

VMware Horizon® 7 on VMware vSAN™ Best Practices

TECHNICAL WHITE PAPER

Table of Contents

Introduction	3
Purpose	3
Audience.....	3
Technology Overview and Best Practices	3
Overview.....	3
VMware vSAN	4
Introduction	4
All-Flash vs. Hybrid Architecture	4
Storage Hardware	6
Deduplication and Compression	7
Storage Policies	9
Swap Thin Provisioning.....	11
Native Encryption	11
vSAN Encryption vs. VM-level Encryption	12
VMware Horizon 7	13
Introduction	13
Cloning Technology.....	13
Full Clones	13
Linked Clones	14
Instant Clones	15
VMware View Storage Accelerator and vSAN Client Cache	18
References	20
White Papers.....	20
Product Documentation.....	20
About the Authors.....	21

Introduction

Purpose

As more virtual desktop infrastructure customers are embracing hyper-converged infrastructure (HCI) technology to provide cost-effective, highly scalable, and easy-to-manage solution, they are looking for more information and recommendations for how these products work in conjunction.

This white paper provides best practice recommendations when running VMware Horizon® 7 on VMware vSAN™ for a virtual desktop infrastructure (VDI) environment. This document is not meant to be a complete best practice guide on Horizon 7 or on vSAN. Excellent solution architectures are already available (links provided in the Reference section). This document focuses on the specific intersection points between the VDI platform and the storage platform and covers areas such as cloning, deduplication, storage consumption, etc.

Note that Horizon 7 is the full name of the VMware desktop and application management platform and does not denote any specific product versions.

Audience

This reference architecture is intended for customers—IT architects, consultants, and administrators—involved in the early phases of planning, design, and deployment of VDI solutions using VMware Horizon 7 running on vSAN. It is assumed that the reader is familiar with the concepts and operations of VMware vSphere, vSAN and Horizon 7 technologies.

Technology Overview and Best Practices

Overview

This section provides an overview of the technologies that are used in this solution as well as best practices when using these technologies:

- VMware vSAN™
 - All-Flash and Hybrid Architecture
 - Deduplication and Compression
 - Storage Policies
 - Native Encryption
- VMware Horizon® 7
 - Full Clone Technology
 - Linked Clone Technology
 - Instant Clone Technology



VMware vSAN

Introduction

VMware vSAN™ is a hyper-converged infrastructure platform that is fully integrated with VMware vSphere. vSAN aggregates locally attached disks of hosts that are members of a vSphere cluster to create a distributed shared storage solution. Seamless integration with vSphere and the VMware ecosystem makes it the ideal storage platform for Horizon 7 VDI. vSAN provides scale-out storage within a Horizon 7 environment, enabling a grow-as-you-go model, with scaling up by adding disk drives in each host, or with scaling out by adding hosts to the cluster.

All-flash vSAN configurations provide the highest levels of performance with very low latencies for the most demanding virtual desktop workloads. Space efficiency features such as deduplication, compression, and RAID-5/6 erasure coding minimize capacity consumption, which reduces the cost per gigabyte of usable capacity.

Hybrid configurations use both flash and magnetic disks to provide a cost-effective platform for enterprise-class performance and resiliency.

Per-virtual machine (VM) storage policy-based management lowers operational expenditures by enabling administrators to manage performance, availability, and capacity consumption with ease and precision. Native data-at-rest encryption, with FIPS 140-2 validation, can be enabled without the need for specialized hardware, which provides regulatory compliance without the typical costs associated with procuring and maintaining self-encrypting drives.

Many deployment options are available for vSAN. These options range from 2-node clusters for small implementations to multiple clusters each with as many as 64 nodes--all centrally managed by VMware vCenter Server. vSAN stretched clusters can easily be configured to enable cross-site protection with no downtime for disaster avoidance and rapid, automated recovery from entire site failure.

All-Flash vs. Hybrid Architecture

vSAN provides two different configuration options:

- An all-flash configuration
- A hybrid configuration that uses both flash-based devices and magnetic disks

The all-flash configuration uses flash for both the caching layer and capacity layer. All-flash vSAN is an optimized platform for high performance and delivers greater and more consistent overall performance vs. hybrid configurations.



All-flash vSAN aims at delivering extremely high IOPS with predictable low latencies. In all-flash architecture, two different grades of flash devices are commonly used in the storage hardware configuration:

- Lower capacity and higher endurance devices for the cache layer
- More cost-effective, higher capacity, and lower endurance devices for the capacity layer

The hybrid configuration uses:

- Server-based flash devices to provide a cache layer for optimal performance
- Magnetic spinning disks to provide capacity and persistent data storage

Hybrid vSAN configurations delivers both enterprise-ready levels of performance and a resilient storage platform.

