VMware vCloud® Architecture Toolkit™
for Service Providers

Architecting the Digital Workspace with VMware Horizon® 7

Version 2.9

January 2018

Ray Heffer

VMware, Inc.
3401 Hillview Ave
Palo Alto, CA 94304
www.vmware.com

Table of Contents

List of Tables

List of Figures

Executive Summary

The VMware Cloud Provider™ Program is a global network of over 4,200 service providers who have built their cloud and hosting services on VMware software solutions. These service providers deliver world-class cloud and hosting services to their customers across the globe.

The VMware vCloud® Architecture Toolkit™ for Service Providers (vCAT-SP) provides architectural guidance and best practices on VMware-based solutions. Using real world implementations, use cases, and actual customer requirements, this document enables service providers to define their service offering, ready for the demands of enterprise-class customers.

VMware Horizon® is available as VMware Horizon DaaS®, VMware Horizon Air, VMware Horizon Air Hybrid-Mode, and VMware Horizon 7, which is targeted at the enterprise. Horizon DaaS (desktop as a service) is targeted at service providers since it provides a multi-tenancy desktop platform, meaning a single instance of the solution can provide desktops and applications (on-demand) to multiple customers across multiple data centers. This allows service providers to offer a monthly consumption cost model, allowing each tenant to scale up or scale down the number of desktops they require as business demands change, while minimizing initial CAPEX investment. In the realms of cloud computing, this is often referred to as elasticity, and is one of the primary drivers behind DaaS.

Horizon 7 provides a holistic EUC solution architecture designed for a single tenant and unique capabilities often in demand by the modern enterprise.

One of the distinctions is Cloud Pod Architecture, which allows tenants to scale within the data center, and even across multiple data centers to meet requirements of a mobile workforce or even disaster recovery (DR).

Service providers can leverage Horizon 7 as a dedicated “managed” service offering for tenants with requirements that include, but are not limited to:

- Customer managed desktop pools for multiple business units or use-cases
- Departmental Role Based Access Control (RBAC) for organizational compliance
- Extension of existing virtual desktop infrastructure to cloud service provider (Cloud Pod)
- Linux virtual desktops
- GPU powered desktops with vGPU
- Organization retains control of desktop provisioning, application packaging and monitoring
- Integration with VMware Identity Manager™ to deliver a user-facing application portal and single sign-on to SaaS and Windows applications.

Scope

This document is intended for architects and consultants involved in planning, design and deployment of Horizon 7 in order to provide service providers with a digital workspace solution offering. This document is not a reference architecture or validated design. The purpose of this document is to provide guidance for designing the digital workspace specifically for service providers, adopting guidance and real-world best practices.

Horizon 7 uses a “building block” approach to allow enterprise customers to scale, while minimizing support costs and deployment risks. Service providers, however, must be able to provide each “building block” or instance as a managed service to its customers, while maintaining a distinct advantage to enterprises using Identity Management, Single Sign-On (SSO), VMware User Environment Manager™, and application delivery mechanisms.

The “building block” approach is based on real-world implementations from many of the largest VMware deployments in production today. This document will break these down into single-tenant “Pods” that can be leveraged by service providers offering a managed service, and providing a solution that includes:

- Fully managed digital workspace architecture allowing tenants to provide single sign-on, SaaS applications and Windows applications via a dedicated, branded organization application catalog.
- Repeatable design that can be delivered to multiple enterprise tenants and tailored to organizational requirements.
- Horizon 7 foundational building block design with add-on services for Identity Manager, VMware App Volumes™, User Environment Manager, and VMware ThinApp®.
- Tenant solution monitoring and reporting with VMware vRealize® Operations™ for Horizon

Introduction to Horizon 7

VMware Horizon 7 is an enterprise end-user computing platform for delivering virtual desktops and applications, scaling from a handful of desktops to tens of thousands across multiple data centers. Traditional VDI (Virtual Desktop Infrastructure) contained a central component called the broker, which is responsible for "brokering" connections to the desktop, tracking the state (logon, logoff, power on, shutdown), in addition to maintaining desktop entitlements through user and group membership. Horizon 7 still retains the desktop broker (Connection Server), but there are many other components that deliver the entire Horizon solution stack.

3.1 Defining the Digital Workspace

Delivering Desktop-as-a-Service (DaaS) can often be a complimentary offering to tenants that already host server workloads with a service provider. DaaS has traditionally been associated with Virtual Desktop Infrastructure (VDI) and not the digital workspace that brings a full suite of applications (Windows, SaaS, and mobile) to the end-user. However, it is important to distinguish that the VMware vision of the "digital workspace" is not exclusive with Horizon 7. To make the digital workspace possible, products such as VMware Identity Manager, App Volumes and User Environment Manager are all able to co-exist with either Horizon 7 or Horizon DaaS and even other desktop platform offerings.

The digital workspace represents a consumer orientated, self-service platform that can be used on any device, delivering an organizations portfolio of applications. It addresses a new generation of end-user computing that is no longer confined to the virtual desktop.

Horizon 7 combined with Identity Manager, App Volumes and User Environment Manager marks a suite of products targeted at delivering the digital workspace for enterprises. The benefit to service providers is an enterprise scale platform that can finally be offered to enterprise cloud tenants, bridging the gap between existing "VDI" desktop offerings and on-premises deployments.

For service providers, it is important to take a "bird's eye" view of the entire solution stack that serves the digital workspace with Horizon 7. This document breaks down the digital workspace into five distinct layers, which have a direct correlation to tenant-facing functionality, service provider boundaries (for instance, firewall ports, user portal integration), core and management infrastructure.

Figure 1. Digital Workspace Solution Layers

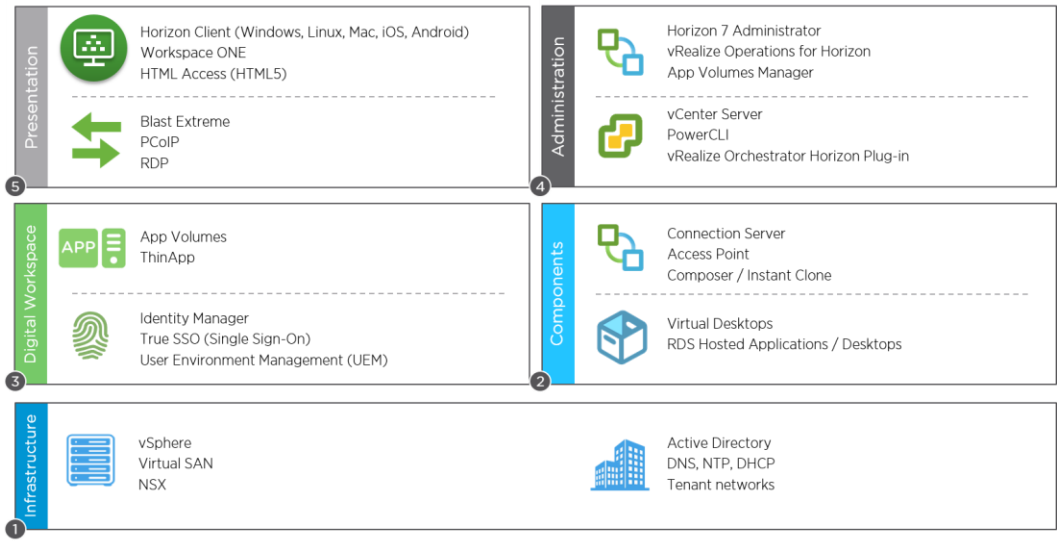

1. Infrastructure Layer

   This layer decouples datacenter constructs into the software-defined data center (SDDC) and provides virtual infrastructure compute resources, networking, and storage. This also includes tenant infrastructure: Active Directory, DNS, DHCP, and customer networks.

2. Component Layer

   This layer contains the individual components from Horizon 7, including the virtual desktop and Remote Desktop Session Host (RDS Host) infrastructure that hosts applications and desktops.

3. Digital Workspace Layer

   For the scope of this document, the digital workspace layer consists of identity management and application delivery.

4. Administration Layer

   Management of the platform requires access to these components.

5. Presentation Layer

   This layer consists of the client-facing user interface (UI) and protocols that provide access to the digital workspace.

VMware Horizon Client Architecture

Starting with the presentation layer, client access uses the VMware Horizon Client™, HTML Access using HTML5 and remote display protocols, including Blast Extreme.

Figure 2. Horizon Presentation Layer

4.1 Horizon Client

The VMware Horizon Client provides end-users with access to both desktops and applications in a Horizon 7 environment, and is available on desktop PC's, Thin Clients, and mobile devices on multiple operating systems (Windows, Mac, iOS, Linux, and Android).

The connection flow of the Horizon Client is the same with Horizon 7, Horizon Air or Horizon DaaS. External access to Horizon 7 is either facilitated using VMware Access Point™ virtual appliances in the DMZ, or Security Server which runs on Windows Server.

For the purposes of this document, Access Point will be used as the preferred security gateway since it has many advantages for service providers. Unlike Security Server which must be "paired" with a Connection Server, Access Point can be independently scaled, with multiple instances residing behind a load-balancer without the need for pairing. In addition, Access Point handles up to 2,000 active sessions per appliance, and it is recommended at least two are deployed for availability and load distribution.

When the user launches the chosen desktop or application pool, Access Point will communicate on HTTPS (TCP 443) to receive the desktop VM IP from the Connection Server. The role of the PCoIP Gateway on the Access Point appliance is to then forward the PCoIP connection to the IP address of the Horizon Agent. When the Blast Extreme protocol is used, it uses a secure WebSocket on HTTPS.

Note Security Server uses JMS, IPsec and AJP13 (see Section 6.1.1.6 Security Server Framework), but Access Point does not use these protocols (JMS is still used on the Connection Servers). If you refer to the network ports diagram (see Figure 21. NSX Firewall Policy Example Rule-Set), you'll see this resides in a dotted line to illustrate this.

4.2 HTML Access

HTML Access is a client-less feature that leverages the Blast protocol and allows users to connect to desktops or applications using an HTML5 compliant web browser.

Figure 3. HTML Access

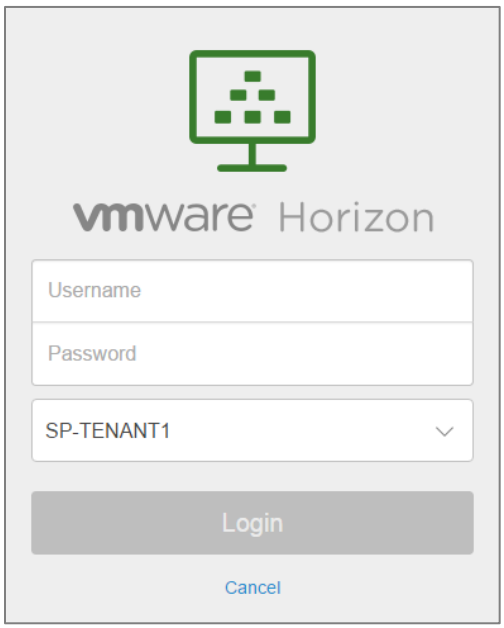

HTML Access uses the Blast protocol, and unlike Blast Extreme that leverages the Horizon Client for full feature functionality, HTML Access is TCP only. One of the benefits for service providers is the ability to present desktops and applications to tenants without requiring end-users to download and install the Horizon Client.

