

BEST PRACTICES GUIDE

Nimble Storage for VMware View VDI



Overview

Virtualization is an important infrastructure technology in many IT solutions. The scope of virtualization initiatives has grown dramatically over the recent years. IT projects have gone from a few – often test and development – servers to virtualizing mainstream applications – like Microsoft Exchange and SQL. The latest push on leveraging virtualization technologies is to begin replacing the desktop and laptop environments in an organization's IT infrastructure. This area is called VDI (Virtual Desktop Infrastructure) and presents a new set of challenges, tools and solutions to address the successful virtualization of this part of a company's infrastructure. The VDI solution is the focus of this guide. In particular, we will be examining the VMware View and Nimble Storage solution.

The challenges for storage solutions to support a VDI environment stem from two key aspects of the overall virtualization architecture that gets deployed. As hundreds or even thousands of end-user desktops and laptops are centralized into the IT data center, the individual workloads need to be merged and user-experience maintained. A positive end-user experience depends on storage performance, and picking the right storage for VDI is a critical factor for success.

The workloads present an interesting challenge in that there are common everyday tasks that need to be performed by the user – let's call this steady-state – as well as operational tasks that need to be performed that can disrupt a shared environment. Things like booting, logging in, running AV scans, performing updates, or running backups of the users systems can introduce storage I/O stress.

The typical measure of throughput (MB/s) or IOPS of the storage layer may not be sufficient to understand the right solution. To ensure end-user experience is maintained, the latency or response time of the storage system becomes a critical performance measure as well.

In short, what is needed is an adaptive storage solution than can handle peak loads of disruption or activity as well as service the steady-state operations in a cost effective and easily administered package. Nimble Storage provides such a solution for VDI environments.

Audience

This guide is intended for solution architects, storage engineers, system administrators and IT managers who want to understand how to best deploy VMware View based VDI solutions using Nimble Storage. It is assumed that the reader has a working knowledge of VMware vSphere system administration, Microsoft® Windows® desktop administration, iSCSI SAN network design and basic Nimble Storage operations. Some knowledge of VMware View desktop management solutions is useful.

Nimble Storage Solutions

Nimble has developed an entirely new approach to storage that seamlessly combines SSDs with high-capacity drives. Nimble's breakthrough [CASL™ architecture](#) is designed from the ground up to address the performance and cost challenges of VDI.

Nimble Storage iSCSI arrays provide a complete application-aware data storage solution that includes flash accelerated primary storage, cost effective capacity, instant application-aware backup, and replication. This allows you to consolidate management of primary, secondary, and off-site disaster recovery storage within a single storage solution.

A Nimble Storage array provides iSCSI target volumes (LUNs) to VMware hosts. Volumes created on Nimble Storage arrays are highly optimized for virtual machines by providing the following benefits:

Performance and Capacity Efficiency

- CASL's innovative use of flash enhances read performance and addresses boot-storm and login-storm situations (i.e., multiple users booting or logging in at the same time). Serializing random writes enhances write performance – effectively meeting peaks in write IO's triggered by virus scans and OS upgrades.
- CASL compresses all data and provides savings 30-75% capacity savings. Zero-copy cloning ensures duplicate images don't consume valuable disk space.
- Most importantly, Nimble can respond rapidly to workload fluctuations. In comparison, traditional tiered architectures require hours or even days to respond to changes in IO patterns.
- Thin Provisioning: Efficiently stores actual data written rather than reserved space.

Integrated Data Protection

- Highly efficient snapshots allow frequent recovery points with 60-90 days of retention of virtual desktop images and user data. Backups take seconds and do not impact application or storage performance.
- Zero-Copy Cloning: Preemptive de-duplication to eliminate storage footprint of repetitive data.
- Built-in replication ensures VDI data is protected off-site in the event of a site failure.

Pain-Free Operations

- The intuitive Nimble UI eliminates the complexity of provisioning and managing separate storage, backup, and disaster recovery devices.
- The Nimble vCenter plug-in simplifies the process of provisioning, replicating, backing up, monitoring, and managing virtual desktops.

- Frequent heartbeats monitor system health to identify potential failures before they occur. High availability and resiliency features combined with truly non-disruptive upgrades virtually eliminate downtime.

VMware View VDI

Virtual desktop solutions can range from dozens of desktops (<100) to hundreds or even thousands of user desktops in one infrastructure. Whether adding the desktop infrastructure to an existing Nimble Storage based virtualization solution or building a dedicated VDI deployment from the ground up, many of the considerations for a successful deployment will be the same.