All incoming writes are performed at the cache layer and then de-staged to the capacity layer. All data in the cache layer must be eventually de-staged, which happens asynchronously to achieve maximum efficiency. This helps extend the usable life of lower endurance flash devices in the capacity layer and lower the overall cost of the solution. All-flash configurations are required for storage efficiency capabilities such as deduplication, compression and RAID-5/6 erasure coding, all of which minimize raw capacity consumption.

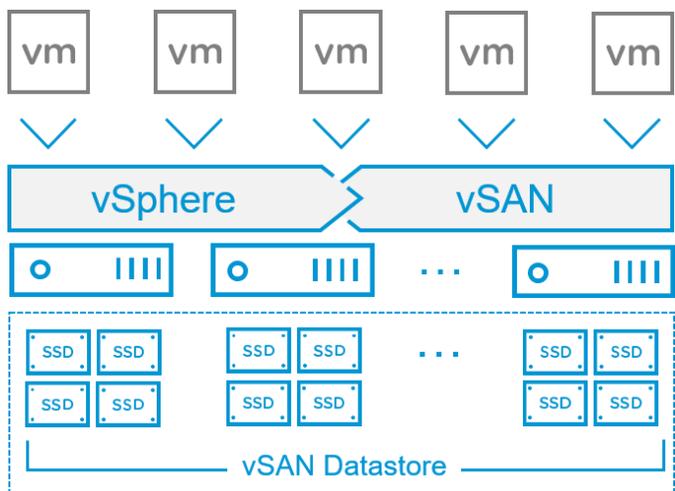


Figure 1. vSAN All-Flash Datastore

Hybrid vSAN configurations use both flash and magnetic disks to provide a cost-effective platform for enterprise-class performance and resiliency.



Hybrid configurations offer the lowest TCO due to the inherent lower cost of magnetic disks when compared to flash disks for the capacity layer.

However, it is important to know that properly designing and sizing a vSAN hybrid configuration is extremely important to deliver predictable performance. Correct sizing of the cache device is the chief consideration, with sizing of the magnetic disk subsystem behind the cache being the secondary consideration. Hybrid configurations do not support storage efficiency capabilities such as deduplication, compression or RAID-5/6 erasure coding.

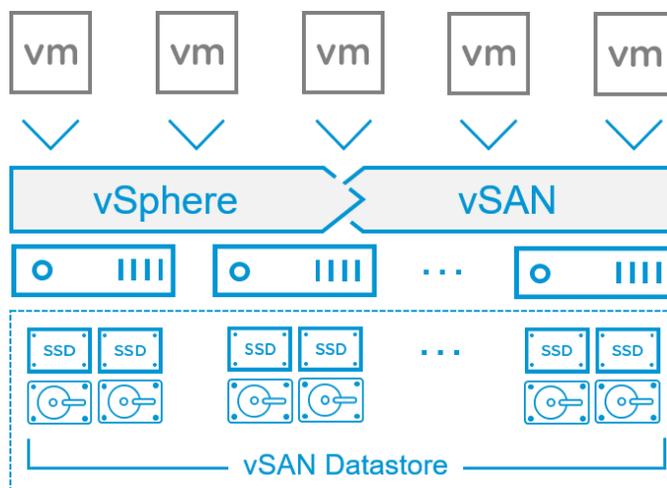


Figure 2. vSAN Hybrid Datastore

Storage Hardware

vSAN hosts that contribute storage can be configured with between one and five disk groups for the storage of vSAN objects. Disk groups require at least a single flash disk drive used for the cache tier, and between one and seven disk drives for the capacity tier. In all disk group configurations, a flash device is used for cache. In hybrid configurations, the capacity devices are comprised of SAS or NL-SAS magnetic disks. In all-flash configurations, the capacity devices may be flash SATA, SAS, PCIe, or NVMe.

Devices such as SAS, NL-SAS, or SATA are attached to a Host Bus Adapter (HBA) or RAID controller for consumption of vSAN. These devices should be connected in pass-through mode and not RAID0 mode, depending on the HBA/RAID controller. For controllers that do not support pass-through mode, each device must be presented as an individual RAID0 device. While RAID controllers may support drive mirroring, striping, or erasure coding, these are not supported, nor required by vSAN. vSAN is an object-based storage



system and distributes data across hosts in the cluster, which removes the need for these hardware-level mirroring, striping, or erasure coding. Instead, data protection and performance properties are defined logically using the Storage Policy Based Management (SPBM) framework instead.