HTML Access also supports:

- Single sign-on (SSO) / True SSO
- RSA SecurID
- RADIUS
- Location-based printing

Note For more details on HTML Access, please refer to the document titled “Using HTML Access” in Section 10, References.

4.3 Blast Extreme Protocol

Blast Extreme is an enhanced remote session protocol introduced with Horizon for Linux desktops, Horizon 7 and Horizon DaaS. One of the benefits of Blast Extreme is that it can be used with the Horizon Client, in addition to HTML Access using HTML5. One of the other significant features of Blast Extreme is the ability to offload the CPU cycles used for decoding by using H.264 hardware decoding.

Blast Extreme is a TCP-based protocol, but it can also be configured to use UDP, unlike PCoIP which is UDP only. TCP is the most versatile as it's not likely to be blocked on customer firewalls, whereas UDP is sometimes filtered in some locations (for example, public WiFi or guest networks). See Figure 21 for a full list TCP/UDP ports.

Note Consider that zero-client support might be limited, and older zero-clients might only support PCoIP. If zero-clients are to be used, tenants must be sure that the Blast Extreme protocol is supported.

Figure 4. Horizon Client Blast Extreme Connection Flow

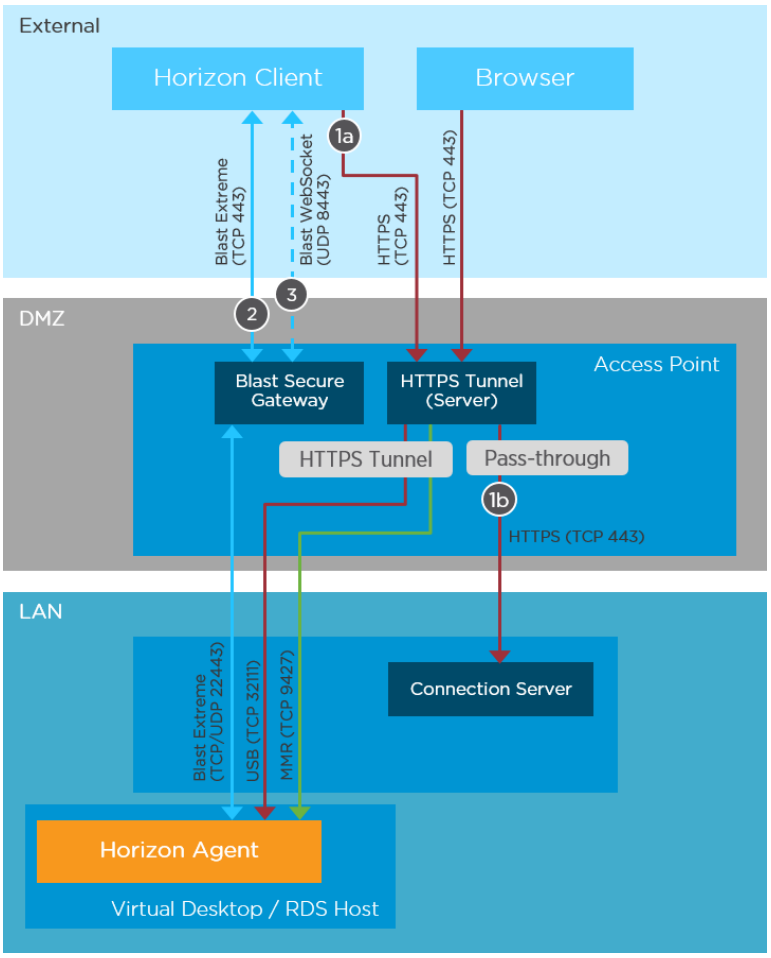

1. The Horizon Client sends authentication credentials using XML-API over HTTPS to the external URL on the Access Point appliance (or Security Server). This is typically via a load-balancer VIP (Virtual IP).
   a. HTTPS Authentication data is passed-through from Access Point to the Connection Server. Any entitled desktop pool(s) are then returned back to client.

Note If Security Servers are used, they must be paired with a Connection Server. Access Point does not require pairing with Connection Servers, which provides greater flexibility without the need to dedicate Connection Servers for external or internal access.

2. The user selects a desktop or application pool entitlement, and a session handshake occurs over HTTPS (TCP 443) to Access Point / Security Server. A secure WebSocket is then established (TCP 443) for the session data between the Horizon Client and the Access Point / Security Server.
3. If configured to use UDP, the Blast Secure Gateway service will attempt to establish a UDP WebSocket connection on 8443. If this fails, due to a firewall blocking the UDP port, then the initial WebSocket TCP 443 connection will be used instead.

Client Drive Redirection (CDR) and Multimedia Redirection (MMR) are encapsulated using HTTPS (TCP 443) from the Horizon Client to Access Point or Security Server. The HTTPS Secure Tunnel connects to the Horizon Agent on TCP 9427 for MMR and CDR traffic.

The client to server port can be configured to use a side channel by configuring the following registry key on the guest OS: `HKEY_LOCAL_MACHINE\SOFTWARE\VMware, Inc.\VMware TSDR\tcpSidechannel`

`tcp` - CDR over TCP Sidechannel
`vvc` - CDR over VVC sidechannel in Blast & PCoIP – Default (Horizon Agent 7.0.2)
`PCoIP` - CDR over TCP sidechannel in Blast & PCoIP
`vchan` - CDR over VVC/PCoIP sidechannel.
`none` - CDR over main channel

4.4 PCoIP Protocol

PCoIP by Teradici, is a real-time protocol that is very well established with VMware Horizon, and is also supported by a large number of third-party zero-client devices. As a protocol, it is supported on both Horizon DaaS and Horizon 7 environments.

Figure 5. Horizon Client PCoIP Connection Flow

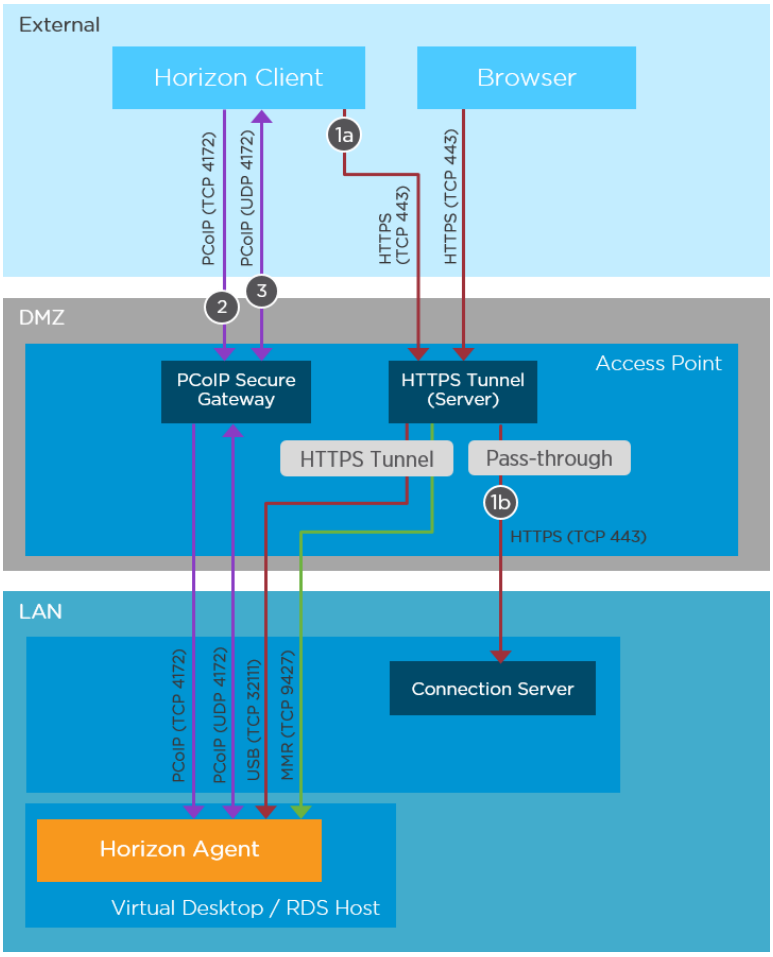

1. The Horizon Client sends authentication credentials using XML-API over HTTPS to the PCoIP external URL on the Access Point appliance. This is typically via a load-balancer VIP (Virtual IP).
   a. HTTPS Authentication data is passed-through from Access Point to the Connection Server. Any entitled desktop pool(s) are returned back to client.

Note If Security Servers are used, they must be paired with a Connection Server. Access Point doesn't require pairing with Connection Servers, which provides greater flexibility without the need to dedicate Connection Servers for external or internal access.

2. The user selects a desktop or application, and the connection is initiated on TCP 4172 to Access Point / Security Server. This is the PCoIP session handshake.
3. A bi-directional PCoIP connection is then established on UDP 4172 for the session data between the Horizon Client and the pcoipExternalUrl for Access Point / Security Server. The PCoIP session is forwarded between Access Point and the brokered virtual desktop (Horizon Agent).

Note `pcoipExternalUrl` is used for Access Point. When Security Servers are used in a Horizon solution, the PCoIP External URL configured on the paired Connection Server will be used.

4.5 Device Redirection

Horizon 7 allows for the redirection of client devices and multimedia traffic. This is not an exhaustive list of device redirection available with Horizon 7, others not listed here include Real-Time Audio-Video, Flash Redirection, Scanner Redirection, Serial Port Redirection, URL Content Redirection and Flash URL redirection.

4.5.1 Client Drive Redirection

Client Drive Redirection (CDR) is useful for users that wish to redirect a local or network drives from the client to the virtual desktop or application. File type association also allows for local files to launch the virtual application rather than local instance of the application. This setting can either be enabled/disabled by the client or enforced by the administrator via a client Group Policy Object (GPO), at a server level or farm wide. CDR is supported on Windows, Mac OS and Linux clients.

Figure 6. Accessing local drives with CDR

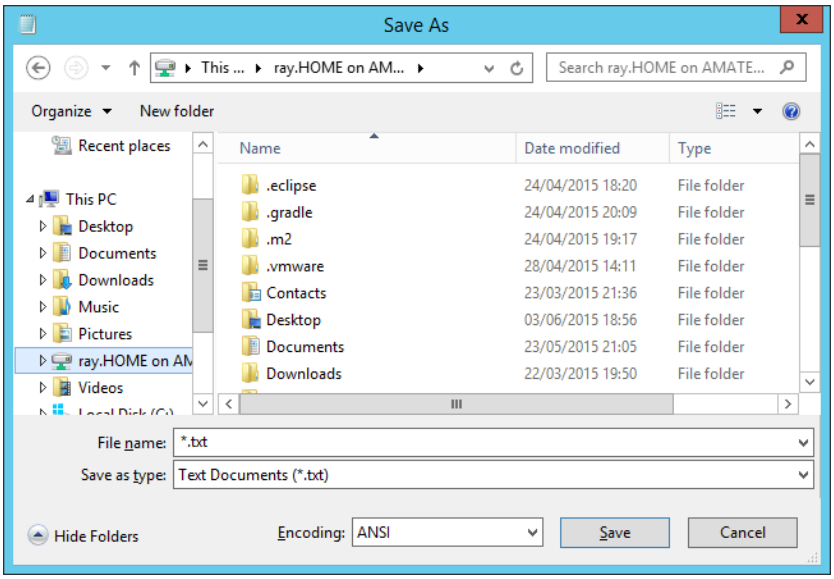

As described in the previous section, when Blast Extreme is used as the remote / client protocol, CDR will be established inside a View Virtual Channel (VVC) over HTTPS (TCP 443). If HTTPS Secure Tunnel is not configured, CDR will use TCP 9427 between the client and server.

