HPE Nimble Storage Deployment Considerations for VMware Horizon View

Contents

Notices

The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Confidential computer software. Valid license from Hewlett Packard Enterprise required for possession, use, or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Links to third-party websites take you outside the Hewlett Packard Enterprise website. Hewlett Packard Enterprise has no control over and is not responsible for information outside the Hewlett Packard Enterprise website.

Acknowledgments

Intel®, Itanium®, Pentium®, Intel Inside®, and the Intel Inside logo are trademarks of Intel Corporation in the United States and other countries.

Microsoft® and Windows® are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

Adobe® and Acrobat® are trademarks of Adobe Systems Incorporated. Java® and Oracle® are registered trademarks of Oracle and/or its affiliates.

UNIX® is a registered trademark of The Open Group.

Publication Date

Wednesday March 7, 2018 16:04:30

Document ID

xsj1492018119413

Support

All documentation and knowledge base articles are available on HPE InfoSight at *https://infosight.hpe.com*. To register for HPE InfoSight, click the *Create Account* link on the main page.

Email: *support@nimblestorage.com*

For all other general support contact information, go to *https://www.nimblestorage.com/customer-support/*.

Overview: The Challenges of VDI

Virtualization is an important infrastructure technology in many IT solutions. In recent years, the scope of virtualization initiatives has grown dramatically. IT projects have gone from virtualizing a few servers (often, ones used for test and development) to virtualizing mainstream applications such as Microsoft Exchange and SQL Server.

In the latest push to leverage virtualization technologies, organizations are beginning to replace the physical desktop and laptop environments in their IT infrastructure with virtual desktops. This virtual desktop infrastructure (VDI) presents new challenges that require new tools and solutions to address the successful virtualization of this part of a company's infrastructure.

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

The range of workloads that VDI environments must be able to handle also presents a challenge. Users must perform common, everyday tasks (the steady-state workload). In addition, administrators must carry out operational tasks that can disrupt a shared environment. Processes such as booting, logging in, running antivirus scans, performing updates, or running backups of users' systems can introduce storage I/O stress.

The typical measure of throughput (Mb/s) or IOPS of the storage layer might not be sufficient to clarify what the right solution is. To ensure that the 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 a storage solution that can handle peak loads of disruption or activity and steady-state operations in a cost-effective and easily administered package. HPE Nimble Storage arrays provide such a solution for VDI environments.

Audience

The VDI solution described in this guide uses VMware Horizon View software with HPE Nimble Storage arrays. Hewlett Packard Enterprise (HPE) customers and partners can use the information in the guide to take advantage of an infrastructure that is built to deliver IT efficiency and enable IT innovation. The intended audience includes, but is not limited to, sales engineers, field consultants, professional services staff, IT managers, partner engineering, and customers.

The guide assumes that the reader has a working knowledge of the following areas:

- VMware vSphere system administration
- Windows desktop administration
- SAN design
- Basic HPE Nimble Storage operations

Some knowledge of VMware Horizon View desktop management solutions will also be useful.

VMware Horizon View and the Goals of a VDI Solution

Virtual desktop solutions can range from dozens of user desktops (<100) to hundreds or even thousands of desktops in a single infrastructure. Whether you add the desktop infrastructure to an existing HPE Nimble Storage based virtualization solution or build a dedicated VDI deployment from the ground up, many of the considerations for a successful deployment are the same.

The plan for virtualizing an organization's desktop systems has to take many variables into account. You must consider the quantity and variance of all of the different end-user systems to include as well as the direct interaction with the user population, which makes this a challenging job.

One driving value of using a desktop virtualization approach for consolidation is that each user desktop and the related user experience is deployed on a standalone virtual machine (VM) and a dedicated operating system (OS) instance. VMware Horizon View delivers desktops and applications through a single platform. It transforms static desktops into secure, digital workspaces that can be delivered on demand. You can provision virtual or remote desktops and applications through a single VDI and application-virtualization platform to streamline management and easily entitle end users.

Figure 1: Virtual desktop experience

The task of defining a desktop virtualization solution must account for diverse areas:

- Safeguarding the user experience and ensuring accessibility
- Providing application support and licensing
- Managing the servers that run the virtual desktops
- Providing network connectivity to the users, the other servers, and the storage
- Managing the storage that holds the infrastructure, the desktops, and the user data

Figure 2: Horizon View topology

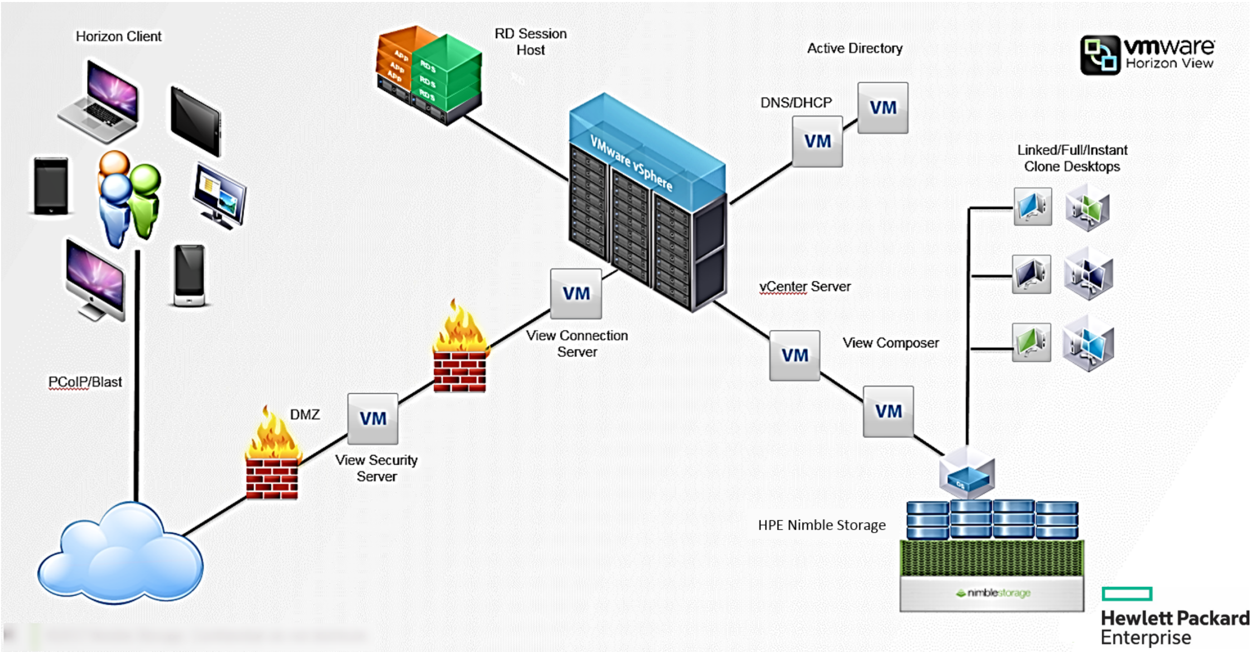

The topology diagram showcases a typical Horizon View deployment:

- Users can log in from any mobile device or computer that uses the Horizon Client software.
- For additional security, users enter through a security server that is placed inside a DMZ.
- User credentials are validated against Microsoft Active Directory before the user's appropriate assigned resources are presented to the Horizon client.
- When the user clicks a resource, such as the Windows 10 desktop, the connection server brokers the user to a Windows 10 desktop from the vCenter and the corresponding VMware ESXi host.

Depending on the provisioning methodology that is used, the desktop can be either a full, instant clone or a linked clone. HPE Nimble Storage arrays provide the datastores (the volumes) for all infrastructure servers (connection servers, VMware View Composer servers, and database servers) as well as for desktop VMs.

HPE Nimble Storage Technical Overview

All-Flash Array Hardware

The HPE Nimble Storage AF-Series array family starts at the entry level with the AF1000 model and expands through the AF3000, AF5000, and AF7000 models up to the AF9000 model at the high end. All-flash arrays can be upgraded nondisruptively from the entry level all the way to the high-end array model.

For more information, see *HPE Nimble Storage All Flash Arrays*.