Just as compute and networking must be on the [VMware Compatibility Guide](#), vSAN storage devices, such as Host Bus Adapters (HBA), RAID controllers, and storage devices must be on the [VMware Compatibility Guide for vSAN](#) to be supported. It is also important that these devices are running a supported firmware version as detailed in the HCL.

With regards to availability, consider choosing hosts that have sufficient disk drive slots to accommodate more than one disk group, for both hybrid and all-flash configurations. Having multiple groups will increase availability by reducing the storage failure domain per host. In other words, for hosts with a single disk group and a single cache device, a cache device failure will result in failure of the entire host. However, for hosts with two disk groups with one cache device each, a single cache device failure in one disk group will not impact data being served from the remaining disk group. In addition, when deduplication and compression is enabled, the loss of a single capacity disk, in any disk group, will also result in failure of that entire disk group.

With regards to performance, choosing hosts with multiple disk groups will improve overall performance for both front-end VM traffic and back-end vSAN traffic. Back-end vSAN traffic occurs after a disk device or host goes offline, fails, or when the capacity utilization of any disk exceeds 80%. Having multiple disk groups per host enables greater parallelism in these operations.

Recommendation: *Configure hosts with more than one disk group to achieve the highest levels of vSAN availability and performance. In addition, for hosts that are configured with many disk drives and multiple disk groups, distribute the storage I/O path across more than one HBA controller.*

Deduplication and Compression

vSAN deduplication and compression provides enterprise-class storage efficiency by minimizing the space required to make data persistent in the capacity layer. Deduplication and compression are always enabled or disabled together at the cluster level using a simple drop-down menu. It is not possible to enable vSAN deduplication or compression individually or for individual VMs. All-flash vSAN is required to use deduplication and compression. Note that a rolling reformat of all disks in the vSAN cluster is required, which can take a considerable amount of time depending on the amount of data. However, this process does not incur VM downtime and can be done online, usually during an upgrade.

Recommendation: *If vSAN deduplication and compression is part of the design decision, enable the service before any virtual desktops are deployed to the vSAN datastore. This will expedite the time required to enable the service.*



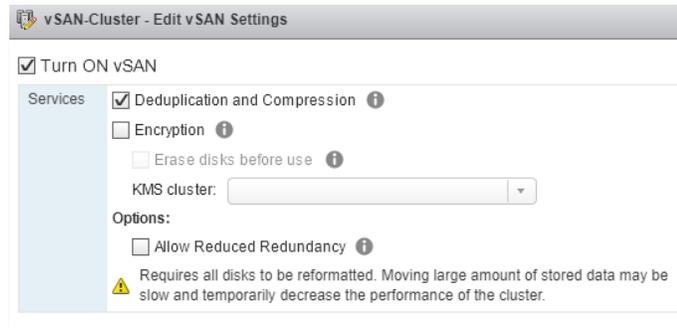


Figure 3. Enabling vSAN Deduplication and Compression

Deduplication occurs when the data is de-staged from the cache tier to the capacity tier. The deduplication algorithm utilizes a 4K-fixed block size and is performed within each disk group. In other words, redundant copies of a block within the same disk group are reduced to one copy, but redundant blocks across multiple disk groups are not deduplicated. Upon writing a 4K block, it is hashed to find whether an identical block already exists in the capacity tier of the disk group. If there is one, only a small metadata is updated. If no such identical block is available, compression is then applied to the 4K block. If the 4K block can be compressed to 2K or less, vSAN persists the compressed data to the capacity tier. Otherwise, the 4K block is persisted to the capacity tier uncompressed.

Deduplication and compression are applied to data in the capacity tier, commonly accounting for approximately 90% of all data on a vSAN datastore. Storing this data in 4K blocks enables effective deduplication and compression with minimal resource overhead for these operations. Deduplication and compression are not applied to data in the cache tier, which serves as a write buffer in an all-flash vSAN configuration. Naturally, the cache tier is being written to much more frequently than the capacity tier.