For external client connections (internet) it is recommended to use the VVC for CDR since it uses the existing HTTPS 443 connection, eliminating the need to open another firewall port.

4.5.2 USB Redirection

This feature allows end-users to connect their removable devices, such as USB flash drives, cameras, and headsets to the virtual desktop session.

When the HTTPS Secure Tunnel is configured, see Figure 4 and Figure 5, USB traffic is tunneled using HTTPS (TCP) port 443. Where tenants do not require external (internet) access and direct (WAN) connections are used, no Access Point will be required and USB traffic between the Horizon Client and Horizon Agent (desktop) will use TCP 32111.

4.5.3 Multimedia Redirection (MMR)

By default, MMR is disabled, however the Tenant Administrator can enable this using the global policy in Horizon 7, allowing for enhanced multimedia playback on the local client device rather than host. This can reduce CPU usage on the host server during multimedia playback.

MMR will adapt to network conditions on Windows 10 and Windows Server 2012 or 2012 R2 or later (Windows Server 2016 is tech-preview at the time of writing this document). If the network latency between Horizon Client and the remote desktop is less than 30 milliseconds, the video is redirected with MMR. If the network latency is 30 milliseconds or higher, the video is not redirected, instead it is rendered on the ESXi host.

Note MMR uses TCP port 9247 for direct connections where Access Point is not used, however it is recommended that HTTPS Secure Tunnel is configured which places MMR traffic in HTTPS (TCP) 443.

Horizon Pod and Block Design Methodology

A VMware Horizon Pod is a single instance of Horizon (View) that is configured for a cloud tenant. The Pod and block approach is especially useful for implementations that start small, but plan to scale out in the future. For larger implementations, Pod and block allows for massive scale, up to 50,000 concurrent connections across the entire Cloud Pod.

A single resource (desktop) block is delineated by a vCenter Server®, and contains one or more vSphere resource clusters containing virtual desktops or Microsoft RDS (Remote Desktop Session) hosts. A separate management cluster hosts all Horizon server components including:

- vRealize Operations for Horizon
- VMware vCenter Server Appliance™
- VMware Identity Manager™
- Enrollment Server (True SSO)
- VMware App Volumes™ Manager
- Connection Servers
- Security Server (not required if Access Point is used)
- Access Point
- VMware NSX® Manager™
- Microsoft SQL
- Active Directory
- File Servers (User Environment Manager, ThinApp repository, home directory file shares)

5.1 Deployment Models

Pod and block methodology is an architecture model (see the following figure) often used in enterprise deployments of Horizon 7. The Pod and block architecture model can be adapted from an enterprise deployment, to a service provider managed implementation. In this case, each tenant would represent a resource block, each block delineated by their own vCenter Server, and one or more vSphere clusters providing desktop and application resources, in addition to VMware vSAN™ storage and NSX networking.

Figure 7. VMware Horizon Pod Example

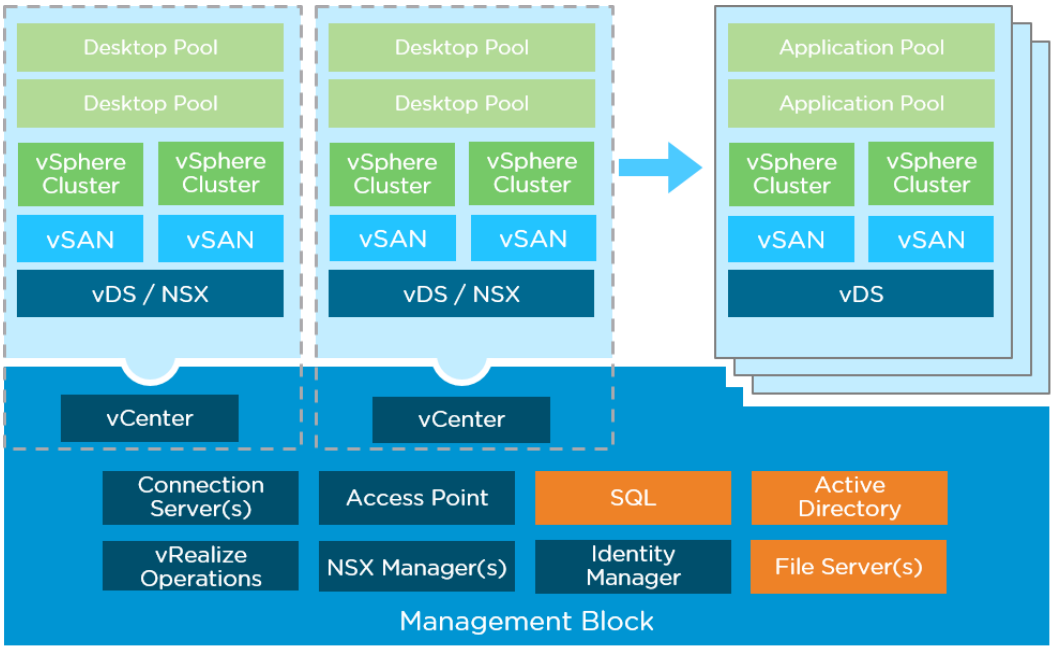

Note It is critical that a single "Pod" resides on the same well-connected LAN / data center, and it must not span a wide area network (WAN). This is so that the message bus, which leverages JMS (see Section 6.1.1.4, Message Bus), is not hindered with high latency. For multi data center deployments, use Cloud Pod Architecture. (See Section 5.2, Cloud Pod Architecture.)

5.1.1 Management Block

Service providers can offer either a shared or dedicated management block, depending on the tenant service type.

While it is technically feasible to run the entire solution (management and desktop resources) on a single vSphere cluster, this would only be recommended for Proof-of-Concept (PoC) or test environments. For production workloads, it is important to maintain separation of management and workload resources.

There are many reasons to maintain separation of the management cluster to production desktop and application resource clusters. In the event of a power outage or other disaster, the management block is the first environment that is powered on. This provides a high level of certainty that the environment can be managed, and will not be impacted by resource demands and tasks used for desktop virtual machines (for instance, hundreds or thousands of VMs powering on). Even during normal day-to-day operations, the separation of management and resource clusters provides a better guarantee of performance, resulting is a greater quality of service.

5.1.1.1 Dedicated

A dedicated management block is provisioned for a single tenant and is more suited to tenants where the administrator requires access to the management vCenter Server, and is responsible for administering all components including the vSphere infrastructure.

5.1.1.2 Shared

Using this model, a shared management block is managed in its entirety by the service provider, and only the service provider would have access to the management block vCenter Server. Within the management cluster, resource pools separate each tenant's virtual machine workloads (see the following figure) which includes the tenant Connection Servers and Access Point appliances. In addition Network I/O Control (NetIOC) is recommended to prioritize management network traffic between tenants.

The tenant management block hosts each vCenter Server responsible for their dedicated resource blocks. The dedicated resource blocks host the tenant desktop and application workloads.

Figure 8. Shared Management vSphere Cluster

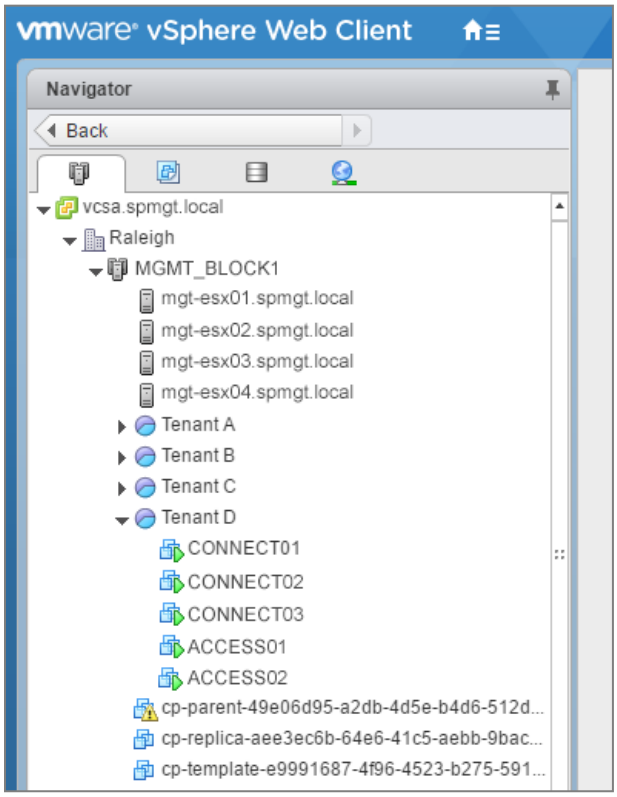

Consider, however, that using this approach means that unexpected downtime or maintenance of the entire management block would impact all tenant management functions. Availability of both management and resource blocks are critical in this model. If a particular tenant requires a strict Service Level Agreement (SLA) for availability, then a dedicated Management block may be deployed for that tenant.

For both service providers and enterprises alike, vSAN Ready Nodes are ideal and an online configurator (http://vsanreadynode.vmware.com/RN/RN) is available to help select the most suitable vSAN Ready Node hardware for your environment.

Note A good starting point for the number of tenants per management block is 4. This allows for ease of management and lowers the impact of outages and maintenance to the management cluster, whether planned or unplanned.

5.1.2 Resource Block

Sizing for each tenant resource block may vary, ideally between 4 and 32 hosts (see the following figure). A single resource block is dedicated to a tenant regardless of the management block model.

For both service providers and enterprises alike, vSAN Ready Nodes are ideal and an online configurator (http://vsanreadynode.vmware.com/RN/RN) is available to help select the most suitable vSAN Ready-Node hardware for your environment.

Figure 9. Horizon Resource Block

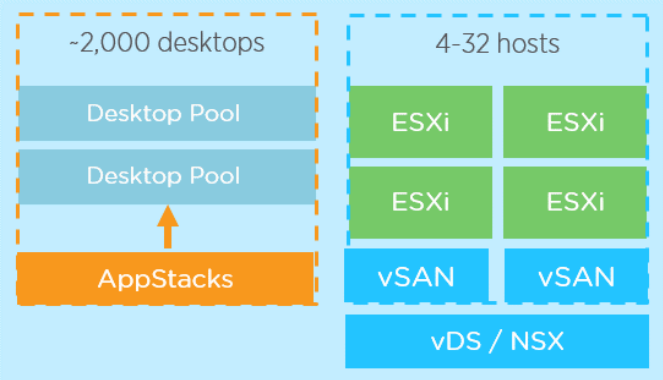

A resource block is delineated by a dedicated vCenter Server, which contain one or more vSphere clusters. While vCenter Server supports up to 10,000 powered-on virtual machines, the Pod and block design methodology divides a Horizon "Pod" into multiple blocks with a maximum of 10,000 concurrent sessions per Pod. Pod and block designs of five blocks (2,000 desktops per block) have proven to be the most effective in meeting manageability requirements and scale.

Note vSphere 6.5 supports up to 64 hosts in a single HA/DRS cluster, and the current supported maximum for Horizon 7 is 32 hosts. If there is a need for a larger cluster size, it is recommended that you contact your VMware representative and file an RPQ (Request for Product Qualification). This is subject to change and is correct as of the date of this document.