Figure 3: The HPE Nimble Storage all-flash array family

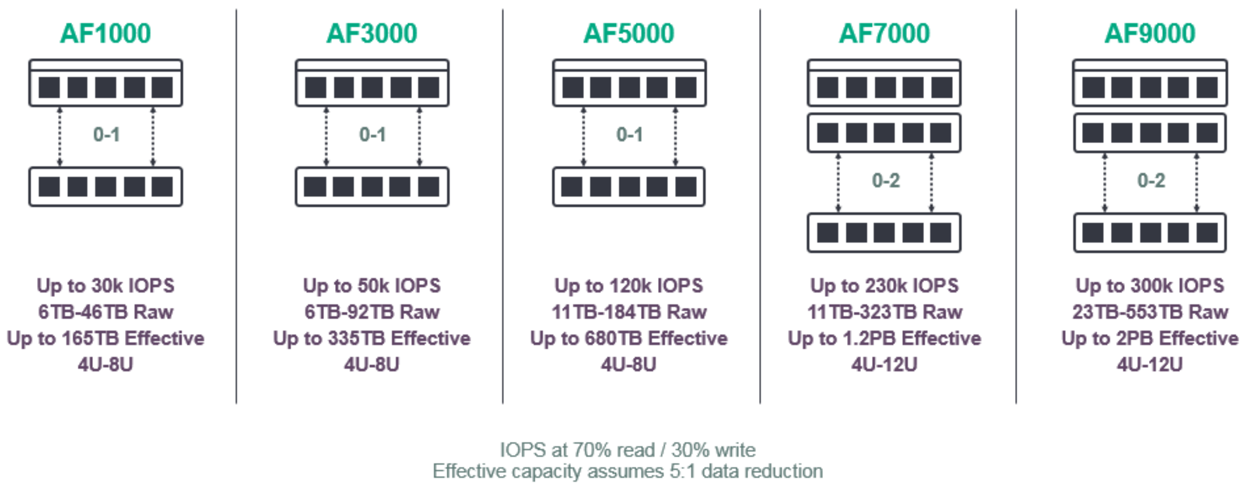

Figure 4: All-flash array front view

Figure 5: All-flash array rear view

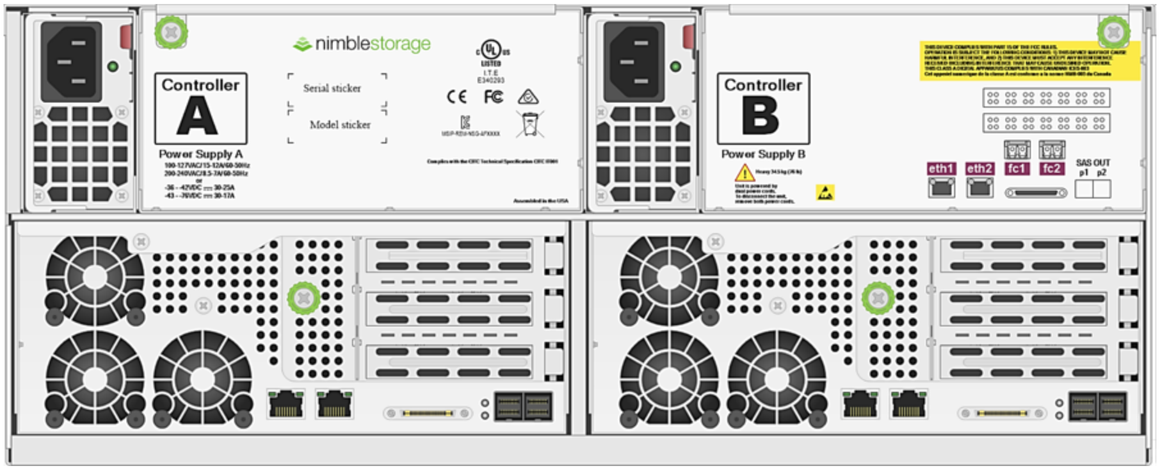

Predictive Flash Platform

The HPE Nimble Storage predictive flash platform enables enterprise IT organizations to implement a single architectural approach that dynamically caters to the needs of varying workloads. Driven by the power of predictive analytics, predictive flash is the only storage platform that optimizes across performance, capacity, data protection, and reliability within a dramatically smaller footprint.

Predictive flash is built on the HPE Nimble Storage Cache Accelerated Sequential Layout (CASL) architecture, the NimbleOS operating system, and HPE InfoSight, the HPE cloud-connected predictive analytics and management system. CASL and HPE InfoSight form the foundation of the predictive flash platform, which enables the dynamic and intelligent deployment of storage resources to meet the growing demands of business-critical applications.

Universal Hardware Architecture

HPE Nimble Storage arrays are built on a universal hardware platform with modular components. All components, including controllers, can easily be upgraded nondisruptively by the customer or an HPE representative. The universal hardware architecture spans both all-flash and adaptive flash arrays, giving HPE customers maximum flexibility and reuse of their array hardware.

Figure 6: HPE Nimble Storage universal hardware architecture

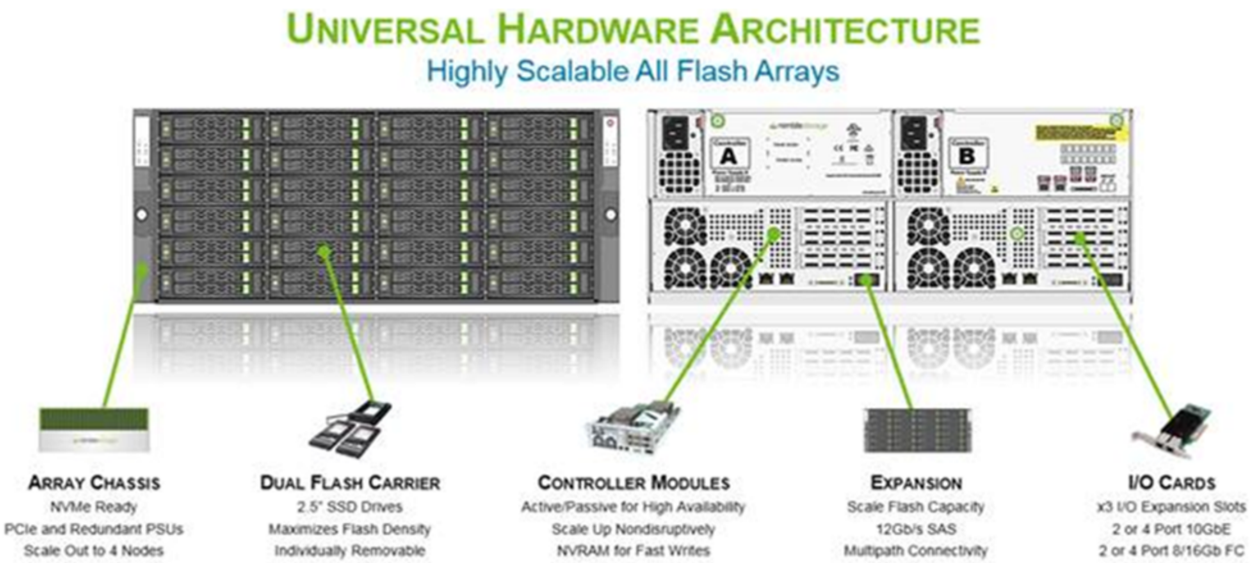

NimbleOS Architecture

The NimbleOS operating system is based on the patented HPE Nimble Storage CASL architecture. CASL leverages the unique properties of flash and disk to deliver high performance and capacity—all within a remarkably small footprint. CASL scales performance and capacity seamlessly and independently.

Unified Flash Fabric

The HPE Nimble Storage unified flash fabric is a single consolidated architecture that enables flash for all enterprise applications. In the past, enterprises have been forced to choose between hybrid flash and all-flash arrays. That is no longer the case with the HPE Nimble Storage unified flash fabric. The unified flash fabric enables flash for all enterprise applications by unifying HPE Nimble Storage all-flash and adaptive flash arrays into a single consolidated architecture with common data services.

Figure 7: HPE Nimble Storage unified flash fabric

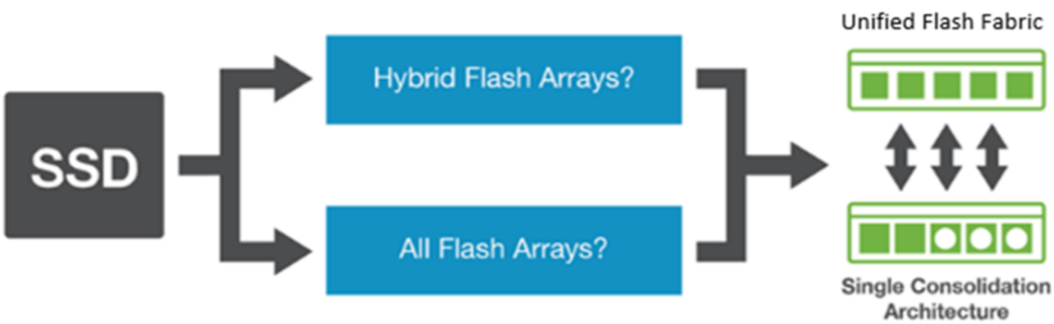

The unified flash fabric provides many significant benefits:

- **Thin provisioning and efficient capacity utilization.** With CASL, capacity is consumed only as data is written. CASL efficiently reclaims free space on an ongoing basis, preserving write performance with higher levels of capacity utilization. This strategy avoids the fragmentation issues that hamper other architectures.