When looking at the task of virtualizing the desktop systems of an organization, there are many things to take into account. The quantity and variance of all of the different end-user systems to include as well as the direct interaction with the user population make this a challenging job.

One of the driving values of using a desktop virtualization approach for consolidation is that each user desktop and related experience is deployed on a stand-alone virtual machine and dedicated OS instance. This is in contrast to terminal server based solutions that combine multiple users onto a single shared OS platform. Using a VDI approach lends itself to less conflict between applications, user settings and workloads for each of the users. With a full private copy of Windows, there is a reduced chance for one user to impact another user's experience or platform stability.

The task of defining the desktop virtualization solution must deal with several diverse areas that include:

1. User experience and accessibility
2. Application support and licensing
3. Server resource management
4. Network connectivity
5. Storage management

This paper will focus on the storage aspects of the VMware View VDI solution.

A VDI environment has the following basic components:

- Servers to run the virtual desktops
- Networking to the users, other servers and the storage
- Storage to hold the infrastructure, desktops, and user data
- Operating systems and applications
- VDI components to manage and run the desktop virtualization operations

Figure 1 shows the layers involved in a typical VDI architecture. For some environments, there may be multiple ESX servers connected to a single Nimble Storage array supporting larger collections of desktop users.

Sizing

One of the key considerations for VDI deployments is proper sizing of the total solution to meet the user needs and expectations as well as budget and operational needs of the IT organization.

Sizing exercises to best understand your entire VDI deployment are outside the scope of this paper. Instead we will focus on the key considerations needed to best architect the appropriate storage solution to support the VDI configuration.

When looking at the measurements to help size and organize the storage for the desktop, many factors need to be considered. It is a combination of storage IOPS, throughput, latency and capacity. In addition, the variance in I/O profile of the workloads from a large collection of Windows desktops will have to be considered. While there are no hard and fast bounds for these measurements, we have found there are some characteristics useful for selecting and architecting the best storage solution. The compound effect of a large number of desktop users on the storage is also affected by the types of users, how they are connected and what other IT desktop management operations need to be supported.

For each desktop user, we have seen a range in IOPS from a low steady-state level of 3-5 per user to a high boot/login requirement of 50-100. For example, a simple, single Windows7 user login event can generate on the order of 2500-3000 low-level I/O operations. In each case, the VDI layers (server/network/storage) will need to expand to absorb the momentary workloads as well as support the target steady-state situations. Proper management of the desktop VM state and effective user profile control can help alleviate some of the extreme fluctuations in the workloads that must run on the servers and storage layers.

As an example, the following graph depicts the workload variations that could occur during a typical day in a VDI environment of hundreds of desktop VMs. There is a fairly wide spread in performance needs between steady-state operations during the course of the work day and the potential peak workloads that must be accommodated by the storage solution selected. The Nimble Storage CASL architecture is well suited to adapting to this type of shifting workloads.