5.2 Cloud Pod Architecture

A Pod represents a single instance of Horizon, which is made up of one or more resource blocks (vCenter Server) and a collection of Connection Servers. As previously described in Section 5.1 Deployment Models, it is not recommended or supported to span a single Pod across a wide area network (WAN) due to the latency impact on the JMS message bus. Cloud Pod Architecture addresses this by federating multiple “Pods” (up to 25) in a single Cloud Pod. See the following figure.

Figure 10. Cloud Pod Architecture

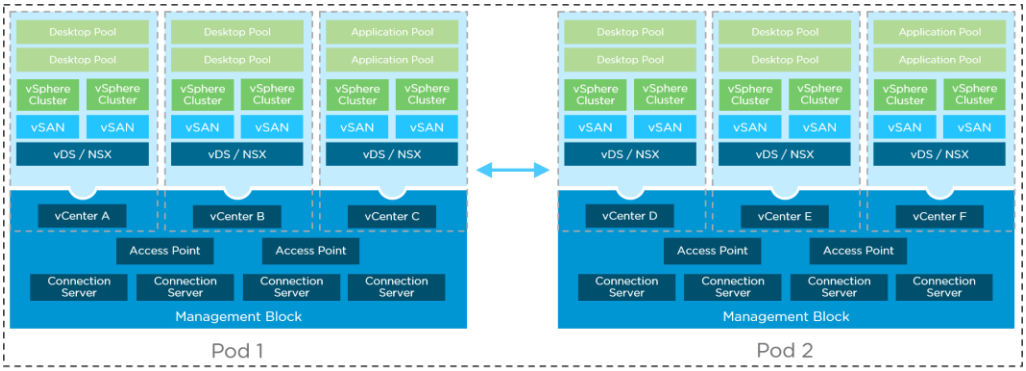

For cloud tenants, a Cloud Pod Architecture (CPA) implementation may be useful when extending an existing on-premises deployment of Horizon to the service provider. To provide CPA functionality, the first Pod must be initialized and then additional Pods can be joined to create a Pod federation.

It is also possible to initialize CPA for a single Pod. This enables a useful feature called Global Entitlements, which allows for a single entitlement to be assigned multiple desktop or application pools. In order to understand this, it is important to understand how local and global entitlements are presented to the end user.

Table 1. Cloud Pod Architecture Maximums

Description	Value
Maximum number of sites	5
Maximum number of tenant Pods in a federated Cloud Pod	25
Concurrent connections per Pod	10,000
Concurrent connections per Cloud Pod	50,000

Figure 11. Horizon Client Pool Entitlements

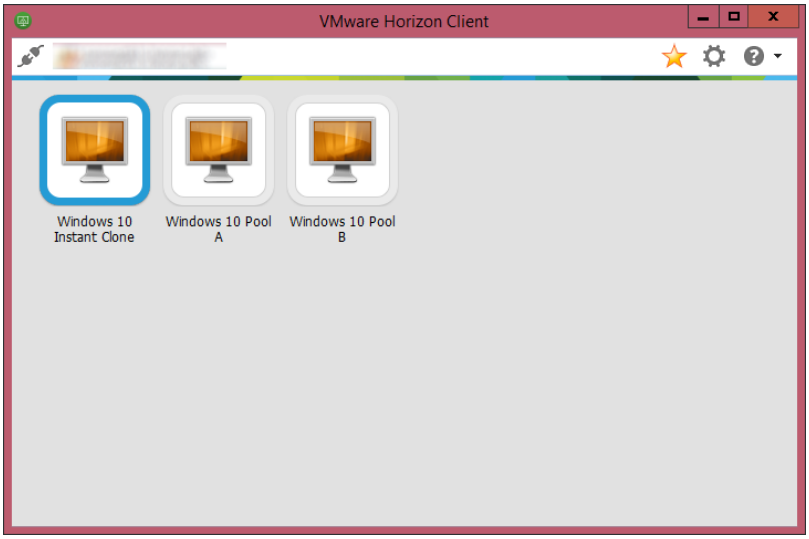

A local entitlement allows the Tenant Administrator to entitle access to a desktop or application pool to an Active Directory user or group object. When the user launches the Horizon Client or HTML Access, they will see the desktop or application pool entitlement. If multiple desktop pools are added, these will each show separately to the end user. For example, if a user is entitled to desktop pools A and B, they will see both (see Figure 11).

A global entitlement appears the same as a local entitlement to the end user, however, multiple desktop pools from the same Pod or multiple Pods can be added.

When a user connects to a global entitlement, if the first pool has no spare desktops available it will try the next pool within the local Pod. If there are still no desktops available, it will continue to the next Pod (same site) then eventually Pods in other sites providing they are part of the same global entitlement in the Pod federation.

Table 2. Cloud Pod Architecture Terminology

Term	Description
Site	Typically, a data center site. A site will contain a logical grouping of well-connected resources, exhibiting low-latency, high bandwidth, reliable network connections.
Global Entitlement	A global entitlement of desktop or applications pools across the federated Pod.
Scope (Within Pod, Site or Any)	A scope is defined as part of the global entitlement and can be set to Within Pod, Within Site or All Sites. When a global site load-balancer places the user connection, the scope then determines how to handle traffic across sites. Within Pod places the desktop in the local (current) Pod only. Site, will place the desktop connection to any available Pod in the same site. Any, will place the desktop connection anywhere in the federated Cloud Pod.
Home Site	A home site is associated with a user or group and is selected for a Global Entitlement. For example, a user or group's home site can be set to London for a Global Entitlement. If a home site is associated with a user, this will override any group home site assignments.

5.2.1.1 Load Balancing

Cloud Pod Architecture is not a load-balancing solution, and therefore if multiple sites are configured in a federated Cloud Pod, a global site load balancer must be implemented. The global load balancer must be configured with a single namespace for the tenant (for instance, `tenantpod1234.myvmware.com`) which will direct the user connection to one of the Access Point appliances in either site.

Local load balancing must be in place for each service provider data center, in order place connections on an available Access Point appliance for a given site. Depending on the load balancer used, a number of algorithms can be configured, but in most cases “least connections” is preferred.

It is also recommended that the load-balancer monitors the health of each Access Point appliance by sending an HTTP GET /favicon.ico request. This will check for the presence of /favicon.ico served by the Access Point web server. If a "200 OK" response is received, the load balancer will consider the Access Point node to be heathy. If there is no response, then the load balancer will consider it offline and not send further connections to the offline Access Point appliance.

Some of the most common methods for providing a load balanced external connection is to use a single virtual IP (VIP), provide multiple VIPs, or use the single VIP with port forwarding.

For more information on load balancing please refer to “Load Balancing with VMware Access Point” (see Section 10, References).

Digital Workspace Platform

While Horizon 7 plays an integral role in delivering the digital workspace, it is the additional functionality of the other components listed in this section that provides the entire solution. Referring back to Figure 1, these components can be separated into 5 distinct layers, and this section will focus on layers 2 (Components), 3 (Digital Workspace) and 4 (Administration).

6.1 Horizon 7 View Connection Servers

An instance of Horizon (a Pod) consists of Connection Servers which are responsible for the brokering of desktops and applications, entitlements, Tenant Administrator UI (VMware Horizon 7 Administrator) and tracking state change information between all nodes in the Pod.

Figure 12. Connection Server Architecture

As illustrated above, a Connection Server contains gateway services, in addition to the framework and communication components.

6.1.1.1 PCoIP Secure Gateway

The PCoIP Secure Gateway runs as a service on each Connection Server. When Access Point (or Security Server) is used, the PCoIP Secure Gateway forwards connections to the Horizon Agent on the virtual desktop.

6.1.1.2 Blast Secure Gateway

The Blast Secure Gateway service also runs on each Connection Server as server-side Node.js (JavaScript runtime environment) processes.

6.1.1.3 VDMDS

This is a directory service that provides replication across the Pod (not between Pods) using Active Directory Lightweight Directory Services.

6.1.1.4 Message Bus

The message bus leverages JMS (Java Messaging Service) which is a Java API that allows Connection Servers to create, send and receive messages using reliable, asynchronous low-latency communication. Tenants can configure the message security mode, allowing for message signing and encryption and/or SSL for all server communication.

6.1.1.5 View Framework Node Manager

The View Framework Node Manager (wsnm.exe) is a process that runs on each Connection Server, and provides event logging, security, and COM+ framework services.

6.1.1.6 Security Server Framework

To facilitate secure communication between the Security Server(s) and Connection Servers, JSM, AJP13 (Apache JServ Protocol version 1.3) and IPSEC (optional) are used. These protocols are not required with Access Point as it uses standard HTTPS communication.

6.2 Desktop Pools

Tenant administrators have the ability to create desktop pools using either Linked Clones, which requires a View Composer Server, or Instant Clones. Service providers can make one or both of these technologies available, however there are some limitations to be aware of:

6.2.1 Linked Clones

Linked Clones provide administrators with the ability to clone a pool of desktops from a master image. In order to enable this functionality, a View Composer server is deployed, which consists of a web service (SIM), the Universal File Access (UFA) service and other components, including the vCenter Server API. View Composer receives instructions from Horizon as XML messages, and initiates the linked clone creation using vCenter Server API calls and the SIM service (Scalable Image Management). It is essentially a workflow engine.

When a Linked Clone desktop pool is created, Composer will run a number of parallel clone creation tasks, defined by the maximum setting in “Max concurrent View Composer provisioning operations” (default is 8).

6.2.2 Instant Clones

Instant Clone technology is a rapid clone mechanism available with Horizon 7 Enterprise Edition. One of the benefits of Instant Clones is that is doesn’t rely on the Composer Server, instead it uses vmFork, supported with vSphere Version 6.0 Update 1 or later. Despite some of the limitations listed in Table 3, it is significantly faster than View Composer Linked Clones.

Table 3. Linked vs Instant Clones

Feature	Linked Clones	Instant Clones
Windows 7	✓	✓
Windows 8	✓	
Windows 10	✓	✓
Windows Server 2008 / 2012 / 2012 R2	✓	
RDSH	✓	
Floating Desktop Pools	✓	✓
Dedicated Desktop Pools	✓	
App Volumes	✓	✓
Persistent or Disposable disks	✓	
Linux		
Multi-display	✓	✓**
vDGA/vGPU Support	✓	
IPv6	✓	
Multi-VLAN (network labels)	✓	
Local host datastores	✓	
VMware PowerCLI™ support	✓	
Reboots required for cloning	2	0

* Applies to Horizon 7.0, current at the time of this document. Check release notes of future releases to determine whether this has changed.

** Instant Clones are limited to 2 monitors with a maximum resolution of 2560 X 1600

6.3 Access Point

Access Point is a hardened Linux-based virtual appliance primarily known for its role in Horizon DaaS implementations as an internet-facing gateway for cloud tenants. In previous versions of Horizon, View Security Servers facilitated remote access to the VDI environment. Security Servers need to maintain a pairing with a View Connection Server, so typical implementations consist of at least two Connection Servers for internal connections, and two Connection / Security Server pairs for internet-facing connections. One of the limitations encountered with this architecture, which makes Access Point far more desirable, is the maximum of 7 Connection Servers in a Pod (See Table 5. Horizon 7 Configuration Maximums).