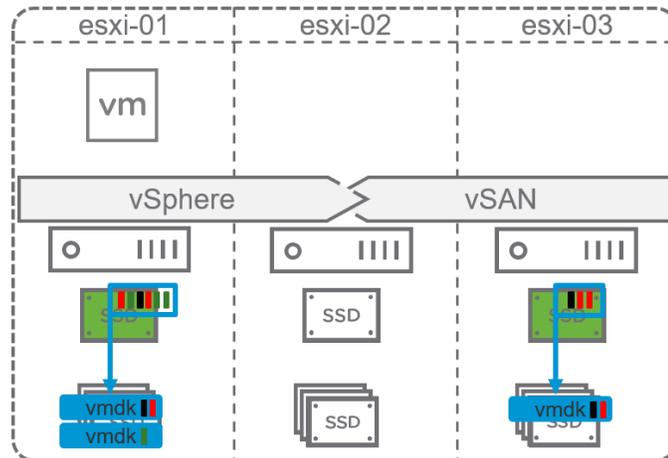


Figure 4. Deduplication and Compression Space Efficiency

The processes of deduplication and compression on any storage platform incur overhead and potentially impact performance in terms of latency and maximum IOPS. vSAN is no exception. However, considering deduplication and compression are only supported in all-flash vSAN configurations, these effects are predictable in the majority of use cases. The extreme performance and low latency of flash devices easily outweigh the additional resource requirements of deduplication and compression. Enabling deduplication and compression consumes a small amount of capacity for metadata, such as hash, translation, and allocation maps. The space consumed by this metadata is relative to the size of the vSAN datastore and is typically around 5% of the total capacity. Note that the user interface displays the percentage of used capacity, not total capacity (used and free space). In addition, enabling deduplication and compression consumes minimal CPU overhead – typically around 5% of the total cluster processing capacity.

Recommendation: *If using an all-flash vSAN configuration, enable deduplication and compression for Horizon 7 linked clone environments for both storage efficiency and accurate reporting of storage utilization. For instant clones, only enable deduplication and compression for improved reporting of storage utilization.*

Storage Policies

Per-VM storage policy-based management is a foundational benefit of vSAN hyper-converged infrastructure. Unlike traditional storage solutions which must apply storage policies on a LUN or volume which may contain several VMs, vSAN enables precise control on a per-VM level. Administrators can manage performance, availability and capacity consumption with ease and precision for each VM in the environment.

Typically, vSAN storage policies are created and managed using the vSphere Client. Storage policies can be assigned to entire VMs or individual VMDKs



within those VMs. Storage policies are either applied to VMs at the time of deployment or reassigned if the application requirements have changed. These modifications are performed with no downtime and without the need to migrate VMs from one datastore to another. It is important to note that changing the vSAN default storage policy or a global policy that applies to many VMs will require temporary storage overhead and may take a long time to complete depending on the scope of changes.

Recommendation: Only apply storage policy changes to small groups of VMs at any one time to minimize temporary storage overhead and overall resynchronization activity.

For Horizon 7 virtual desktop infrastructure, default storage policies are automatically created during desktop pool creation, depending on the type of pool you create. Horizon 7 creates vSAN storage policies for linked clone desktop pools, instant clone desktop pools, full clone desktop pools, or an automated farm per Horizon 7 cluster. Once these storage policies are created for the desktop pool; they will never be changed by Horizon 7. An administrator can edit these storage policies in vCenter, similar to a regular vSAN policy if Horizon 7 was not in use. Any new default storage policies enacted by Horizon 7 will not impact existing desktop pools. Each VM maintains its storage policy regardless of its physical location in the cluster. If the storage policy becomes non-compliant because of a host, disk, network failure or workload changes, vSAN reconfigures the data of the affected VMs and load balances to meet the compliance of the storage policy.

The screenshot shows the vSAN storage policy configuration interface. The 'Placement' section is expanded, revealing the following settings:

- Storage Type: VSAN
- Primary level of failures to tolerate: 1
- Failure tolerance method: RAID-1 (Mirroring) - Performance
- Number of disk stripes per object: 1
- Flash read cache reservation (%): 0
- Object space reservation (%): 0

At the bottom of the configuration area, there is a button labeled '<Add rule>'.

Figure 5. Default vSAN storage policies configured by Horizon 7

The default policy settings that Horizon 7 automatically configures are similar to the default vSAN storage policy settings that are configured for all vSAN deployments. These settings provide the baseline vSAN capabilities and are appropriate for many use cases unless the environment requires higher levels of availability, performance or storage efficiency.