- **Accelerated write performance.** After writes are placed in NVDIMM (made persistent and mirrored to the passive partner controller), they are acknowledged back to the host and sent to solid-state disk (SSD) at a later time (generally when there is a full stripe to be written). As a result, writes to an HPE Nimble Storage array are acknowledged at memory speeds.
- **Maximized flash write cycles.** By sequencing random write data, NimbleOS sends full stripes of data to SSD. Compressing and deduplicating the data inline minimizes the data footprint on the disk. In addition, because the data sent to disk is of variable block size, NimbleOS does not have to break it into smaller, fixed-sized chunks to be placed on SSD. Therefore, it is sent to SSD efficiently. This efficiency enables HPE Nimble Storage arrays to maximize the deployable life of a flash drive by minimizing write wear on the flash cells.
- **Read performance.** Because all reads come from SSD, NimbleOS and HPE Nimble Storage all-flash arrays deliver submillisecond read latency and high throughput across a wide variety of demanding enterprise applications.
- **All-flash arrays.** HPE Nimble Storage all-flash arrays use only SSDs to store data. As a result, all read operations come directly from the SSDs themselves, delivering extremely fast read operations. All writes are also sent to SSD, but because of the NimbleOS architecture and the use of NVDIMMs to store and organize write operations, all writes are acknowledged at memory speeds (just as with the HPE Nimble Storage adaptive flash arrays).

  HPE Nimble Storage all-flash arrays use triple-level cell (TLC) SSDs, which enable maximum flash storage density. Traditional SSD problems with write wear, write amplification, and so on are not an issue for the variable block NimbleOS architecture, which minimizes write amplification through its intelligent data layout and management in the file system.

- **Efficient, fully integrated data protection.** All-inclusive snapshot-based data protection is built into the adaptive flash platform. Snapshots and production data reside on the same array, eliminating the inherent inefficiencies of running primary and backup storage silos. With its intuitive dashboards and proactive notifications about potential issues, HPE InfoSight ensures that customers' data protection strategies work as expected.
- **Inline compression.** CASL uses fast, inline compression for variable-application block sizes to decrease the footprint of inbound write data by as much as 75%. When enough variable-sized blocks are accumulated to form a full write stripe, CASL writes the data to disk. If the data being written is active, it is also copied to SSD cache for faster reads. Written data is protected by triple-parity RAID.
- **Inline deduplication.** HPE Nimble Storage all-flash arrays include inline data deduplication in addition to inline compression. The combination of inline deduplication and inline compression delivers a comprehensive data reduction capability that allows NimbleOS to minimize the data footprint on SSD, maximize usable space, and greatly minimize write amplification.
- **Thin, redirect-on-write snapshots with SmartSnap.** HPE Nimble Storage snapshots are point-in-time copies that capture just-changed data, making it easy to store three months of frequent snapshots on a single array. Data can be restored instantly because snapshots reside on the same array as the primary data.
- **Efficient replication with SmartReplicate.** SmartReplicate sends only compressed, changed data blocks over the network for simple and WAN-efficient disaster recovery.
- **Zero-copy clones.** HPE Nimble Storage snapshots enable the quick creation of fully functioning copies, or *clones*, of volumes. Instant clones deliver the same performance and functionality as the source volume, which is an advantage for virtualization, VDI, and test or development workloads.
- **Application-consistent snapshots.** The HPE Nimble Storage VDI solution uses the Volume Shadow Copy Service (VSS) framework and VMware integration to create instant application-consistent and VM-consistent backups, using application templates with pretuned storage parameters.
- **Flexible data encryption with SmartSecure.** NimbleOS enables individual volume-level encryption with little or no performance impact. Encrypted volumes can be replicated to another HPE Nimble Storage target, and data can be securely shredded at the volume level of granularity.

Transparent Application Migration

Enterprises can deploy HPE Nimble Storage all-flash arrays, adaptive flash arrays, or a combination of both to meet the varying needs of all applications. Because both types of arrays run the same NimbleOS, management and functionality are identical, and arrays can be clustered and managed as one.

The transparent application migration feature enables the administrator to transparently migrate volumes that were set up on an adaptive flash array to an all-flash array or vice-versa. If, for example, infrastructure volumes are housed on an adaptive flash array, and they need to be migrated to an all-flash array, moving them through transparent application migration can reduce the time required to rebuild the entire infrastructure. VDI infrastructure volumes can be migrated across arrays without affecting the consistency of the infrastructure.

Figure 8: Transparent application migration

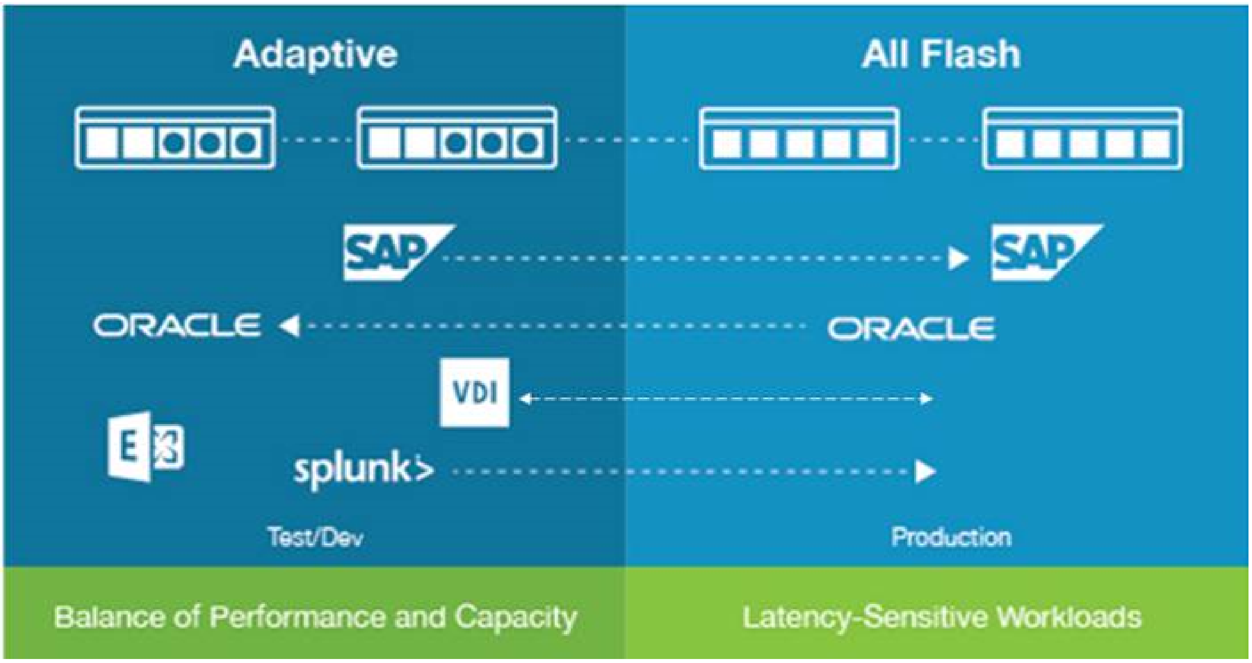

HPE InfoSight

HPE InfoSight leverages the power of deep data analytics to offer customers precise guidance on the optimal approach to scaling flash, CPU, and capacity for changing application needs while maintaining peak storage health. Using systems modeling, predictive algorithms, and statistical analysis, HPE InfoSight solves storage administrators' most difficult problems, ensuring that storage resources are dynamically and intelligently deployed to satisfy the changing needs of business-critical applications.

At the heart of HPE InfoSight is a powerful engine that applies deep data analytics to telemetry data that is gathered from HPE Nimble Storage arrays deployed across the globe. More than 30 million sensor values are collected per day for each HPE Nimble Storage array. The HPE InfoSight engine transforms the millions of gathered data points into actionable information that helps customers realize significant operational efficiency in many ways:

- Maintaining optimal storage performance
- Projecting storage capacity needs
- Proactively monitoring storage health and receiving granular alerts
- Proactively diagnosing and automatically resolving complex issues, which frees up IT resources for value-creating projects
- Ensuring a reliable data protection strategy with detailed alerts and monitoring

- Expertly guiding storage resource planning, by determining the optimal approach to scaling cache and IOPS to meet changing SLAs
- Identifying latency and performance bottlenecks through the entire virtualization stack
- Delivering transformed support experience from level-3 support

For more information, see *HPE InfoSight Predictive Analytics*.

VMVision for Hypervisor and VMware Monitoring

VMVision provides a granular view of the resources that are used by every VM connected to an HPE Nimble Storage array. This feature makes it possible to correlate the performance of VMs in a datastore with insights into hypervisor and host resource constraints such as vCPU, memory, and network.

