never designed to support. However, the flash-optimized hybrid Nimble Storage architecture provides a solid platform for virtual desktop infrastructure environments. Nimble Storage provides adaptive performance that eliminates extensive management of traditional storage architectures without the cost of flash-only solutions.

Nimble Storage zero-copy cloning optimizes the storage footprint, using compressed, shared blocks that reduce deployment and management costs. At the same time, integrated data protection provides instant snapshot backup of VDI images and user data and replication of critical files such as golden images and user data, eliminates traditional backup overhead, and facilitates affordable disaster recovery.

Nimble Storage SmartStack takes the guesswork out of VDI by leveraging the most efficient storage solution for VDI implementations along with the best-of-breed Cisco UCS compute and networking and Citrix and VMware VDI software. Get started on your journey to a virtual end-user environment by contacting your local Nimble Storage representative or authorized value-added reseller. To learn more, visit www.nimblestorage.com or connect with us on our online community connect.nimblestorage.com.



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



© 2013 Nimble Storage, Inc. All rights reserved. CASL is a trademark or registered trademark of Nimble Storage. All other trademarks are the property of their respective owners. WP-FOS-VDI-0613