Recommendation: *If storage efficiency is part of the design decision, consider using RAID-5/6 erasure coding instead of the default RAID-1 mirroring. If virtual desktops have already been deployed using the default policy settings, make a clone of the existing policy and then change the failure tolerance method to RAID-5/6 of the cloned policy. Then, apply this new storage policy to small groups of desktops at one time to minimize the impact of vSAN policy reconfiguration. In addition, consider using FTT=2 for the replica VM to increase availability.*

Swap Thin Provisioning

vSAN storage policies allow configuration of the VM or VMDK object space reservation, which is synonymous with enabling thick-provisioning on vSAN. When the administrator (or Horizon 7) configures an object space reservation of 0%, the VM or VMDK is thin-provisioned. However, this is not applied to the VM swap file (.vswp) in versions prior to vSAN 6.7. In these earlier versions, the .vswp file always has an object space reservation of 100%, even if the storage policy specifies 0%. This behavior can be disabled by configuring the advanced host setting “SwapThickProvisionedDisabled”, so that the .vswp file is thin provisioned for these versions of vSAN.

Recommendation: *Since swap files are thin provisioned in vSAN 6.7 by default, manually enable swap file thin provisioning in vSAN versions prior to 6.7 using the above advanced setting. It is important to only use swap file thin provisioning in environments where physical memory is not overcommitted, or where storage efficiency is part of the design decision.*

Native Encryption

vSAN native encryption for data-at-rest further improves security and provides compliance with increasingly stringent regulatory requirements. vSAN encryption uses an AES 256 cipher and is FIPS 140-2 validated. vSAN encryption is hardware-agnostic, meaning it can be deployed on any supported hardware in all-flash or hybrid configurations. Self-encrypting drives (SEDs) are not required. vSAN encryption is enabled and configured at the datastore level. In other words, every object on the vSAN datastore is encrypted when this feature is enabled. Note that a rolling reformat of all disks in the vSAN cluster is required, which can take a considerable amount of time depending on the amount of data. However, this process does not incur VM downtime and can be done online, usually during an upgrade.

Recommendation: *If vSAN encryption is part of the design decision, enable the service before any virtual desktops are deployed to the vSAN datastore. This will expedite the time required to enable the service.*



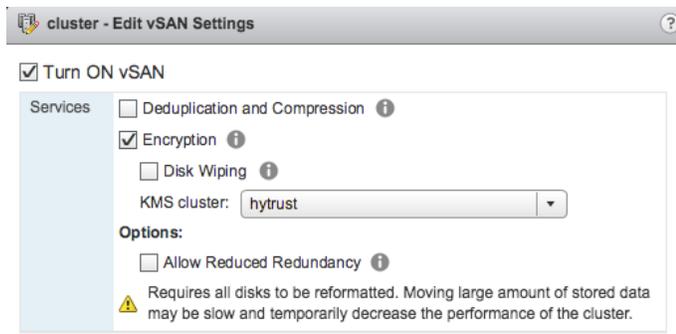


Figure 6. Enabling vSAN encryption

Data is encrypted when it is written to persistent media in both the cache and capacity tiers of a vSAN datastore. Encryption occurs just above the device driver layer of the storage stack, which means it is compatible with all vSAN features such as deduplication and compression, RAID-5/6 erasure coding, stretched cluster configurations. All vSphere features including VMware vSphere vMotion, VMware vSphere Distributed Resource Scheduler (DRS), VMware vSphere High Availability (HA), and VMware vSphere Replication are supported.

A Key Management Server (KMS) is required to enable and use vSAN encryption. Nearly all KMIP-compliant KMS vendors are compatible, with specific testing completed for vendors such as HyTrust®, Gemalto®, Thales e-Security®, CloudLink®, and Vormetric®. These solutions are commonly deployed in clusters of hardware appliances or virtual appliances for redundancy and high availability. Encryption keys are transferred to vSAN hosts using the Key Management Interoperability Protocol (KMIP). Industry standards and compliance with regulations often require the generation of new keys on a regular basis. This reduces the risk of a key being exposed or compromised by brute force. Generating new keys is performed in the vSAN UI with just a few clicks.



Figure 7. KMS configured for use with vCenter Server

vSAN Encryption vs. VM-level Encryption

VMware vSphere and vSAN provide two different methods of encrypting data, and it is important to understand the differences between the two solutions.



- vSAN provides native data-at-rest encryption for the entire datastore, as covered previously in this section.
- VMware vSphere provides VM-level encryption, which is not associated or related to the vSAN encryption capabilities.

VM-level encryption can be used by non-vSAN users. vSAN encryption is enabled one time for the entire datastore, whereas VM-level encryption is enabled through policy-based management on a per-VM basis.