VMVision helps in determining VM latency factors, whether from the storage, the host, or the network. It also helps in taking corrective action on noisy-neighbor VMs and reclaiming space from underused VMs. Every hour, through the heartbeat mechanism, the correlated statistics are sent to HPE InfoSight for processing. No additional host-side agents, tools, or licenses are necessary for the feature to work.

Figure 9: HPE Nimble Storage VMVision

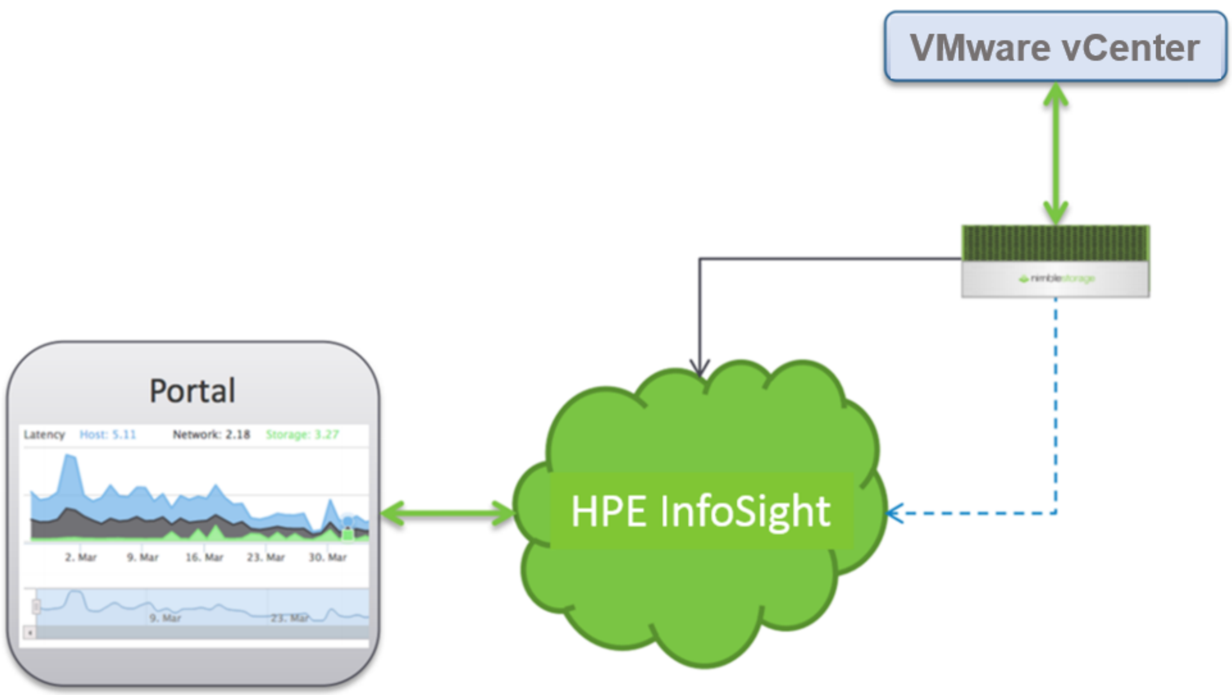

Deployment Considerations

You can leverage the following storage considerations when you deploy HPE Nimble Storage arrays in a VMware Horizon View environment.

Volume and Datastore Provisioning

The basic guidelines for configuring volumes on an HPE Nimble Storage array are the same as those for provisioning HPE Nimble Storage arrays in any other part of a VMware virtualization solution. The choices are straightforward:

- The performance policy should be set to VMware VDI, which automatically picks the optimal block size (4 KB) and enables compression, deduplication (all-flash array), and caching.
- HPE Nimble Storage volumes are thin provisioned by default.
- Volume thresholds and reserves of 10% to 20% can be used to provide notification from the HPE Nimble Storage arrays, although most management and monitoring is performed from the VMware infrastructure (for example, ESXi, vCenter Server, Horizon View Administrator, the HPE Nimble Storage vCenter plugin, and so forth).
- The HPE Nimble Storage volume is provisioned in the VMware environment as a datastore with appropriate multipath settings that are based on the network connectivity of the ESXi host.

Another key decision is to consider what should go into the datastore. A VDI implementation contains a great many desktop VMs. Although in some implementations and reference architectures thousands of desktops are provisioned from the same datastore (or volume) without suffering any limitations, this arrangement is usually not optimal. For ease of management, monitoring, resources, and recoverability, it makes sense to break the total number down into more manageable working sets.

Current HPE Nimble Storage guidelines recommend building VMware Virtual Machine File System (VMFS) datastores from single HPE Nimble Storage volumes to hold approximately 100 to 150 full-size desktop VMs. The alternative approach is to provision datastores (volumes) per Horizon View desktop pool. This method enables the volumes to be managed and backed up individually.

HPE Nimble Storage supports hosting the datastore replicas either on the same volume as the desktop VMs or on a different volume. However, there are two important guidelines:

- Specify only one separate replica datastore for a pool.
- If a replica datastore is shared, make it accessible from all ESXi hosts that have access to the desktop pool.

Figure 10: Storage policy management

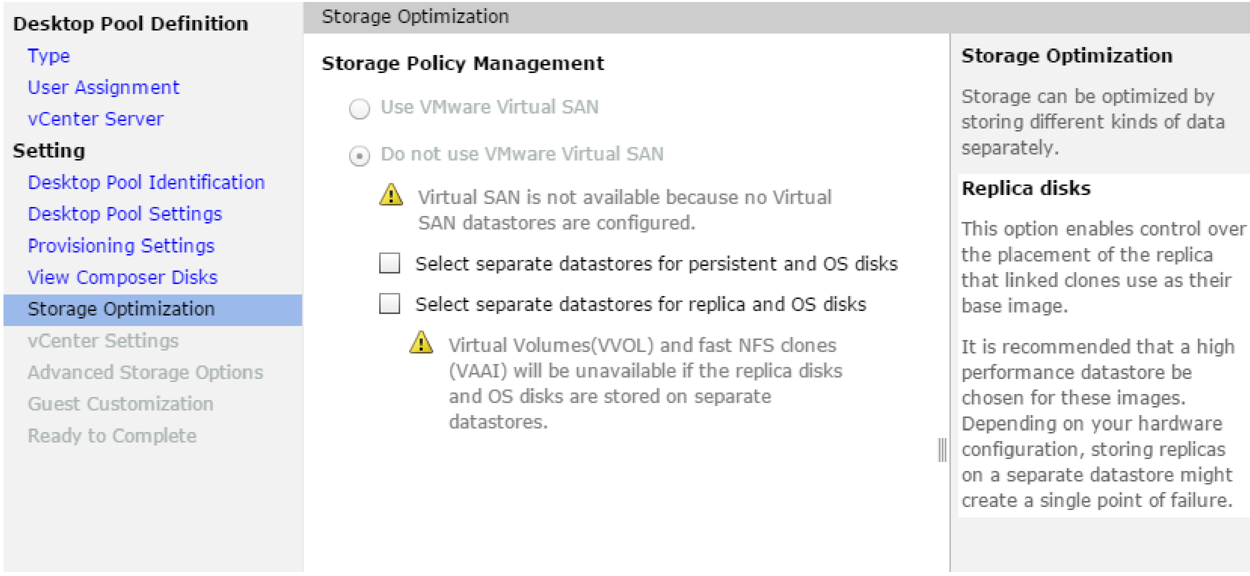

Given the benefits of HPE Nimble Storage data savings (deduplication, compression, and thin provisioning), it is possible to overallocate storage within conservative bounds. For example, if you provision a 2 TB volume and end up needing only a few hundred gigabytes of actual storage, there is no waste because the storage was not consumed. If, however, the VMs grow to their potential defined disk sizes, or if VMware or HPE Nimble Storage snapshots are used, then extra room is available to accommodate the needed space.

HPE recommends limiting individual datastores to just a few terabytes in size. Both lab testing and real-world experience have shown that a size of 2 TB to 4 TB is a good guideline to follow, and it matches well with the guidelines for numbers of desktop VMs per datastore.

Another way to provision storage is to categorize it by purpose. HPE recommends provisioning separate datastores for infrastructure components such as Microsoft Active Directory, vCenter Server, applications, user shares, and so forth.

HPE also recommends provisioning separate datastores and HPE Nimble Storage volumes for templates and linked-clone golden image desktops (not to be confused with the master replica). The reason is not for performance but for protecting and monitoring the environment. You can set volume protection policies differently for each volume or volume collection. For example, you might protect critical data such as user data, templates, and golden images more aggressively, and other volumes, such as linked clones, less aggressively. Less protection might be acceptable for data that can be easily reconstructed, such as a linked clone.