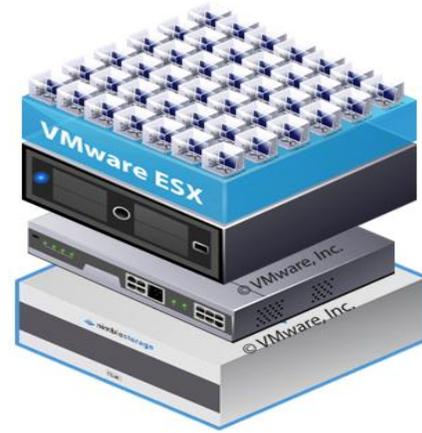


Figure 1: Basic VDI Architecture

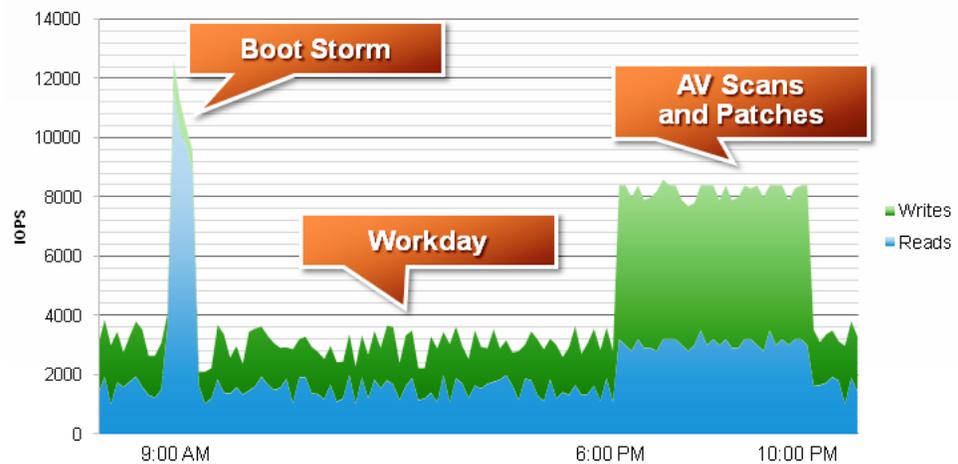


Figure 2: Typical VDI Day

To support a VDI workload, the storage solution needs to deal effectively with a very random workload of I/O activity that ranges in size from 4K-64K (or greater) and has a high write or read activity levels at different times. The following table: Workload Examples, lists two different example desktop user workloads and the projected throughput and rate requirements they might require.

User Type	IOPS (Avg)	IOPS (Peak)	Throughput (Avg)	Throughput (Peak)
Light	5	20	0.3 MB/s	0.8 MB/s
Heavy	20	50	0.6 MB/s	1.6 MB/s

Table 1: Workload Examples

Throughput requirements for the VDI desktop storage components are tied more to operating system, application requirements and temporary space and not related to accessing user data – which by design should be moved out of the virtual desktop and into centralized, shared user resources. The operating system and application (e.g., Microsoft Word) throughput per user is quite low and many times can be measured in KB/s rates.

Low latency or fast I/O response times is another key measure of a good user experience, and as such is the key element and design goal of building a good desktop virtualization solution. Latency of both reads and writes must be accounted for in the storage layer. Observations and analysis of actual user desktop workloads show read/write ratios to be about 50/50 during steady-state situations, although in some environments, write rates have been seen as high as 80% of all I/O. Boot/login storms tend to be much more read intensive.

Note: latencies are usually measured at the user level in seconds and at the storage level in milliseconds. The VDI solution should target latency results of less than 1-2 seconds for most user operations and less than 20 msec for the majority of operations at the storage layer.

The Nimble Storage CASL (Cache Accelerated Serialized Layout) architecture is unique and shows that with the right mix of technology, both performance and capacity can effectively be tied together. The Nimble Storage architecture provides a platform uniquely suited for random I/O workloads with good response times for reads through caching and effective write response times through write coalescing. For more details, visit the Nimble Storage web site at <http://www.nimblestorage.com/products>.

The Nimble Solution breaks the barrier between measuring and configuring storage in terms of spindles and focuses more on service levels (IOPS, throughput, latency) of the storage layer. With the Nimble Storage solution, we do not have to worry about the specific physical disk or RAID layout of the storage. All Nimble Storage volumes benefit from the CASL architecture. This simplifies storage configuration decisions and tasks. When setting up Nimble Storage volumes, only a few policy oriented criteria need to be addressed. Focus can be on management and measurability with less concern on layout to address performance issues.

The main thing to address with the storage implementation is the organization of the pieces of the desktop storage components. These pieces include: iSCSI LUNs (Nimble Storage Volumes), VMware VMFS datastores, VM templates, and the virtual machines themselves (VMDK). The templates and virtual machines will define the underlying virtual disk structure depending on how they are deployed.

With VDI solutions, we have to deal with an order of magnitude or more of components (e.g., desktops) to be virtualized. A good VDI design will strive to reach a balance between the number of desktop components related to any one storage component and the proliferation of too many pieces of the storage infrastructure to manage. The goal is to make sure not to overload a dependency on any single part (e.g., server, network, datastore, template, pool, etc.) or create too many points of management that reduce operational benefits. One aspect of this management challenge is the provisioning of storage volumes to the VMware environment.

Configuring the volumes on the Nimble Storage array follow the same basic guidelines as provisioning Nimble storage for any other part of a VMware virtualization solution. The choices are pretty straightforward. The performance policy should be set to VMware which automatically picks the optimal block size and enables compression and caching. Nimble volumes are thin provisioned by default so they do not need to be treated as thin provisioned inside of VMware. Volume thresholds and reserves of 10-20% can be used to provide notification from the Nimble Storage system although most of the management and monitoring will be performed from the VMware infrastructure (e.g., ESX, Virtual Center, View Administrator, Nimble Virtual Center plug-in, etc).