Other than the granularity of encryption, the primary differences are when the data is encrypted and if storage efficiency capabilities are supported. With vSAN, data is transmitted unencrypted until it reaches the datastore, where it is then encrypted. vSAN encryption can co-exist and benefit from deduplication and compression capabilities. With VM-level encryption, data is encrypted in upper layers before it is transmitted to the underlying datastore, however this feature cannot take advantage of vSAN deduplication and compression.

Recommendation: *If storage efficiency or cluster-wide encryption is part of the design decision, only enable vSAN data-at-rest encryption. If data-in-flight encryption or per-VM encryption granularity is more important, use VM-level encryption instead.*

VMware Horizon 7

Introduction

VMware Horizon® 7 delivers virtualized or hosted desktops and applications through a single platform to end users. These desktop and application services—including Remote Desktop Services (RDS) hosted apps, packaged apps with VMware ThinApp®, software-as-a-service (SaaS) apps, and even virtualized apps from Citrix—can all be accessed from one digital workspace across devices, locations, media, and connections without compromising quality and user experience. Leveraging complete workspace environment management and optimized for the software defined data center, Horizon 7 helps IT control, manage, and protect all of the Windows resources end users want, at the speed they expect, with the efficiency that business demands.

Cloning Technology

A clone is a copy of a master VM or golden image with a unique identity of its own, including a MAC address, UUID, and other system information. VMware Horizon 7 provides three types of cloning technologies to provide customers choice and flexibility. Persistent virtual desktops can be deployed using full clones. Non-persistent virtual desktops can be deployed using linked clones or the newest cloning technology, instant clones.

Full Clones

A full clone is an independent copy of a VM. It shares nothing with its master VM or golden image, and it operates entirely separately from the golden image used to create it. Since each full clone VM is almost identical to the



golden image, this means that there is high degree of duplication across a pool of full clone VMs.

Recommendation: *If using an all-flash vSAN configuration, always enable vSAN deduplication and compression to reduce the duplication across multiple full clone VMs*

Linked Clones

A View Composer linked clone uses significantly less storage space than a full clone because it accesses software on shared virtual disks. Because of this sharing mechanism, a linked clone must always have access to the disk used for cloning.

To make a linked clone, you take a snapshot of the golden image and then the Horizon 7 cloning process creates a replica VM to use for cloning. The linked clone shares virtual disks with the replica VM. The differential—the bits of software that are unique to the linked clone—is stored in a diff disk or redo disk. The differential is called delta disks. This arrangement allows the linked clone to occupy a smaller amount of physical disk space than the golden image, but still access the software installed on the shared virtual disks. You can create hundreds of linked diff disks from one replica, reducing the total storage space required.

Linked clones are generated on Horizon 7 by the View Composer server. In the process of creating delta disks, two vmdks are created for each linked clone: .vmdk and checkpoint.vmdk.

- The .vmdk disk is a snapshot of the state of the delta disks at the time of creation. It cannot be modified in any way by the user or by the system. View Composer creates this snapshot and persists it so that you can rapidly revert to a pristine copy of the delta disks during a refresh or recompose operation.
- The checkpoint.vmdk disk is where all the system and user changes are written. As such, it will grow as the virtual desktop is used. When a linked clone is refreshed or recomposed, the checkpoint.vmdk disk is deleted and recreated, but the .vmdk disk remains.

vSAN logical units of space are allocated in 4MB blocks, whereas VDI data is often written to the filesystem in smaller blocks (e.g. 512KB). This can result in free space allocated in each 4MB block. This behavior explains why at the time of initial linked clone creation, the .vmdk and checkpoint.vmdk disks appear larger than they actually are in the physical layer. This also explains why these disks will appear much larger than linked clones deployed on traditional VMFS storage.

However, the empty spaces will be utilized for future writes as users begin to use their virtual desktops. As additional data is written these empty spaces are consumed, the eventual storage utilization become comparable to VMFS. By design, enabling vSAN deduplication and compression will provide the



best storage efficiency at the logical and physical layers. This is due to removal of redundant copies of data (including empty spaces) and compression of data after it has been deduplicated. Even if there are no deduplication savings (i.e. data is completely unique), enabling this mode will report the actual physical storage consumption to the logical layer, post-compression.

Recommendation: *If using an all-flash vSAN configuration, enable deduplication and compression for Horizon 7 linked clone environments for both storage efficiency and accurate reporting of storage utilization.*

As the end-user creates and deletes content on their desktops, Windows automatically creates and deletes system files. When the end-user and OS delete data, the corresponding data are marked for deletion but are not immediately deleted on the physical storage hardware. This behavior occurs for any storage system and may cause storage bloat.