Separating datastores by functionality is also helpful in monitoring the environment. It is easier to spot consumption, bottlenecks, and resource utilization when you can monitor datastores individually.

Another reason to break down storage volumes at this level into workable-size datastores is to facilitate the desktop instance management capabilities of VMware Horizon View. In Horizon 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 datastores to hold the desktop VMs.

A desktop pool can use more than one datastore, and datastores can be shared between desktop pools. This flexibility provides a convenient method of organizing and scaling the storage under a larger collection of desktop images. If more storage is required, you can add another datastore to the desktop pool. If another type or group of desktops is desired, that group can share a datastore until it becomes full. You can use the desktop pool to differentiate different types of desktop users. If you want to monitor a user profile independently, you can associate that pool with one or more specific datastores.

The number of desktops that can be managed in a single desktop pool can be large—up to 2000. HPE recommends using multiple datastore collections as the desktop pool size grows and following the guidelines presented here.

Storage DRS

The VMware storage DRS (SDRS) feature does not support the use of VMware vSphere linked clones. Because the SDRS model does not detect the sharing relationship between two linked-clone VMs, there is potential for loss of space efficiency when moving linked clones. To protect the integrity of linked-clone VMs, SDRS rejects placement operations that specifically create linked clones.

Another problem is that the database entry for each linked clone includes the full paths of both linked VMs. If SDRS were to move the VMs, the paths would be broken. However, the native Horizon View rebalance feature offers an alternative way to move linked VMs.

Sizing

A key consideration for VDI deployments is the proper sizing of the total solution to meet both the needs and expectations of users and the budget and operational requirements of the IT organization.

Sizing exercises to help you best understand your entire VDI deployment are outside the scope of this guide. Instead, this paper focuses on key considerations for best architecting the appropriate storage solution to support the VDI configuration. For detailed sizing guidelines, see the following resources:

- *Nimble Storage Alert #: EXT-0036: Case Automation Updates for Proactive Hybrid Flash Array SSD Replacement based on Wear Level*
- *Sizing Questionnaire and Guidelines for VDI Implementations*

The storage requirements for an environment are affected by many factors. You can look at the following measurements to size and organize the storage for the desktops:

- Storage IOPS
- Throughput
- Latency
- Capacity
- Variance in I/O profile of the workloads from a large collection of Windows desktops

Although these measurements have no hard-and-fast boundaries, in testing we found that certain characteristics are 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 following factors:

- The types of users
- How they are connected
- What other IT desktop management operations need to be supported

During testing, for each desktop user, we saw a range in IOPS from a low steady-state level of 5 to 20 per user to a high boot/login requirement of 50 to 100. For example, a simple, single Windows 7 user login event might generate on the order of 2500 to 3000 low-level disk I/O operations. In each case, the VDI layers (server, network, and storage) must expand to absorb the momentary workloads while continuing to support the target steady-state situations. Proper management of the desktop VM state and effective control of user profiles can help alleviate some of the extreme fluctuations in the workloads that run on the servers and the storage layers.

The following graph shows an example of workload variations that might occur during a typical day in a VDI environment of hundreds of desktop VMs. The selected storage solution must accommodate a fairly wide spread in performance needs between steady-state operations during the course of the work day and potential peak workloads. The HPE Nimble Storage CASL architecture with NimbleOS is well suited to adapting to this type of shifting-workloads environment.

Figure 11: Typical VDI day

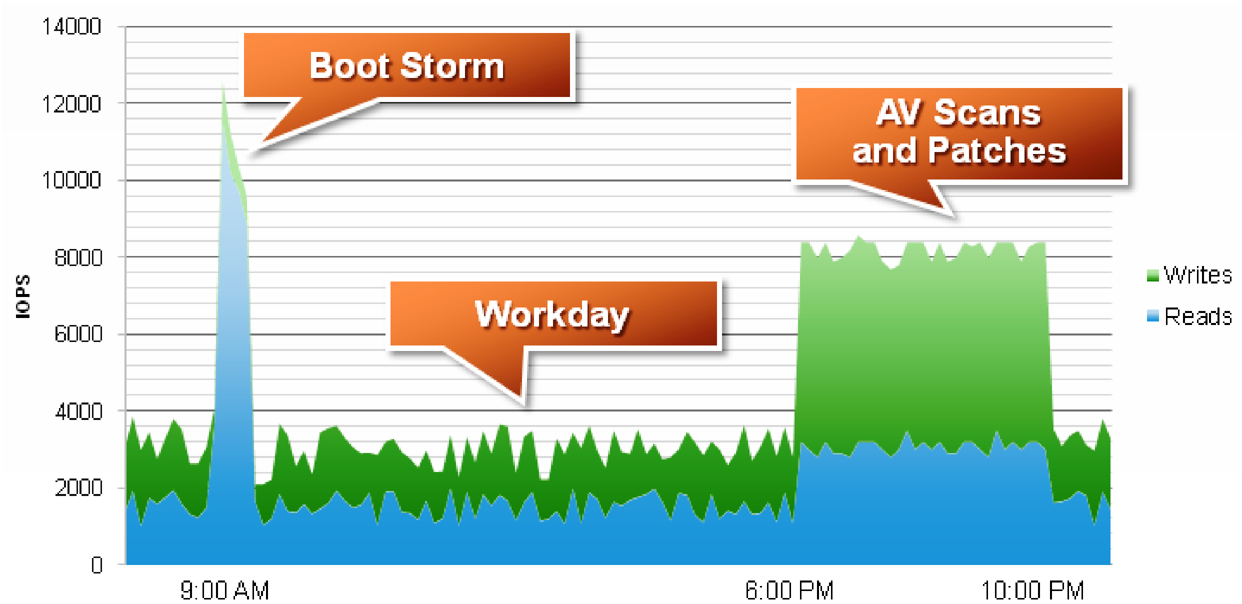

To support a VDI workload, the storage solution must deal effectively with a very random workload of I/O activity that ranges in size from 4 KB to 64 KB or greater and that has high write or read activity levels at different times. The following table lists the projected throughput and rate requirements for two examples of desktop user workloads.

Figure 12: Workload examples

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

Throughput requirements for the VDI desktop storage components are tied primarily to operating system, application requirements, and temporary space. They are 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 (for example, Microsoft Word) throughput per user is quite low and often can be measured in KB/s rates.

Low latency or fast I/O response time is another key measure of a good user experience. For that reason, it is the key element and the 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 30/70 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 in seconds at the user level and in milliseconds at the storage level. The VDI solution should target latency results of less than 1 to 2 milliseconds for most user operations and less than 1 millisecond for most operations as measured at the storage layer for all-flash arrays.

SE-Sparse and View Storage Accelerator

In early versions of VMware Horizon View and virtual desktops deployed with linked clones, the delta file used was based on 512 byte boundaries when it was carved from the storage system. Because the actual block size of the I/O is very mixed, this limit created a problem for users when they tried to align the storage blocks with block boundaries. It often resulted in additional write and read I/O on all storage platforms.

VMware Horizon View 5.2 introduced a new file system called space-efficient sparse virtual disks, or SE-sparse disks. This file system can change the I/O characteristics of the traditional VMFS-sparse file system that was used in earlier versions of VMware View. SE-sparse helps to eliminate this problem by using 4 KB block boundaries, and it reduces the truncation or misalignment of data blocks.

The SE-sparse disks feature also provides I/O alignment. The space-efficient virtual disk format supports a 4 KB block size as opposed to a redo log format, which is based on 512 byte blocks. This block size ensures that all I/Os from the hypervisor to the storage system (that is, the physical disks) are always aligned on 4 KB block boundaries.

Space reclamation has a multiplier effect, which translates to large capacity savings on centralized storage arrays. When multiple VMs all leverage the same datastore, the space savings on each VM translate to space savings on the entire datastore and on the storage array. This is another benefit of using SE-sparse disks.

In Horizon View 5.2, SE-sparse disks are automatically chosen by default when customers provision pools based on View Composer linked clones. The SE-sparse disk format applies only to the linked-clone portion of the VM, not to the master image or the replica.