One of the key strategies behind Access Point is to adopt a secure virtual appliance architecture that can be used across both Horizon DaaS and Horizon 7, since these both have very similar external gateway requirements. For service providers, Access Point also allows for horizontal scaling and importantly deployments can be automated (See Section 10, References for further information on deployment), in addition to configuration using RESTful API.

Since service provider deployments of Horizon 7 with Access Point are not bound to Connection Servers, this allows for a very flexible approach to scaling out or scaling back. Access Point appliances can be deployed or destroyed as required.

Finally, security is an important consideration for enterprise tenants. Many tenants will not allow any authentication traffic to enter the tenant environment, and Access Point, residing in a DMZ, allows for multiple authentication methods including RSA SecurID, RADIUS and smart card.

Note It is recommended that Access Point deployment and configuration is automated. This can be achieved using a PowerShell script so that Access Point is production ready on first boot.

6.4 Desktop and Application Virtualization

There are two ways to present applications to end-users with Horizon 7. There is traditional VDI, where the desktop exists as a virtual machine and applications are loaded as part of the base image, and there is RDSH (Remote Desktop Session Host), also referred to as session-based computing. RDSH allows applications installed on the RDS host to be presented to users using the Horizon Client, HTML Access or through VMware Workspace™ ONE™.

6.4.1 App Volumes

App Volumes allows for the lifecycle management and delivery of applications to end-users in an enterprise tenant environment. This allows cloud Tenant Administrators to manage, deliver, upgrade and maintain applications throughout an applications lifecycle. Application delivery is performed near real time, which is also referred to as Just-in-Time (JIT) delivery.

Figure 13. VMware App Volumes

As illustrated, applications can be pre-installed inside one or more AppStacks. Each AppStack is a read-only virtual disk (VMDK), that can be assigned to multiple virtual machines (RDSH or a VDI desktop). In addition, a writable AppStack otherwise known as a writable volume, can be delivered to individual users using 1:1 mappings to allow user-installed applications and customization.

Table 4. App Volumes Architecture Components

Component	Description
App Volumes Manager	Provides a single-tenant dashboard console for user assignments and configuration. The App Volumes Manager is also responsible for agent monitoring and brokering.
App Volumes Agent	The agent redirects file system and registry read operations to the AppStack, in addition to sending changes to a writable volume (if configured). The agent reports its current status to the App Volumes Manager.
AppStack	An AppStack is a read-only volume containing pre-installed data and applications. A single AppStack can be mounted to one or more virtual machines.
Writable Volume	A writable volume is a virtual disk assigned to a single user virtual machine, allowing the user to install their own applications and capture application settings in the user profile.
Provisioning Virtual Machine	This is a clean Windows virtual machine, installed with the App Volumes Agent, and is used to capture the installation or update of applications to an AppStack.
Policy file (snapvol.cfg)	The policy file (snapvol.cfg) resides on the root folder of each AppStack and writable volume, and allows for the exclusion of file and registry locations. This is created by default, but can be customized.

Note At the time of this document, App Volumes 2.11 is the recommended version for Horizon 7, and this includes support for Instant Clone desktops.

6.4.2 VMware User Environment Manager

Using a combination of App Volumes and User Environment Manager, tenants are able to deliver a minimal desktop image, yet provide a persistent desktop experience to end-users. One of the main advantages of User Environment Manager is the minimal infrastructure footprint required.

As you can see in the logical architecture below, components of User Environment Manager do not require any significant infrastructure and can be fully managed by the tenant.

Figure 14. VMware User Environment Manager Architecture

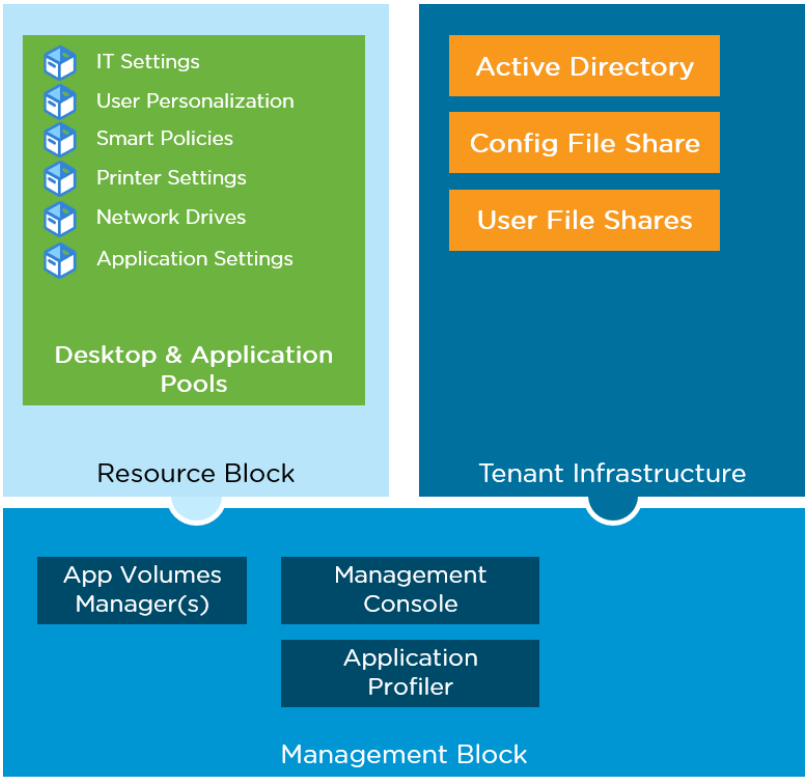

Features such as clipboard access, USB, printing and drive redirection can be restricted based on the user location. For example, Tenant Administrators can allow remote access over the internet, and setup Horizon policies to disable device redirection features unless users connect from an office local area network.

6.5 Digital Workspace Architecture

The digital workspace provides end-users with single-sign on (SSO) access to a variety of applications, such as Windows, SaaS (Software-as-a-Service) and mobile applications. While VMware Horizon provides both desktops and applications to end-users using the Horizon Client or HTML Access, as a stand-alone product it doesn't offer the "digital workspace" experience with features such as True SSO, or SaaS integration unless it is deployed with Identity Manager.

For VMware Cloud Providers, extending beyond the traditional desktop-as-a-service (DaaS) offering can be achieved using Horizon 7, VMware Identity Manager, App Volumes and User Environment Manager.

Note The scope of this document is focused on Horizon 7, Identity Manager, App Volumes and User Environment Manager for service providers offering the digital workspace. Enterprise Mobility Management (EMM) which is part of Workspace ONE and AirWatch is not part of the scope of this document.

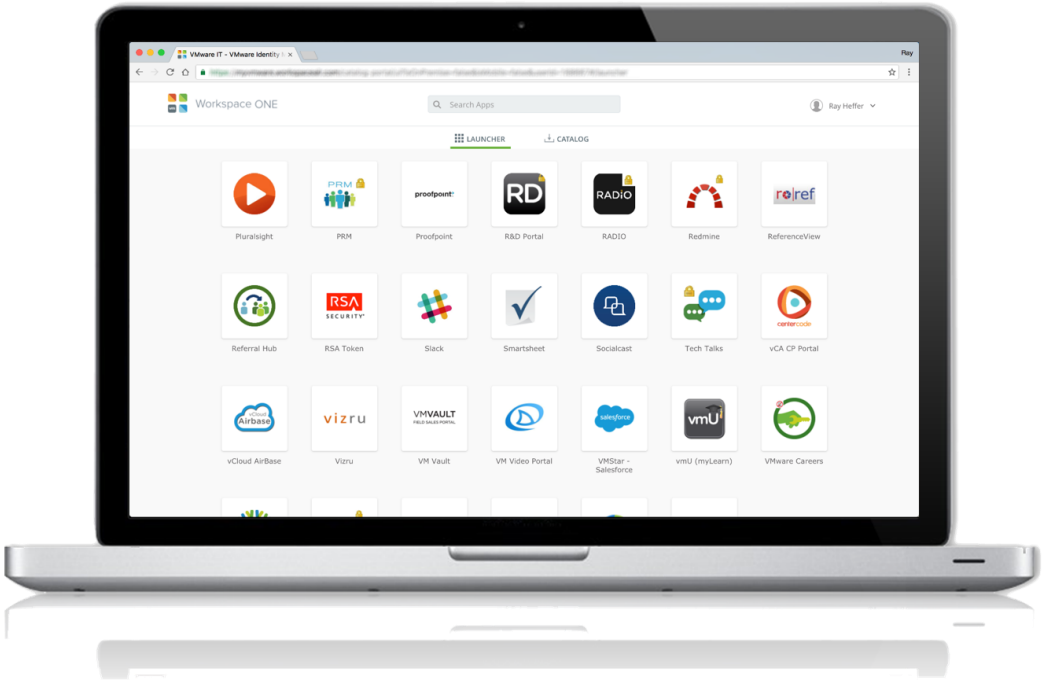

6.5.1 VMware Identity Manager

VMware Identity Manager provides the Workspace ONE portal as shown in the previous section, allowing users single-sign on (SSO) access to SaaS (Software-as-a-Service) applications, RDSH hosted Windows applications and desktop sessions, Horizon virtual desktops, ThinApp and Citrix applications. Identity Manager is available as a VMware hosted SaaS offering (VMware vCloud Air™) or as an on-premises deployment of which this document is based on.

6.5.1.1 Identity Manager Architecture

VMware Identity Manager is provisioned as a hardened Linux virtual appliance (running SUSE Linux Enterprise 11), that consists of several internal components including the connector and web services. In previous versions of Identity Manager (formerly known as Workspace Portal), the connector and web services were provisioned as separate virtual appliances. Identify Manager no longer has this requirement, therefore service providers can truly benefit from horizontal scaling.

Note Please refer to the Identity Manager 2.8 documentation for installation and sizing guidance.

The connector is the default identity provider service responsible for the synchronization of user data between Active Directory and the Identity Manager service. Third-party identity providers such as Google for Work and Microsoft Azure, can be used since they support the SAML 2.0 (Security Assertion Markup Language) protocol. This is otherwise known as SAML JIT (Just-in-Time) user provisioning.

When SAML JIT user provisioning is used with a third-party provider, it will create users in the Identity Manager service dynamically at logon.

In order to provide high availability of Identity Manager, it is recommended that a minimum of three virtual appliances are deployed with two or more Access Point virtual appliances. See the VMware documentation “Recommended Number of Nodes in VMware Identity Manager Cluster” in Section 10, References for more information.

In previous versions of Access Point (2.7 or below), in order to provide gateway services to both Horizon 7 and Identity Manager, two pairs of Access Point virtual appliances were required running version 2.5 (for Horizon support) and 2.7 (for Identity Manager support). This is no longer the case since Access Point 2.7.2, and a single pair of Access Point appliances can support both Horizon and Identity Manager.

As illustrated in the following diagram, a load-balancer is used (DMZ) for external connections, and an internal load-balancer sits in front of VMware Identity Manager appliances and View Connection Servers.