Recommendation: *To avoid consuming unnecessary storage when using linked clones:*

- *Refresh or recompose the pool on a frequent basis*
- *Set logoff policy of the pool to “refresh on logoff”*
- *Use an SDD (System Disposal Disk), on which you redirect the temporary writes. This disk is deleted and recreated on every logoff.*

Instant Clones

Like a linked clone, an instant clone shares virtual disks with the replica VM after the linked clone is created. The process of creating instant clones differs from that used for linked clones in the following way: The cloning process creates a running parent VM from the replica VM. At creation time, the instant clone shares the memory pages of the running parent VM from which it is created.

Instant clones use copy-on-write for memory and disk management. Instant clones are based on a running parent VM, derived from a master VM. At the instant when an instant clone is created from a running parent VM, any reads of unchanged information come from the already existing running parent VM. However, any changes made to the instant clone are written to a delta disk, not to the running parent VM. This strategy preserves security and isolation between the instant clones by ensuring the following:

- Each instant clone is immediately accessible.
- Changes do not affect the shared data and memory of the running parent VM on which all other instant clones are based. Sharing the memory page of a running parent VM at creation time enables instant clones to be created within a few seconds and instantly powered on. With a few exceptions such as vGPU enabled desktop and Linux desktop, an instant clone requires no extra boot when the cloning process is finished.



- After creation, the clone is linked to the replica VM and not to the running parent VM. You can delete the running parent VM without affecting the instant clone.

One of the additions that Horizon 7 makes is to maintain a set of internal VMs for instant clone provisioning. These internal VMs are created and managed by Horizon 7 automatically. There are 3 types of internal VMs (illustrated in diagram below):

- When you create an instant clone pool, Horizon 7 creates one internal template VM and it is used for AD domain join. There is exactly one internal template VM per golden image snapshot. The internal template VM is a thin clone and therefore does not take up much space
- Next, Horizon 7 creates a replica VM for each golden image snapshot per datastore. It is used as a base disk for instant cloning. Since vSAN only has one datastore, Horizon 7 only creates one replica VM per golden image snapshot.
- Horizon 7 then creates a parent VM per host. The parent VM is partially booted and “stunned” and captures the runtime state of the memory. It is used to “fork” instant clones. On each host, Horizon 7 creates a parent VM for each golden image snapshot.

Internal template, replica VM and parent VM are automatically managed by Horizon 7, although they show up in vCenter (with prefix -cp) under protected status. Once a pool of instant clones has been created, they are no longer dependent on the internal VMs. Horizon 7 perseveres internal VMs so that provisioning of additional instant clones from the same golden image snapshot is a rapid procedure. Normally you would not need to delete the internal VMs; Horizon 7 automatically deletes them when there are no pools using those internal VMs. On some occasions when the automatic deletion encounters an error, you can unprotect the internal VM folders in vCenter, and then delete.



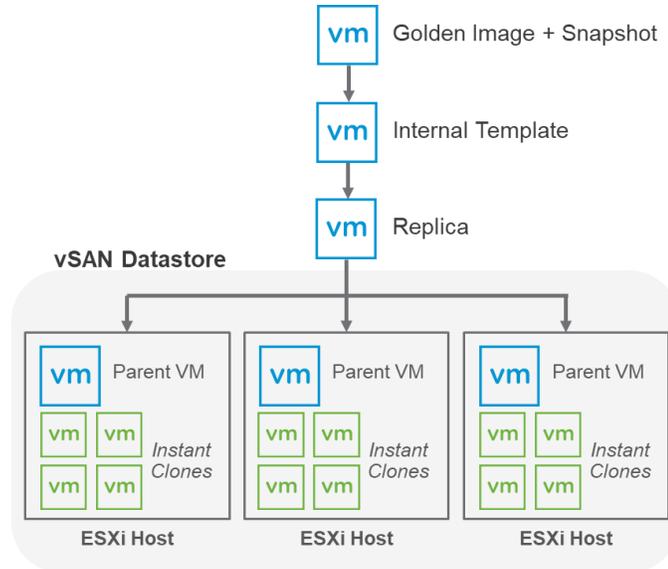


Figure 8. Horizon 7 Instant Clone Architecture on vSAN

The parent VM is pinned to each host. When a host is put into maintenance mode, Horizon 7 automatically deletes the parent VM to enable the host to enter maintenance mode.