Figure 13: Advanced storage options

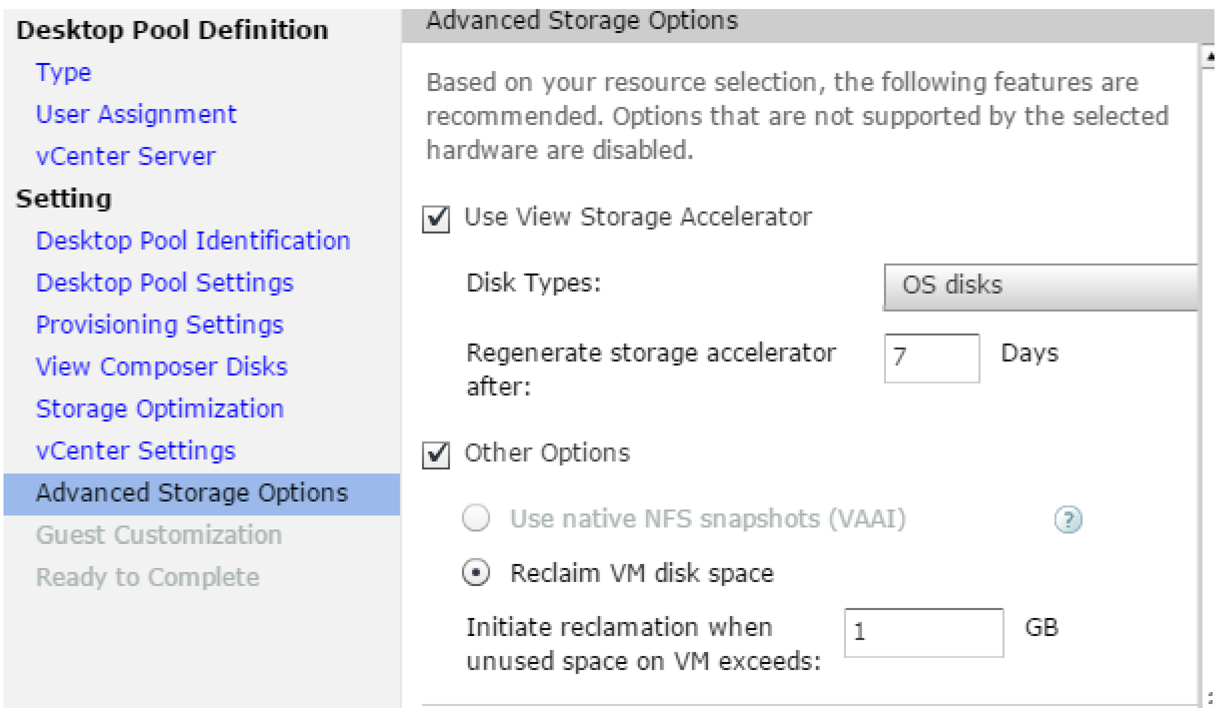

Note The space-reclamation operations in VMs require some CPU cycles. For this reason, HPE recommends scheduling reclamation during periods of inactivity. Space reclamation also generates additional IOPS to the underlying disk storage subsystem. Therefore, it is crucial to select a blackout window during which the space reclamation process can run.

The following graph shows an example of a test run that used linked clones with regular VMFS versus SE-sparse disk on HPE Nimble Storage arrays.The graph of this I/O shows approximately double the amount of misaligned I/O when VMFS-sparse is used as compared to SE-sparse. As the number of desktops running on the HPE Nimble Storage system increases, this reduction in unaligned I/O with SE-sparse has increasingly positive effects on the storage infrastructure.

Figure 14: Unaligned I/O comparison

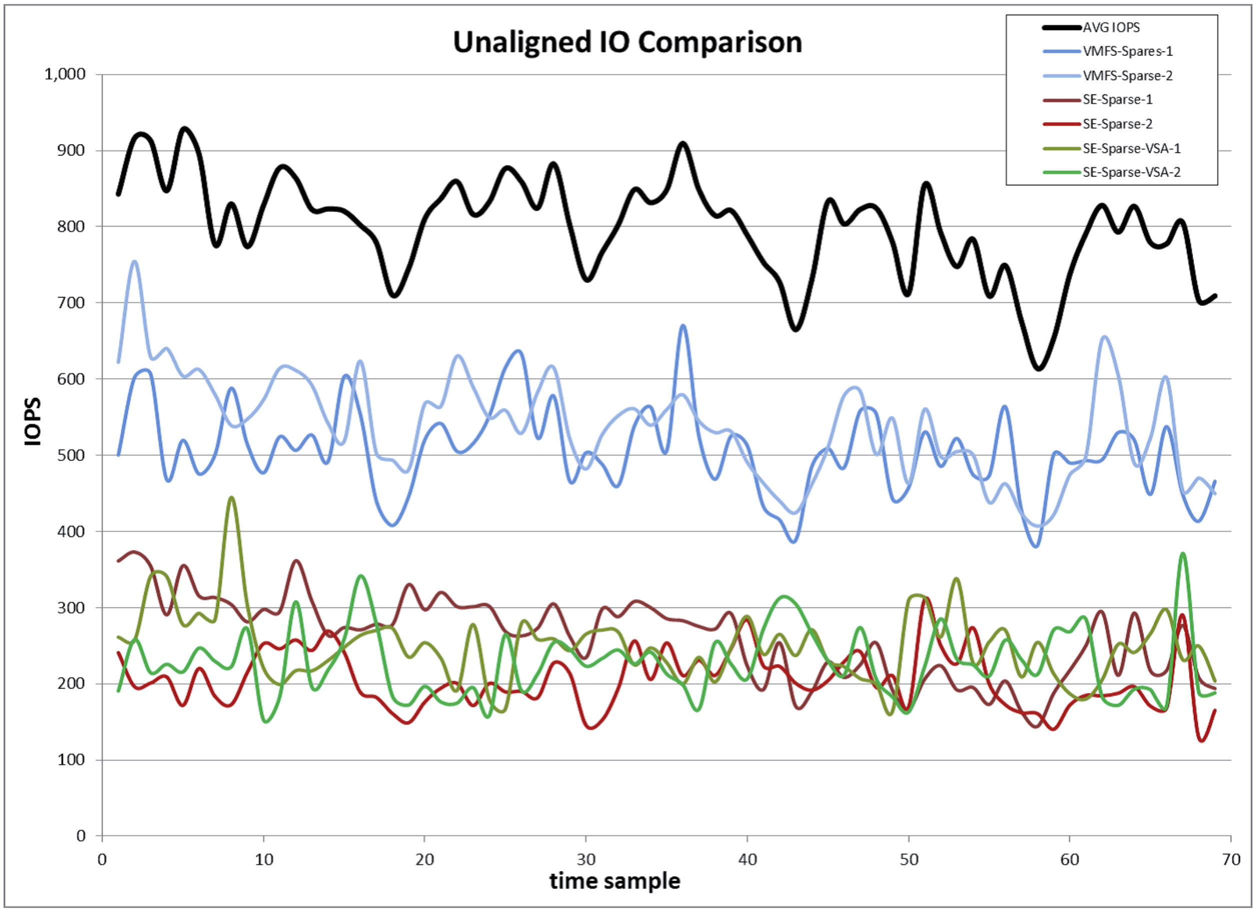

VMware Space Reclamation

The vSphere Storage APIs for Array Integration (VAAI) unmap function was introduced in VMware vSphere 5.0 to enable the ESXi host to inform the storage array that files or VMs had been moved or deleted from a thin-provisioned VMFS datastore. This communication allowed the array to reclaim the freed blocks. VMware software had no way of doing this previously, so many customers ended up with a considerable amount of stranded space on their thin-provisioned VMFS datastores. However, some problems with using this primitive required VMware to disable it temporarily.

The **TRIM** command was introduced with release 5.0 U1 and up to accomplish the same function. The only drawback to using the **TRIM** command is that it must be run manually or be scripted and run recursively. The cleanup might take some time to run, depending on the size of the volume. Also, it is important to run the **UNMAP** command during a blackout period because it might require some CPU and IOPS cycles. For more information, see the following sources:

- *Using the esxcli storage vmfs unmap command to reclaim VMFS deleted blocks on thin-provisioned LUNs (2057513)*
- *VMware Space Reclamation - On All Datastores*

Another way to perform the unmap function is through the guest OS. In this case, the unmap function is supported by using a virtual disk instead of the usual raw device mapping. For more information, see the blog post *vSphere 6.0 Storage Features Part 8: VAAI UNMAP changes*.

As of NimbleOS version 3, the HPE Nimble Storage vCenter plugin offers advice on the best times to run the unmap function.

Figure 15: HPE Nimble Storage vCenter plugin

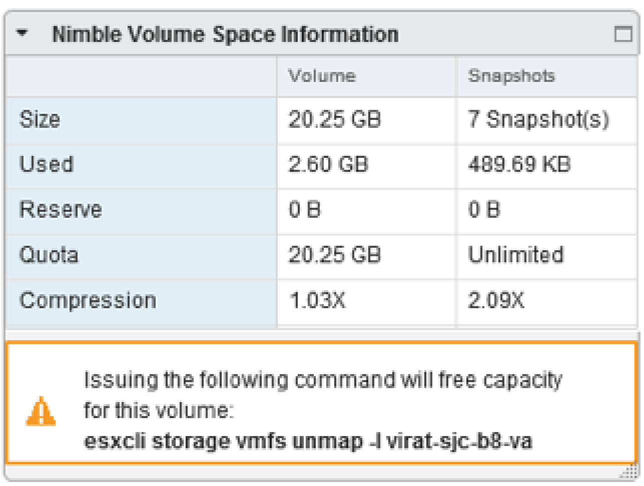

Protection and Recovery

Every production environment needs adequate data protection. Critical infrastructure elements, including desktop templates, golden images, user data, and parts of the VMware platform such as vCenter Server and Horizon View Connection Manager, must be protected with local snapshot protection points and ultimately with remote (replicated) copies. The HPE Nimble Storage VDI solution has highly efficient built-in snapshot mechanisms and the capability to replicate data to other HPE Nimble Storage arrays.