Figure 15. Identity Manager Architecture

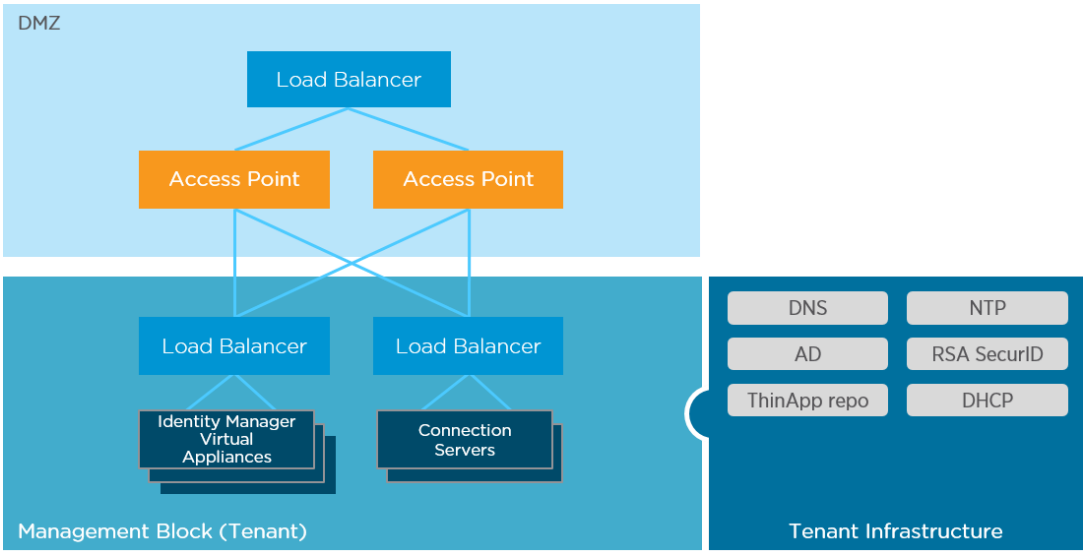

Note Identity Manager supports Transport Layer Security (TLS) 1.1 or 1.2. TLS 1.0 is disabled by default in VMware Identity Manager 2.6 or later. It can be enabled by following KB article 2144805: https://kb.vmware.com/kb/2144805

6.5.2 True SSO

True SSO when integrated with VMware Horizon and Identity Manager, allows users to authenticate with Horizon using biometric fingerprint, RADIUS or RSA SecurID, or using their AD credentials with the Identity Manager portal and then access desktops and applications without being prompted to enter their AD credentials again. Instead, single-sign on (SSO) authentication uses short-lived SSL certificates.

Figure 16. True SSO HA Deployment

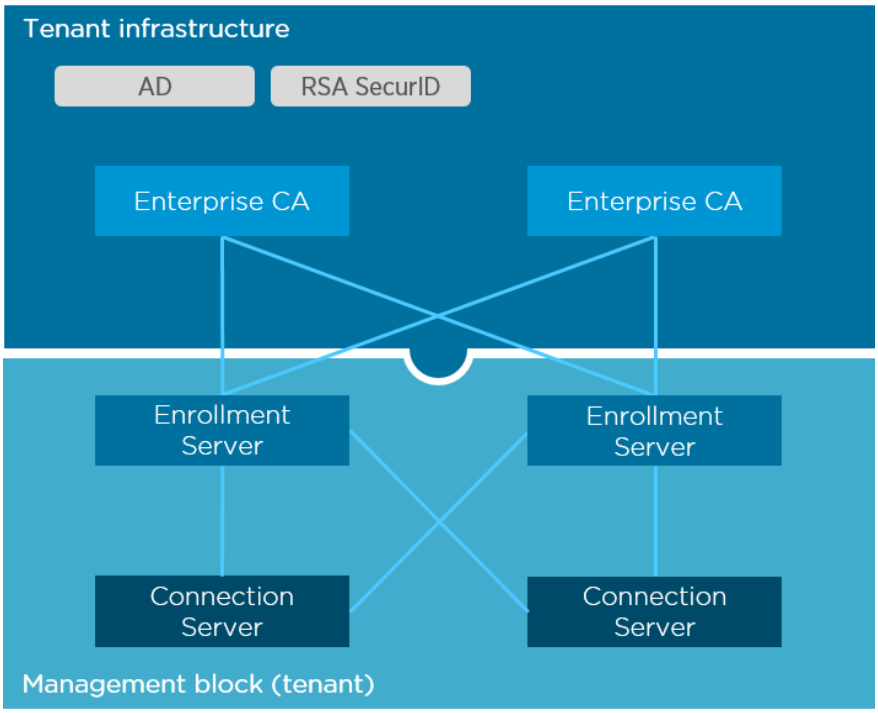

As illustrated in the diagram above, True SSO is not a single component of Horizon. To provide True SSO functionality, the tenant Management Block must contain the Enrollment Service, which consists of one or more Enterprise Certificate Authority servers (CA) and one or more Enrollment Servers. The enrollment service communicates directly with the Enterprise Certificate Authority to obtain the SSL certificate. The certificate is then used to logon to the domain. It is also possible to install the Enterprise CA on the same server as the Enrollment Service,

Note At least one Enterprise CA and Enrollment Server must exist in each tenant Active Directory domain. Enrollment Servers must be paired with a View Connection Server (up to two Enrollment Servers per Connection Server).

6.6 Management

One of the fundamental requirements for service providers is to establish roles and responsibilities for the management of each component in the digital workspace solution. This will provide a clear boundary for tenant and service provider administration, areas of responsibility and security.

Service providers will typically be responsible for:

- Physical server hardware and rack configuration
- Core networking, load-balancers and firewall security
- VMware product licensing
- Tenant infrastructure provisioning
- Backup and recovery
- Infrastructure monitoring
- vSAN configuration
- Maintaining Service Level Agreements (for instance, availability)
- Lifecycle management of vSphere and EUC infrastructure

Tenants will typically be responsible for:

- Application and OS licensing (Windows desktop, applications)
- Desktop image and pool management
- Desktop and application provisioning, entitlement and lifecycle management
- Active Directory (tenant domains) and user administration
- User environment monitoring (desktops, applications, web services)
- File servers
- Application management (AppStack provisioning and ThinApp packing)
- Creating and managing writeable volumes
- End User Support (Tier 1)

Multi-tenant environments adopt Role Based Access Control (RBAC) for each system, and must be clearly defined as service provider or tenant managed. In some scenarios, both the Tenant Administrator and service provider need access to the same components. For example, where the service provider is fully managing the VMware infrastructure access to both management and tenant vCenter Server instances might be required. Where this is a requirement, the tenant may have to provide Active Directory credentials with specific rights in order for the service provider to support the environment.

6.6.1.1 Tenant Operations

Routine operations and management of the solution will be performed by the tenant, using the following management interfaces:

- Horizon 7 Administrator
- Identity Manager Administration Console
- App Volumes Manager
- VMware vCenter Server (Resource Clusters)
- PowerCLI
- vRealize Operations for Horizon

6.7 Monitoring

Monitoring is a critical function for any service provider offering, not only for the data center infrastructure but also for the performance and utilization of user desktops and applications.

vRealize Operations for Horizon is the recommended tool to monitor end-to-end system performance of the Horizon deployment, including applications, user sessions, virtual machines, protocol utilization, and component availability.

Other tools used to monitor and troubleshoot the environment include:

- Horizon 7 Administrator dashboard
- vCenter Server
- VMware vRealize Log Insight™
- View Events database
- Syslog
- App Volumes dashboard

6.7.1 vRealize Operations for Horizon

vRealize Operations for Horizon provides Tenant Administrators complete monitoring of the Horizon 7 infrastructure, including:

- System health and analytics
- Performance monitoring
- Customizable dashboards and alerts
- User experience dashboard
- Out-of-the-box reports
- Report scheduling
- Protocol analysis (including: PCoIP, Blast, Citrix ICA)

Deployment of vRealize Operations for Horizon can be either as a single node or a cluster of nodes for availability and scalability. A cluster is the recommended deployment model. The first vRealize Operations node in the cluster is called the Master Node, and subsequent nodes are called Replica Nodes. Please refer to the documentation for sizing the VMware vRealize Operations Manager™ cluster (See Section 10, References).

Note For sizing guidelines also refer to KB article 2093783: https://kb.vmware.com/kb/2093783

Horizon 7 Sizing and Consumption Model

Architecting service provider implementations of Horizon 7 requires special attention to how each resource block is sized. As described previously, service providers can provide both shared and dedicated management blocks. Sizing for management components such as Connection Servers and Access Point appliances can be predicted based on the number of sessions (See Table 6). However, for enterprise customers, compute resource for desktops and applications, including the number of virtual desktops per host may vary per tenant.

Horizon DaaS implementations adopt pre-defined “desktop models”, giving the tenant a choice of desktop types (session-based RDSH or full desktop, with multiple compute configurations). This gives service providers the ability to offer desktops using a monthly consumption model (number of desktops / per month). Therefore, Horizon DaaS is particularly suited to tenants that don’t have the technical skills or desire to manage their own end-to-end Horizon deployment.

The enterprise desktop is often complex, containing very large numbers of applications, configurations, user customizations and application dependencies. This sizing model allows service providers to implement Horizon 7 in a highly scalable and secure fashion. No tenant will share physical desktop infrastructure with another. The physical infrastructure of management components is fully managed by the service provider. This allows tenant administration teams to focus on managing enterprise desktops and applications.

7.1 Configuration Maximums

The configuration maximums for Horizon 7 need to be observed, and not used for sizing targets. While some configuration maximums are hard limits, others are recommendations so that optimal levels of reliability and performance are achieved.

Table 5. Horizon 7 Configuration Maximums

Setting	Maximum	Details
Hosts per cluster	32	Current supported maximum for Horizon 7*
Virtual desktops/applications per pool	2,000	Recommended maximum
Sessions per resource vCenter Server (block)	10,000	Maximum supported number of virtual machines per vCenter Server
Connection Servers per Pod	7	Maximum number of active Connection Servers
Concurrent sessions per Connection Server	2,000	Maximum concurrent sessions per Connection Server
Concurrent Blast Extreme connections per Access Point	2,000	Maximum recommended concurrent Blast Extreme gateway connections
Concurrent PCoIP connections per Access Point	2,000	Maximum concurrent PCoIP connections via Access Point
Concurrent RDP connections per Access Point	2,000	Maximum concurrent RDP connections via Access Point
Concurrent connections per Pod	10,000	Maximum concurrent desktop or application connections in a single Pod
Concurrent connections per Cloud Pod	50,000	Maximum concurrent desktop or application connections in a federated Cloud Pod
Tenant Pods in a federated Cloud Pod	25	Maximum number of “Pods” federated in single Cloud Pod
Maximum RDS farms per Horizon Pod	200	Current supported maximum for Horizon 7*
Maximum RDS hosts per farm	200	Current supported maximum for Horizon 7*
Concurrent connections to a single App Volumes Manager instance	1,000	Recommended scale limit to a single App Volumes Manager when using multi vCenter Server mode
* Current as of December 2016		

7.2 Assessment

Sizing Horizon 7 for an enterprise deployment requires a full assessment of the desktop estate, uses-cases including desktop configuration, and application workloads. This is not practical for many service providers, therefore resource blocks are provided to each tenant, and the role of Tenant Desktop Administrators will consume the capacity allocated to them for desktops and applications.

It is essential that service providers can scale tenant resources blocks on-demand, which is why SDDC components such as vSAN and NSX provide an excellent platform for the hosted digital workspace.

Note This does not mean that an assessment is not required. Architectural guidelines including a desktop assessment strategy still applies for Tenant Administrators. This may optionally be offered as a service by the provider.