The Nimble volume is provisioned into the VMware environment as a datastore with appropriate multipath settings based on the network connectivity of the ESX host. In most VMware environments, Nimble Storage recommends using Round Robin MPIO for all Nimble Storage volumes. The Nimble volume to VMware datastore relationship is one to one. For more detail on this refer the online resources at: <http://www.nimblestorage.com/resources/overview>.

The next key consideration is what goes into the datastore. For VDI implementations, this will be desktop VMs – lots of them.

The storage provisioning step determines the basic partitioning of the desktop environment. The simplest approach would be to create one large datastore and put all of the desktops in a single container. In practice, this only works well for smaller environments typically less than 50 desktops.

As the size of the environment grows, it makes sense from a management, monitoring, resources and recoverability perspective to break things down into more manageable working sets.

Our current guidelines are to build VMFS datastores from single Nimble Storage volumes to hold up to 50 full size desktop virtual machines. A full size desktop is the size of the image if there were no compression, thin provisioning or linked cloning. For example, a standard VDI desktop size of 20GB is common. Because of the benefits of Nimble Storage compression and thin provisioning, it is possible to over allocate storage within conservative bounds. For example, if we provision a 1TB volume and only end up needing a few hundred GBs of actual storage, there is no waste as the storage was not consumed. If however, the VMs grow to their potential defined disk sizes, or snapshots (VMware or Nimble) are used then we have extra room to accommodate the needed space. Figure 3 shows the information related to a specific datastore used for one of the VMware View Pools.

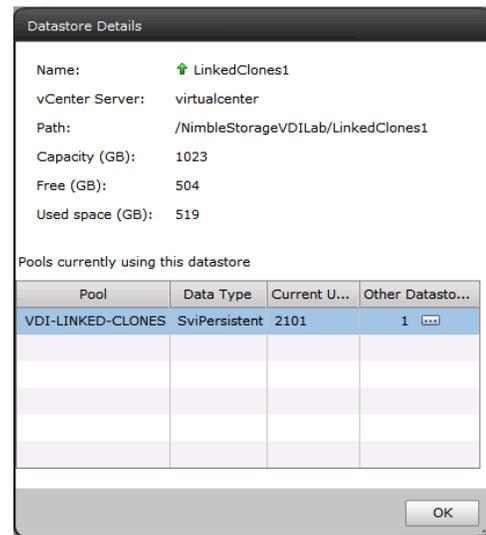


Figure 3: VMware View Datastore Details

The only way to scale the deployment and management of dozens to hundreds of desktop virtual machines is to use some form of VM cloning. Much like the physical world, where a golden image is created and desktops and laptops are deployed to users from a standard base image, VMware View provides a robust management solution for cloning the desktop virtual machine within the VDI environment. There are different methods for producing these cloned virtual machines.

A clone is a copy of an existing virtual machine. When the cloning operation is complete, the clone becomes a separate virtual machine with unique identity of its own. A full clone is an independent copy of a virtual machine that shares nothing with the original virtual machine from which is created once the cloning operation is complete. The new VM has the same storage details as the parent VM.

VMware provides an alternate method of cloning the VM called a Linked Clone. In this approach, the new VM is a copy of a virtual machine that continues to share virtual disks with its parent VM. The parent VM (and its snapshot) represent a VM Replica that is used to link to the new VM. The differential—the data that is unique to the Linked Clone—is stored in a separate location.

With this approach, the linked clone can occupy a smaller footprint of only changed data but always requires access to the parent (Replica) disk in order to operate. As a result, VMware linked clones provide a way to reduce some of the storage requirements and define efficient processes for updating the desktops en masse. For VMware Linked Clone deployment, we recommend keeping the Master

Replica and the Linked Clone in the same datastores to allow for better management of the protection policies and easier measurability of the capacity and performance.

For each clone produced by either method, the VM must be properly prepped for inclusion into the set of desktop virtual machines (e.g., name, IP, licensing, etc.). This is automated through the VMware