Because an instant clone can be created so quickly, an instant-clone desktop does not need to persist after a user logs out. Instead, the instant clone is deleted when the user logs out. Depending on the number of spare VMs configured for the desktop pool, a new instant clone might be created immediately after a used instant clone is deleted. In this manner, users get a newly created desktop whenever they log in. If the master image—the master VM snapshot used to create the pool—has been updated since the last login, the user gets the new image.

Note: The instant clone is deleted when the user logs out, not necessarily when the user disconnects. If the user disconnects the session, the virtual desktop remains, unless the administrator has configured the user to be automatically logged out after disconnecting.

This quick delete and re-provision enables a pool of instant clones to be patched with minimum downtime. While a pool of instant clones is up and running, prepare a new golden image and take a snapshot. When you are ready, schedule a Push Image, Horizon 7 will start creating the internal VMs. At the scheduled time, Horizon 7 will delete all of the unused instant clones and replace them with newly created instant clones from the patched golden image. Then as users log out of the in-use instant clones, Horizon 7 deletes them and replace them with new created instant clones from the patched golden image. Once all the instant clones in the pool have been logged out



once, the pool will have been patched completely. This rolling patching process ensures that your users will always have access to a desktop during patching. For emergency patches, you do have the option to force all users to log off and patch.

Due to naming, there is some confusion as to the difference between vSphere instant clone and Horizon 7 instant clone. They are very different, although one depends on the other. vSphere instant clone is an API that is responsible for the underlying “forking” from a parent VM to children clones. Horizon 7 instant clone calls the vSphere instant clone API and then customizes it for VDI deployment. In the process, Horizon 7 instant clone does three things:

- Adapt vSphere instant clone API for VDI lifecycle management, including initial VM creation as well as patch operations
- Provides Windows customization necessary for Windows Client OS
- Enhancements to scale up to 2,000 instant clones per pool and 10,000 instant clones per Horizon 7 Pod

Similar to linked clones, each instant clone also has a checkpoint.vmdk disk, but it does not have a .vmdk snapshot disk. Since instant clones are deleted and recreated when the user logs out, there are no concerns with storage bloat because of end-user and OS data deletion.

Recommendation: *If using an all-flash vSAN configuration, only enable deduplication and compression for Horizon 7 instant clones for improved reporting of storage utilization.*

VMware View Storage Accelerator and vSAN Client Cache

View Storage Accelerator is an in-memory host caching capability that uses the Content-Based Read Cache (CBRC) feature in ESXi hosts. CBRC provides a per-host RAM-based solution for View desktops, considerably reducing the read I/O requests that are issued to the storage layer. It also improves performance during boot storms when multiple virtual desktops are booted at once, which causes a large number of reads. CBRC is beneficial when administrators or users load applications or data frequently.

vSAN Client Cache is a mechanism that allocates 0.4% of host memory, up to 1GB, as an additional read cache tier. VMs leverage the Client Cache of the host they are running on. VM reads are accelerated since the data can be accessed from the host memory, which is a shorter I/O path than accessing the data from disk. Client Cache extends DRAM caching of CBRC to linked clones, App Volumes, and other non-replica components.



CBRC and vSAN Client Cache are compatible. When data is cached in CBRC, a read will be served out of the CBRC and the request will never hit the vSAN layer. However, when there's a CBRC-miss, the system will check the vSAN Client Cache first before going to disk. In essence, vSAN Client Cache becomes a L2 cache for CBRC. The benefits of using double cache is greater when the VDI workload is read-heavy.

Recommendation: *Always enable CBRC in conjunction with vSAN Client Cache (enabled by default) for optimized client-side caching in Horizon 7 on vSAN environments.*



References

White Papers

For additional information, see the following white papers:

- [VMware vSAN 6.2 Space Efficiency Technologies](#)
- [VMware Horizon 7 on VMware vSAN 6.2 All-Flash](#)
- [VMware Horizon 7 Enterprise Edition VMware Validated Integration Design](#)

Product Documentation

For additional information, see the following product documents:

- [VMware vSAN Documentation](#)
- [VMware vSAN Documentation on StorageHub](#)
- [VMware Horizon 7 Documentation](#)



About the Authors

Angela Ge, Product Line Manager in the VMware End User Computing group

Kristopher Groh, Senior Product Manager in the VMware Storage and Availability group

Sophie Yin, Senior Solutions Architect in the VMware Storage and Availability group

Catherine Xu, Senior Technical Writer in the VMware Storage and Availability group