To facilitate the hosted model, service providers will provide resource infrastructure blueprints as described in Section 7.3, Infrastructure Blueprints.

7.3 Infrastructure Blueprints

Each tenant blueprint defines the sizing constants and specification of hosts in the resource cluster, primarily compute and storage. These values will be used to calculate the resource capacity for each tenant, and will reflect the service levels for a given tenant.

In the following examples, a single blueprint will be used (Blueprint 1), and will be based on the Silver SLA as outlined in Section 9, Service Levels.

Table 6. Blueprint 1 - Sizing Formula Constants

Setting	Value
Virtual desktops/applications per pool	1,000
Sessions per resource vCenter Server (block)	2,000
Maximum users per App Volumes Manager Server	1,000
Maximum users per Identity Manager appliance	30,000
Maximum users per vRealize Operations server	10,000
Minimum / Maximum Access Point appliances per tenant	2-7

Table 7. Blueprint 1 – vSAN Ready-Node Host Specification

Specification	Value
Manufacturer	Dell PowerEdge R730xd, Intel Xeon E5-2667 v3 (vSAN Ready
Host memory (GB)	256
CPU Sockets	2
CPU Cores (per socket)	10
CPU Specification	3.2Ghz (Intel Xeon E5-2667 v3, 20MB cache)
SSD Capacity / Quantity	2 x 400GB
SAS Capacity / Quantity	6 x 1200GB
GPU	None
Minimum number of hosts (single Resource Block)	4
Maximum number of hosts (single Resource Block)	32

Table 8. Blueprint 1 - Performance Specifications (Silver SLA)

Specification	Value	Description
Max host memory utilization	80%	vSphere resource cluster hosts will be sized, not exceeding this value.
Max host CPU utilization	80%	vSphere resource cluster hosts will be sized, not exceeding this value.
Failover hosts	2	Reserve host capacity for 2 hosts to provide availability during a host failure and host maintenance.
Virtual machine memory reservation	50%	Virtual machine memory will be reserved at 50%. See recommendations in Section 7.3.2, vSAN.
vCPU to Physical Core Ratio	10:1	Virtual machine vCPU to physical CPU core ratio (excludes RDSH virtual machines)
RDSH vCPU to Physical Core Ratio	1:1	RDSH vCPU to physical CPU core ratio
RDSH Session RAM (MB)	512	Estimated average Peak RDSH per-session memory utilization
Maximum sessions per RDS Host	24	Maximum sessions per RDS Host

Note Additional host blueprints can be specified for features such as vGPU, vSAN All-Flash vs Hybrid and varying CPU / Memory specifications, in addition to different service levels (Bronze, Silver, Gold, Platinum).

The maximum host memory and CPU utilization stated here is an example. Other environmental factors such as overheads running NSX and vSAN will need to be considered.

7.3.1 vGPU

One of the advantages of Horizon 7 is the ability to offer GPU (Graphical Processing Unit) accelerated workloads for 3D graphical capabilities or machine learning applications which is gaining in popularity.

Infrastructure blueprints may include NVIDIA GPU-enabled clusters using either the GRID 1.0 (K1 or K2) or GRID 2.0 (M60) GPU cards. These can be utilized by tenants with virtual desktops, or using RDSH sessions.

Table 9. NVIDIA GRID vGPU Profiles

Card	Physical GPUs per board	GPU Profiles	Frame Buffer	Max Displays per User	Max vGPUs per pGPU	Max vGPUs per board
NVIDIA GRID K2	2	K280Q	4GB	4	1	2
		K260Q	2GB	4	2	4
		K240Q	1GB	2	4	8
		K220Q	512MB	2	8	16
NVIDIA GRID K1	1	K180Q	4GB	4	1	4
		K160Q	2GB	4	2	8
		K140Q	1GB	2	4	16
		K120Q	512MB	2	8	32
NVIDIA TESLA M60 GRID Virtual Workstation	2	M60-8Q	8GB	4	1	2
		M60-4Q	4GB	4	2	4
		M60-2Q	2GB	4	4	8
		M60-1Q	1GB	2	8	16
		M60-0Q	512MB	2	16	32
NVIDIA TESLA M60 Virtual PC	2	M60-1B	1GB	4	8	16
		M60-0B	512MB	2	16	32
NVIDIA TESLA M60 Virtual Applications	2	M60-8A	8GB	1	1	2
		M60-4A	4GB	1	2	4
		M60-2A	2GB	1	4	8
		M60-1A	1GB	1	8	16
NVIDIA TESLA M10 GRID Virtual Workstation	4	M10-8Q	8GB	4	1	4
		M10-4Q	4GB	4	2	8
		M10-2Q	2GB	4	4	16
		M10-1Q	1GB	2	8	32
		M10-0Q	512MB	2	16	64

Card	Physical GPUs per board	GPU Profiles	Frame Buffer	Max Displays per User	Max vGPUs per pGPU	Max vGPUs per board
NVIDIA TESLA M10 Virtual PC	4	M10-1B	1GB	4	8	32
		M10-0B	512MB	2	16	64
NVIDIA TESLA M10 Virtual Applications	4	M10-8A	8GB	1	1	4
		M10-4A	4GB	1	2	8
		M10-2A	2GB	1	4	16
		M10-1A	1GB	1	8	32

7.3.2 VMware vSAN

vSAN provides scale-out storage which makes it the ideal storage platform for service providers. It allows tenants to start small, and grow as needed by adding additional nodes to vSAN-enabled clusters. Using vSAN as an integral part of the tenant resource cluster blueprint, it provides that scaling up the number of tenants, and scaling out desktops for single tenant, follows a predictable and repeatable design approach.

The following table lists the current supported maximums for vSAN with vSphere 6.5 (current as of the date of this document).

Table 10. vSAN Configuration Maximums

Setting	Maximum
Virtual machines per host	200
Virtual machines per cluster	6,400
Hosts per cluster	64* (32 for Horizon deployments)
Components per host	9,000
Disk groups per host	5
Capacity disks per group	7
SSD disks per group	1
* A maximum of 32 hosts will be used with Horizon 7	

7.3.2.1 Failure Tolerance Method (FTM)

vSAN can be configured in one of two fault tolerance modes. RAID1 (Mirroring) delivers the best performance and RAID 5/6 (Erasure Coding) which is optimized for capacity. Erasure Coding does not utilize witness components, instead it uses data parity. There is an overhead to using Erasure Coding, since it has to perform the parity bit calculations. More information on Erasure Coding can be found in the vSAN 6.2 Design and Sizing Guide (see Section 10, References).

Note An all-flash vSAN configuration is required for RAID5/6 erasure coding, deduplication and compression.

7.3.2.2 Objects and Components

A vSAN object represents a logical block device, typically associated with SCSI block storage. For example, a virtual disk (VMDK), VM swap, and the VM home namespace are all objects. Objects are distributed across the vSAN cluster using components. There are replica components which are full copies of the data, such as the VMDK disk, and witness components that contain only metadata for the object.

RAID1 (Mirroring) will be used in any examples that follow since it offers the best performance. When this mode is used, there is one witness component per object, and a number of replica components that will depend on the FTT (Failures to Tolerate) configuration.

The maximum size of a single component is 255GB, therefore if an object is larger than 255GB, it is split into multiple components. The table below lists the objects for a virtual machine:

Table 11. vSAN Objects

Object	Description
Virtual machine home namespace	VM Configuration (.vmx), log files, and digest files.
Virtual machine swap	Virtual machine swap is equal to the amount of configured VM memory, minus any reservations. Only applies to powered-on virtual machines.
Virtual machine disk (VMDK)	OS VMDK, OS Checkpoint, user data disk, disposable disk, Linked Clone internal disk, App Volumes mounted AppStack, App Volumes Writeable Volume.
Snapshot deltas	Virtual machine snapshot delta files.
Snapshot memory (.vmem)	Virtual machine memory snapshot used to retain the live state of a virtual machine.
Redo-log file	Redo-log files are used with independent non-persistent disks, and are present in the virtual machine name space. App Volumes AppStacks are independent non-persistent (read-only) VMDK objects, with each AppStack mount (concurrent user) requiring a redo-log file object.

By default, virtual swap is provisioned with 100% Object Space Reservation, however, as of VSAN 6.2 virtual swap can be thin provisioned using the **SwapThickPrivisionDisabled** setting.

Note By default, Horizon uses a replica disk policy that has a read cache reservation of 10%. If an all-flash configuration is used, please refer to KB 2109890 (https://kb.vmware.com/kb/2109890). All-flash configurations do not currently support the Flash Read Cache Reservation storage policy attribute, and is not supported.

7.3.2.3 Example Virtual Machine Specification

Using the example in the table below, the number of vSAN objects and components are calculated.

Table 12. Virtual Machine Specification

Specification	Value
Guest OS	Windows 10
Clone type	Linked Clone (Composer)
Configured Memory	4GB
vCPU	2
System disk	40 GB
User data disk	None
App Volumes AppStacks	3
App Volumes Writeable Volume	1*
Snapshots	None
Number of Objects	8

* Only 1 writeable volume per virtual machine is possible.

Using the example in this table, there are the following objects for the virtual machine:

- Virtual machine namespace
- Virtual swap
- OS Disk (VMDK)
- Internal disk (used by Linked Clones)
- AppStack 1 (Redo-log)
- AppStack 2 (Redo-log)
- AppStack 3 (Redo-log)
- Writeable Volume

Each App Volumes AppStack is read-only and accessible to multiple virtual machines, which is achieved by mounting them as independent non-persistent disks. This means that a single AppStack (VMDK) can be mounted on many hundreds or thousands of virtual machines, and any changes that occur while the AppStack is mounted are written to a delay called a Redo-log. While the object count above includes the AppStack, it is the Redo-log file object that will exist in the namespace for each virtual machine.

Consider the following virtual machine namespace file structure:

Figure 17. vSAN Datastore VM Namespace Example

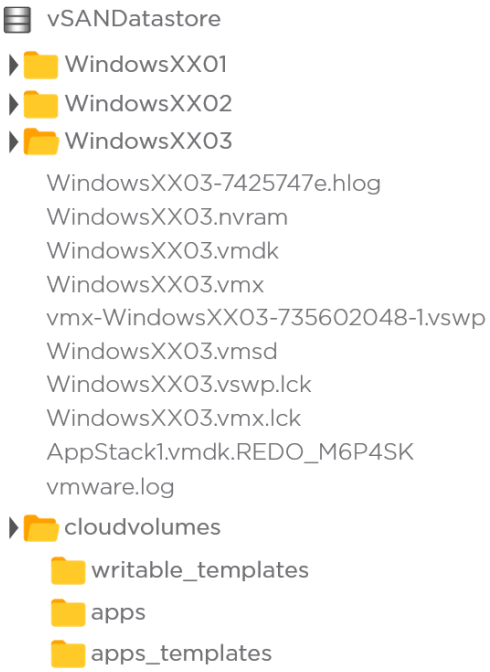

In this example, WindowsXX03 has these objects:

- WindowsXX03 VM namespace
- WindowsXX03 VMDK
- WindowsXX03 virtual swap
- AppStack1 Redo-log

The AppStack objects are only counted once as they reside in the `cloudvolumes` namespace.

7.3.2.4 Failures to Tolerate (FTT)

The number of failures to tolerate (FTT), will determine how many copies of the object will be required. For example: FTT=1 is 2 copies, FTT=2 is 3 copies, and FTT=3 is 4 copies. This count includes the original object itself but not the witness.