For individual desktops, it is sometimes easier and faster to rebuild the desktop system when problems arise. Horizon View provisioning methods for desktop pools can help you automate some of the rebuilding process for certain types of desktop deployment and assignment models.

HPE Nimble Storage arrays also have strong integration with Veeam Availability Suite. With Veeam, you can accomplish the following protection tasks:

- Perform frequent crash-consistent and application-consistent HPE Nimble Storage snapshots and replicate them to a secondary array for low RPO protection.
- Meet strict RTOs and RPOs by performing granular recoveries from HPE Nimble Storage snapshots on both primary and secondary arrays.
- Create backups from HPE Nimble Storage snapshots on primary or secondary arrays, thus reducing (or completely eliminating) the impact from data protection activities on production workloads.

In addition, Veeam Explorer for Storage Snapshots provides granular-level recovery from an entire VM to the individual email or file-level restore.

Figure 16: Veeam and HPE Nimble Storage integration

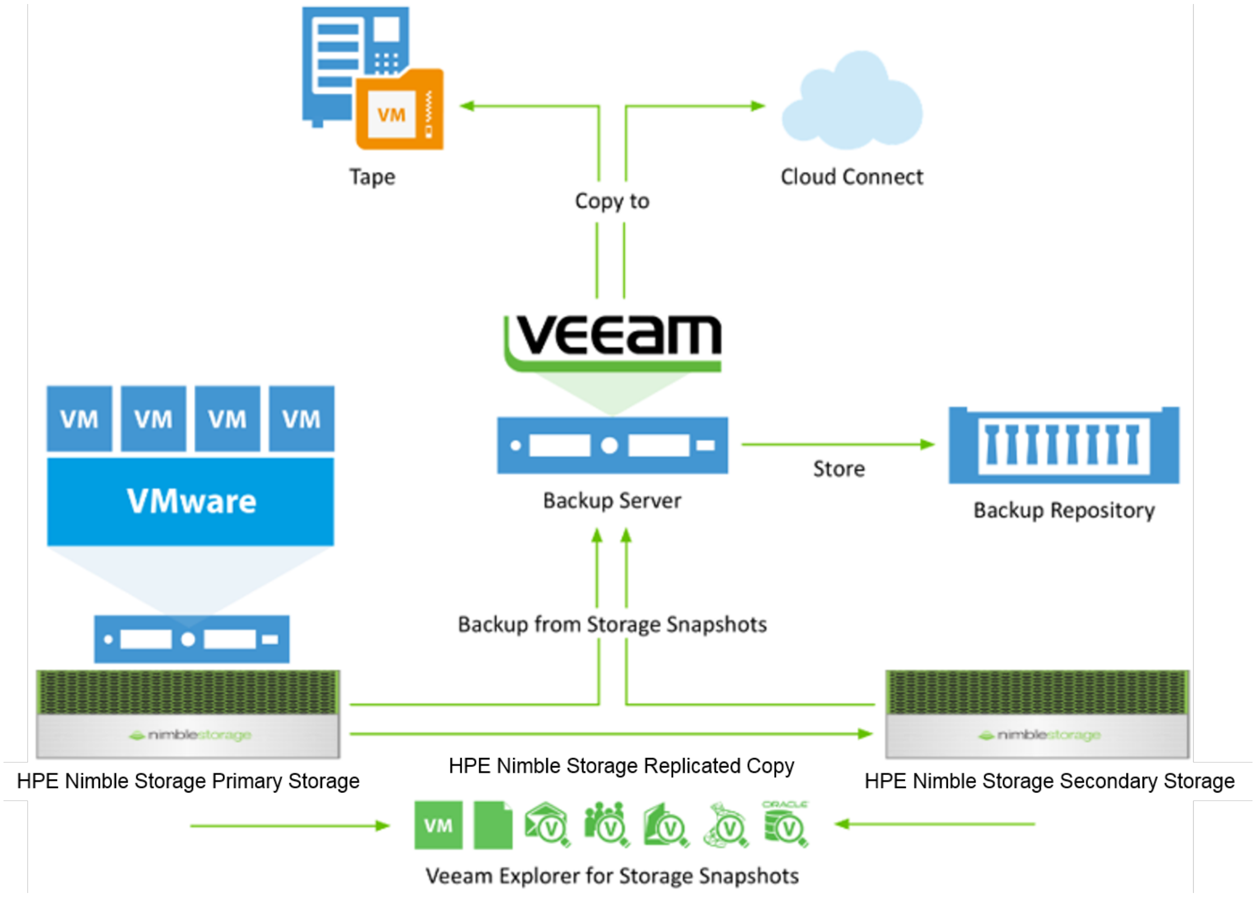

Snapshots

Snapshots offer one way to provide a consistent recovery point within the local storage system. You can store snapshots in the VMware environment, in the storage layer, or in both.

For critical data items, HPE recommends using HPE Nimble Storage snapshots and managing them through volume protection policies. The capabilities of HPE Nimble Storage Protection Manager help you keep the snapshots consistent within the VMware environment.

Because of the large number of desktop VMs that might be associated with a single datastore (the HPE Nimble Storage volume), HPE does not recommend using the VMware vCenter snapshot methods on these volumes. Simple storage snapshots should suffice to create crash-consistent protection for the desktop pools. VMware consistent storage snapshots require that you create a vCenter snapshot of each VM in the targeted volume. If you want such snapshots, be sure to schedule them during periods of little or no desktop activity in those pools.

An additional benefit of leveraging HPE Nimble Storage snapshot functionality is that it helps to reduce the environmental impact caused by increased overhead that is triggered during VMware vSphere snapshots. With the HPE Nimble Storage snapshot functionality, you can create backups and snapshots as often as every 15 minutes with no concerns about creating a performance bottleneck in the VMware vSphere environment.

Replication

The capability of replicating HPE Nimble Storage volumes to a different HPE Nimble Storage array provides another level of infrastructure protection. HPE recommends that you follow VMware replication best practices for your key VDI volumes, just as you would for any other critical virtualization dataset.

If you are replicating data for disaster recovery (DR) purposes, keep in mind that desktop integration into the DR site might need to be managed through a tool such as VMware Site Recovery Manager (SRM). This tool is fully integrated with the HPE Nimble Storage replication through an adapter. Work with your VMware resources to determine which method will provide the greatest flexibility and recoverability of your desktop infrastructure.

Instant Clones

In the Horizon 7 Enterprise Edition, VMware introduced Instant Clone Technology for just-in time-desktops delivery. Instant Clone Technology enables administrators to rapidly clone and deploy VMs.

The following figure shows the difference between the time required to provision View Composer linked clones and the time required for instant clones. Instant Clone Technology (also called VMFork technology) basically cuts the number of steps in half:

After the Master VM is created, a snapshot is taken of it. The snapshot is used to create an internal template VM, which is then used to create a replica VM. Finally, the replica VM is used to create parent VMs for every host in the environment.

Figure 17: Timelines for View Composer clones versus Instant Clone Technology

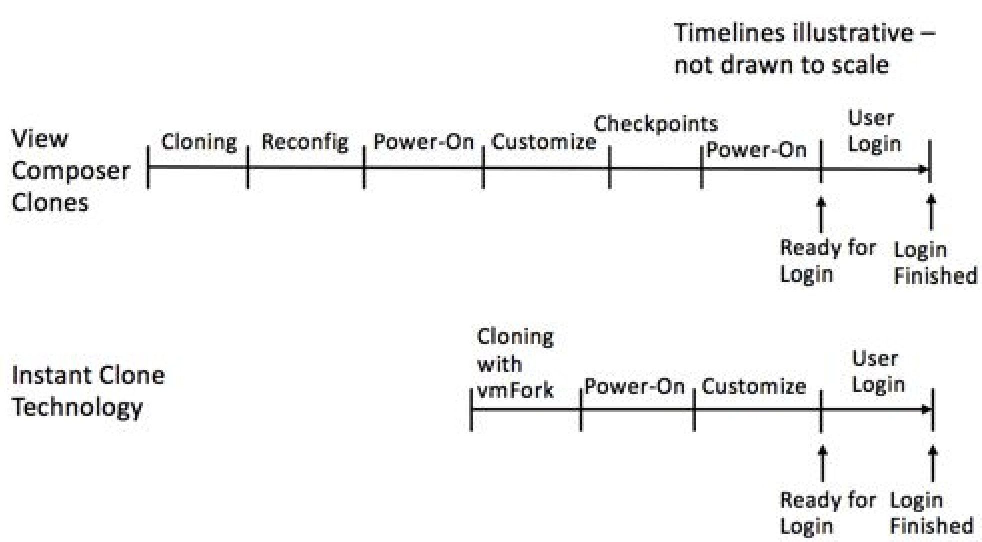

For every host that houses the instant clones, a parent VM is created. This parent VM is used to create and recreate the instant clones. In addition, Instant Clone Technology relies on compute memory. Therefore, when you deploy instant clones, always ensure that enough memory resources are available on the hosts.