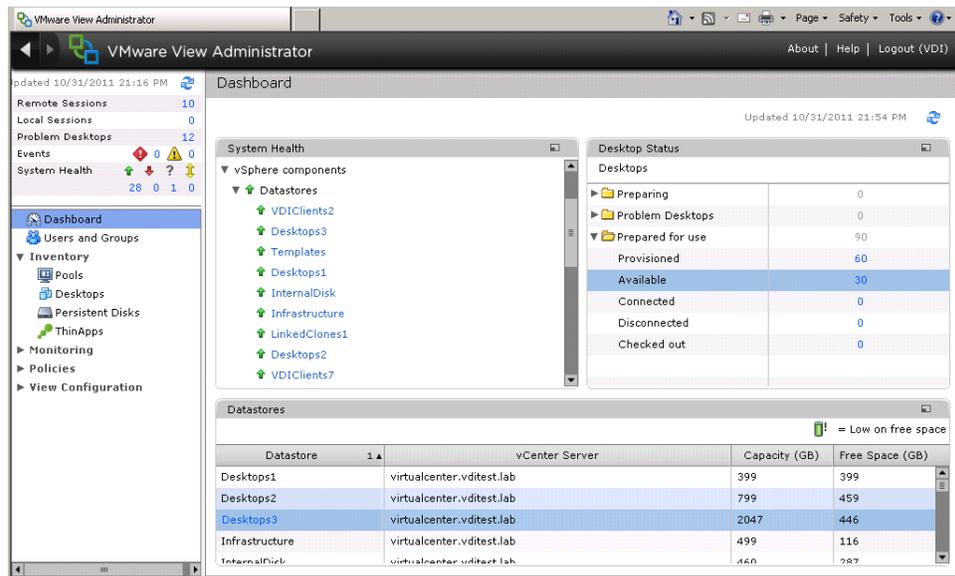


Figure 4: VMware View Dashboard

Virtual Center and VMware View Administrator Capabilities

Figure 4 shows the high level view into managing the pools, desktops and datastore components making up the desktop configurations. From this panel you have a quick overview of the status and management points of the desktop deployment automation functions.

Another consideration for storage provisioning is “purpose”. Separate datastores for infrastructure components (e.g., AD, Virtual Center, applications, user shares, etc) are recommended. In addition, we recommend a separate datastore and Nimble Storage volume for templates and Linked Clone golden image desktops. The reason for this separation does not relate to performance but more to protecting and monitoring the environment. Nimble volume protection policies can be set differently for each volume or volume collection. Critical data like user data, templates and golden images can be protected more aggressively and other volumes like linked clones can be protected less aggressively. This may be acceptable for data that can be easily reconstructed like a linked clone. Separating the datastores by functionality will also help with monitoring the environment. It is easier to spot consumption, bottlenecks and resource utilization when they can be monitored individually.

Another reason for breaking down the storage volumes at this level into workable size datastores is to facilitate the desktop instance management capabilities of VMware View. Within VMware View, the Desktop Pool is a basic method of organizing the production and control of the sets of desktop

instances. The key aspect of the Desktop Pool that ties to the storage layer is the assignment of datastore(s) to hold the desktop VMs.

A Desktop Pool can use more than one datastore and datastores can be shared between Desktop Pools. This provides a convenient method of organizing and scaling the storage under a larger collection of desktop images. If more storage is required then add another datastore to the Desktop Pool. If another type or group of desktops is desired, they can share a datastore until it becomes full. Figure 5 shows the panel used to access the defined pools in a VMware View environment. From here you can add or edit pool properties, controlling the specifics of how desktops are created.