Figure 18: Parent VMs

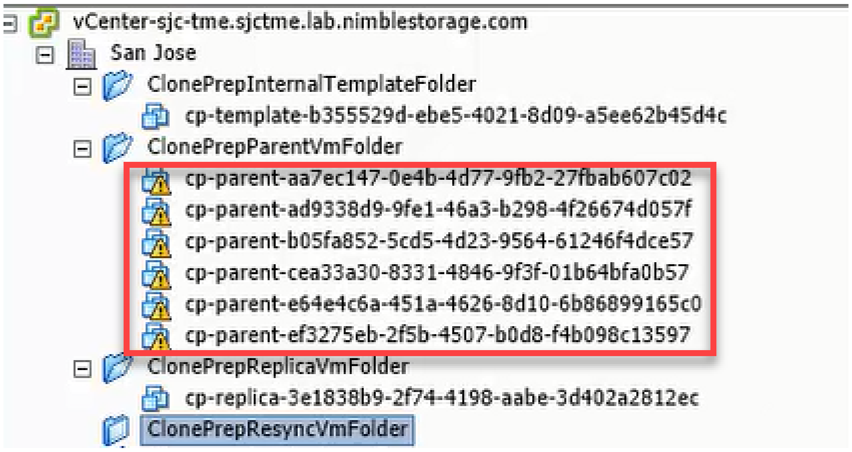

For example, during testing, we deployed 2500 instant clone desktops on an HPE Nimble Storage all-flash array. As the following performance charts show, the average latency was always submillisecond.

Figure 19: Instant clone desktops deployment on HPE Nimble Storage arrays

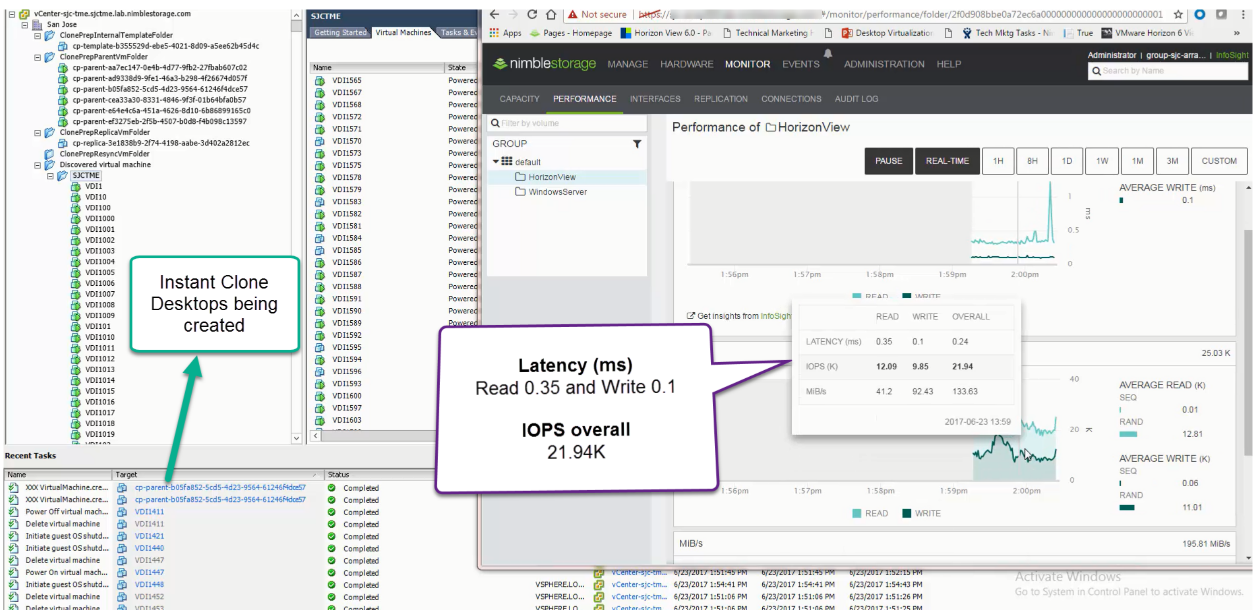

Benefits of Instant Clones over View Composer Linked Clones

At the top of the list of benefits of using Instant Clone Technology is simplified deployment and patching for administrators. Instant clones do not need to be refreshed, recomposed, or rebalanced. When a user logs out of the desktop, that desktop is always deleted and recreated as a fresh image from the latest patch.

Desktop-pool image changes can also be scheduled during the day with no downtime required for the users or for the availability of the desktop pool; the View administrator has full control over when the users receive the latest image. Horizon 7 leverages rolling refresh operations, so users can continue to use their existing logged-in desktop, even during an image-update operation. When a user logs out of a desktop session, a new clone is instantly created, using the latest image, and the desktop is immediately available for login.

However, if you have an urgent patch, you can “force” the user to log out and immediately log in to the latest image.

With Instant Clone Technology, the VMware content-based read cache (CBRC) feature is less useful because the clones are short-lived. They are either deleted or reimaged when the user logs out. However, master VMs and replicas still use CBRC, which is calculated automatically.

SE-sparse wipe-and-shrink operations are not needed. Typically, VMDKs keep growing, and the SE-sparse wipe-and-shrink operation sweeps and frees up unused blocks after they are deleted. With Instant Clone Technology, the operation is not needed because the VMs are short-lived.

Unlike View Composer, Instant Clone Technology does not require a database, which provides multiple advantages:

- It greatly simplifies the Horizon 7 architecture.
- It helps to lower the overall support cost of the solution.
- It reduces the complexity of future environment upgrades.

Summary

When you use HPE Nimble Storage arrays in VMware Horizon View VDI deployments, follow these simple configuration guidelines to get the most from your investment:

- Provision HPE Nimble Storage volumes to VMware datastores by using the default VMware VDI performance policy, which picks the optimal block size. Enable both caching and compression (their default settings).
- Size volumes and VMFS datastores for approximately 100 to 200 full-sized desktops. This range of desktop numbers typically requires approximately 2 TB to 4 TB of data storage, using 20 GB to 40 GB per desktop as a guide.
- Use VMware Horizon View linked clones and instant clones for Horizon View pool configurations to save additional space and speed desktop provisioning.
- Separate the master replica from the linked clones for easier management of datastore and HPE Nimble Storage volume metrics. Separating the master replica and the linked clones does not affect storage performance.
- Use VMware Horizon View desktop pools to partition the desktops by type, location, or manageability. If you want to monitor different desktop types at the storage level, create corresponding HPE Nimble Storage volumes and VMware datastores for the pools.
- Assign multiple datastores to a VMware Horizon View pool to expand capacity for the pool, if needed.
- Leverage SE-sparse disks in the VMware environment, if possible. Also, use the **UNMAP** command to recover space and delete unwanted blocks to increase data savings.
- Define protection templates for critical volumes that cannot be easily redeployed with VMware Horizon View mechanisms. Such volumes include VM templates, golden images, user data, the desktops of critical users, and infrastructure components.

Additional Resources

For additional information, consult the following resources:

- *SEsparse in VMware vSphere 5.5 Performance Study*
- *VMware Instant Clone Technology for Just-In-Time Desktop Delivery in Horizon 7 Enterprise Edition*
- *View Storage Accelerator – In Practice*
- *View Storage Accelerator in VMware View 5.1*
- *Space Efficient Virtual Disks – In Practice*

About the Author

Bharath Ram

Senior Technical Marketing Engineer

Hewlett Packard Enterprise

Bharath Ram specializes in virtualization technologies for servers and desktops and in storage solutions for data centers. He has extensive experience with the VMware and Citrix product portfolios, and he has worked on numerous projects to integrate these products with HPE Nimble Storage offerings. Before joining HPE, Bharath was a solution architect for other storage companies where he implemented virtual desktop and application solutions for the healthcare and insurance domains.

Version History

Version	Release Date	Description
1.0	August 2017	Initial release