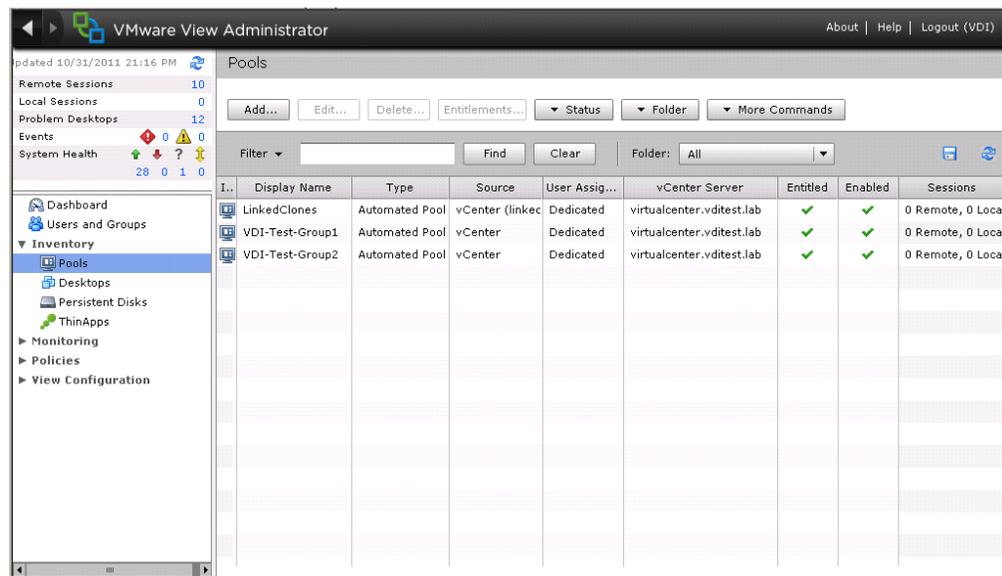


Figure 5: VMware View Desktop Pools

The desktop pool can be used to differentiate different types of desktop users. If the profile of the user is something you may want to monitor independently then associate that pool with one or more specific datastores. For Linked Clones, we suggest using a separate datastore for the Master Replicas and one or more datastores for the Linked Clone portion. For the desktop image, this will place some of the static data (e.g., OS) in one location and the delta (user profile) data in another.

The number of desktops managed in a single pool can be large – as high as 1000. We recommend using multiple datastore collections as the Desktop Pool size grows following the guidelines above.

Protection and Recovery

Any production environment needs proper and adequate data protection methods implemented. The Nimble Storage solution has built in highly efficient snapshot mechanisms as well as replication capabilities to other Nimble Storage arrays.

It is important to make sure critical infrastructure elements are protected with local snapshot protection points and ultimately remote (replicated) copies. These include desktop templates, golden images, and any parts of the VMware environment such as the Virtual Center or View Connection Manager.

With respect to individual desktops, sometimes it can be easier and faster to simply rebuild the single desktop system when problems arise. The VMware View provisioning methods for the Desktop Pools can help automate some of this for certain types of desktop deployment and assignment models.

Snapshots

Snapshots are one method of providing a consistent recovery point within the local storage system. These snapshots can be within the VMware environment, within the storage layer, or both.

We recommend using Nimble Storage snapshots through the Protection Policies for the critical data items as described above. Using the Nimble Protection Manager capabilities, these snapshots are consistent within the VMware environment.

Because of the large number of desktop VMs that may be associated with a single datastore (Nimble volume), we do not recommend using the VMware Virtual Center integrated snapshot methods on these volumes. Simple storage snapshots should suffice for crash consistent protection of the desktop pools. VMware consistent storage snapshots require a VMware Virtual Center invoked snapshot of each VM in the targeted volume. If such snapshots are desired, make sure they are scheduled during periods of little or no desktop activity in those pools.

Replication

The replication of Nimble Storage volumes to another Nimble Storage array provides another level of infrastructure protection. We recommend following VMware replication best practices for your key VDI volumes just like you would for any other critical virtualization data set.

If replication is for DR purposes, keep in mind that desktop integration into the DR site may need to be managed through a tool like VMware Site Recovery Manager (SRM) which is fully integrated into with Nimble Storage replication capabilities through an adapter. This is another situation where critical components are replicated and the VMware View environment is leveraged for recreation of elements in the new context. Work with your VMware resources to determine which method will provide the greatest flexibility and recoverability of your desktop infrastructure.

Summary

When using Nimble Storage for VMware View VDI deployments, following these simple configuration guidelines allows you to get the most out of your investment:

1. Provision Nimble volumes to VMware datastores using the default VMware performance policy which picks optimal block size, and enable both caching and compression (defaults).
2. Use Round Robin MPIO setting for Nimble Storage iSCSI LUNs configured as VMFS Datastores.
3. Size volumes and VMFS datastores for around 50 full size desktops. This will typically be about 500GB using 10-20GB per desktop as a guide.
4. Use VMware View Linked Clones for View Pool configurations to save additional space and speed desktop provisioning.
5. Keep Master Replicas and Linked Clones in the same datastores (Nimble Storage volumes).
6. Use VMware View Desktop Pools to partition the desktops by type, location, or manageability. Keep the ratio through VMware View Pools to around 50-100 desktops per pool. Note: for Linked Clones this will also keep the Master Replica – Linked Clones ratio to the same ratio as the Pool.
7. Assign multiple datastores to a VMware View Pools to expand capacity for the pool if needed.
8. Define protection templates for the critical volumes that cannot be easily redeployed with VMware View mechanisms. These include: VM templates, golden images, user data, critical user's desktops, and infrastructure components.



Nimble Storage, Inc.

2740 Zanker Road, San Jose, CA 95134

Tel: 408-432-9600; 877-364-6253) | www.nimblestorage.com | info@nimblestorage.com

© 2012 Nimble Storage, Inc. All rights reserved. CASL is a trademark or registered trademark of Nimble Storage, Inc.

WP-VDI-0812