In the example used in Table 12, there are 8 objects for the virtual machine, and the following formula is used to calculate the number of components distributed across the vSAN cluster.

Formula: Objects * (copies + witness) = components

Table 13. FTT Formula Result

FTT Setting	Component copies	Formula	Total number of components required for 100 virtual machines
1	2x	8 * (2 + 1) = 24	2,400
2	3x	8 * (3 + 1) = 32	3,200
3	4x	8 * (4 + 1) = 40	4,000

With 100 virtual machines per host and an FTT of 1, there will be a total of 2,400 components distributed across the cluster. Even for a higher service level that would dictate FTT of 2, you can see this is well within the 9,000 per-host component limit.

In the following screenshot, the number of failures to tolerate (FTT) is set to 1, and the Horizon replica disk (VMDK) is selected. In this example, there are 2 replica components (RAID1) and a witness.

Figure 18. Witness and Replica Components

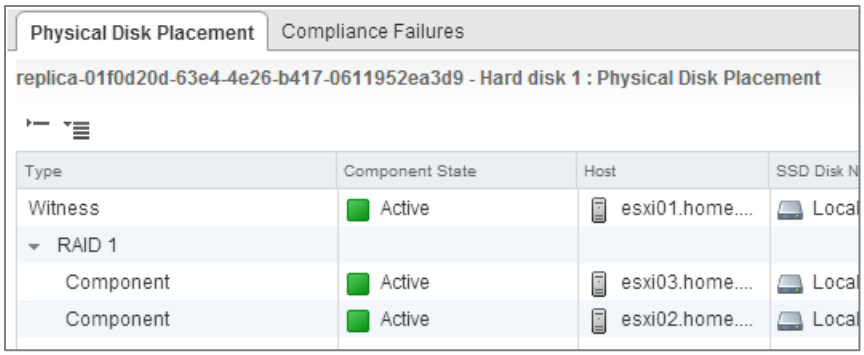

7.4 Sizing Conclusion

This tenant sizing scenario is based on a small deployment of 100 desktop sessions with an expected user session concurrency of 60%. This is a single use-case which is common for a Proof-of-Concept (PoC) with a tenant. Using the sizing and blueprint information in the previous sections, we can now calculate a single tenant deployment based on a shared management block.

To aid with sizing of any Horizon, it is recommended that the VMware Horizon Sizing Estimator is used. The Horizon Sizing Estimator will guide the user through each step of an enterprise deployment with Horizon 7, along with architectural guidelines and recommendations. This is currently available to VMware partners and employees (see the link in Section 10, References).

Table 14. Tenant Use-Case 1 - Engineering PoC

Question	Answer	Design Outcome
Non-core applications	Yes	Applications that are not part of the base-image. App Volumes required.
User-installed applications?	Yes	App Volumes Writable Volumes required or Full Clone Desktops.
Non-core application disk space required (GB)	2GB	Capacity required for App Volumes AppStack.
User-installed application disk space required (GB)	4GB	Capacity required for App Volumes writable volume.
Access to local client drive?	No	Client Drive Redirection not required.
Access to locally attached USB drive?	No	No USB redirection required.
Access to locally attached printer?	Yes	Virtual printing requirement.
User Concurrency	60% 60 users.	Expected concurrent average user connections.

Table 15. Use-Case 1 - Recovery Specifications

Specification	Value	Description
Secondary Pod (Cloud Pod Architecture)	No	Secondary Pod can be used for disaster recovery. Not required for PoC.

Note Security and Management specifications not specified for this PoC use-case.

In this example (Tenant use-case 1), 4 hosts will be used in the resource block as illustrated in the following figure. This will allow for sufficient capacity during a host failure in addition to a host being taken offline for maintenance.

Figure 19. Tenant Resource Block Example

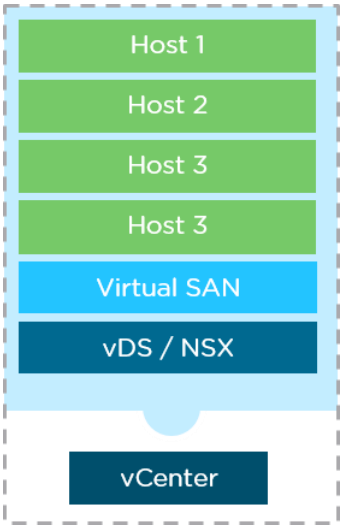

Total Host Logical Cores: **20**
Total Host CPU Clock Speed (GHz): **64Ghz**
Total Host RAM (GB): **256GB**

Available Hosts: **3 (+ 1 failover host)**
Available Host CPU (Ghz) at 80% utilization: **51Ghz**
Available Host RAM (GB) at 90% utilization: **205GB**

Available Cluster CPU Clock Speed (Ghz): 153Ghz

A shared management block has been selected, with the following tenant management components:

Table 16: Tenant Management Components

Component	Qty.	Description
vCenter Server Appliance	2	A pair of vCenter Server appliances (vCenter HA) will be used for the tenant resource block. Tenant administrators will have access for the management of desktop images and vSphere operations.
Access Point	2	A pair of Access Point virtual appliances will be deployed behind an NSX load-balancer. This will be used for resilience and facilitate external (internet) access.
Connection Server	2	Two Horizon View Connection Servers will be deployed for resilience.
NSX Manager	2	Two NSX Manager virtual appliances are deployed for management of the resource cluster networking.
vRealize Operations for Horizon	2	A single vRealize Operations Manager server will be deployed with the Horizon adapter. This will be used for tenant monitoring of the environment.
Identity Manager	3	Three VMware Identity Manager appliances will be deployed for access to Workspace ONE.
Enrollment Servers	2	Two Enrollment Servers will be deployed for True SSO.
User Environment Manager	1*	A single file server will be deployed for User Environment Manager.

* Where a single User Environment Manager file server is deployed, in this example it is not considered critical to operations. For high availability of the User Environment Manager file shares, Microsoft DFS-R (replication) can be used in a hub and spoke replication topology (active-active replication is not supported).

Networking

There are several components that make up the “digital workspace” infrastructure, which includes; Horizon 7, Identity Manager, App Volumes, Access Point and vRealize Operations for Horizon. In the following figure, each of the infrastructure components and the corresponding network ports, direction flow and protocols are listed.

Figure 20. Horizon 7 Network Ports

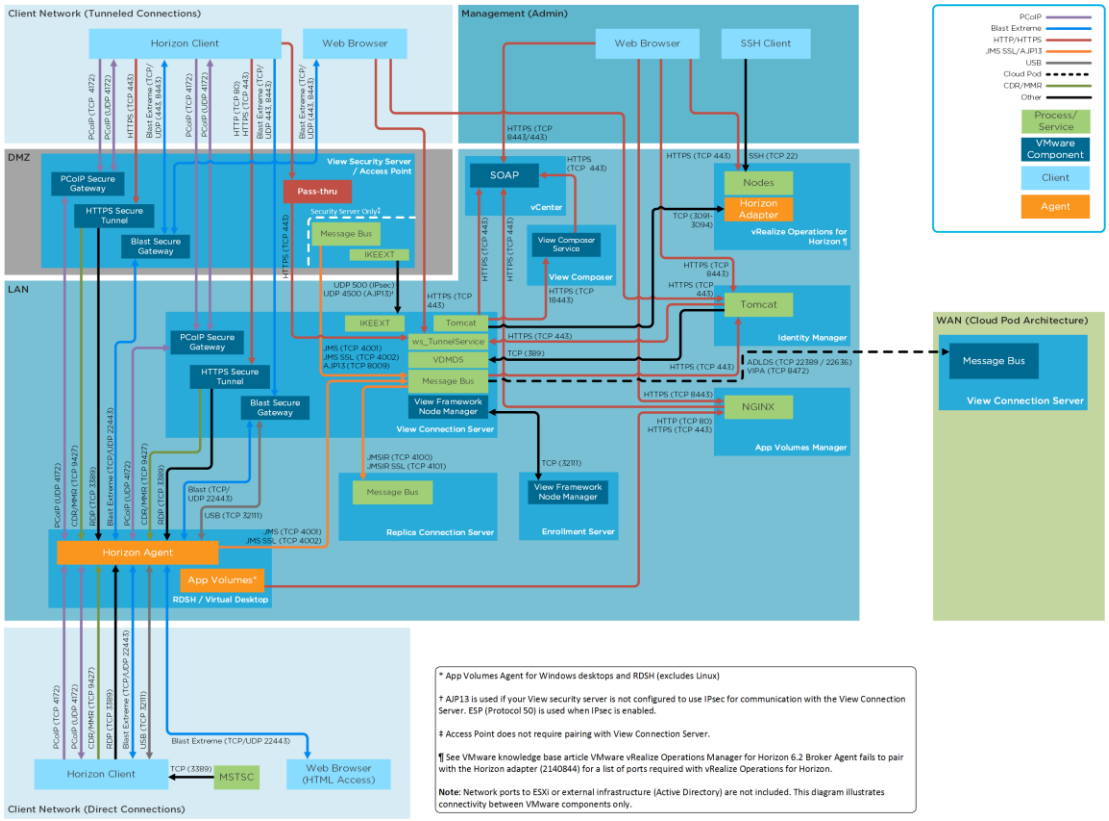

Note See the original (full size) diagram here: http://www.vmware.com/techpapers/2015/vmware-horizon-7-network-ports-diagram-10492.html.

8.1 VMware NSX

Most network data traffic in the tenant desktop environment is “east-west”, meaning that data flows between desktops and server components in the datacenter. For some tenants this is important to understand as it directly impacts compliance and security risk mitigation. Some of the risks associated with east-west traffic include:

- Rogue user behavior
- Zero-day threats
- Compromised websites
- Malware and malicious code (desktop to desktop and desktop to server)

Traditional networking can make this challenging, since it adds complexity with managing multiple VLANs, firewall rule-sets and access control lists.

VMware NSX targets east-west traffic, in addition to north-south, and Tenant Administrators can easily create policies (see the following figure) that dynamically follow desktops.

Figure 21. NSX Firewall Policy Example Rule-Set

No.	Name	Rule ID	Source	Destination	Service	Action	Applied To	Log
▸ SecurityServer - Edge Section (Rule 1 - 2)								
▾ VDI Service - Internal (Rule 3 - 14)								
3	Desktop to Desktop FW ...	1041	VDI Desktops 001	VDI Desktops 001	any	Block	Distributed Firewall	Log
4	View Agent FW Rules	1040	VDI Desktops 001	View Connection Servers...	Horizon View 6x Connect...	Allow	Distributed Firewall	Log
5	InterCS FW Rule	1039	View Connection Servers...	View Connection Servers...	HorizonView6_interCS	Allow	Distributed Firewall	Log
6	Composer FW Rules	1038	View Connection Servers...	Composer Server 001	Horizon View 6x Compos...	Allow	Distributed Firewall	Log
7	Connection Server FW R...	1037	View Clients 001	View Connection Servers...	Horizon View 6x Connect...	Allow	Distributed Firewall	Log
8	Security to Connection S...	1048	Security Servers	View Connection Servers...	Horizon View 6x Security...	Allow	Distributed Firewall	Log
9	SS to Agents	1049	Security Servers	VDI Desktops 001	Horizon View 6x SS to A...	Allow	Distributed Firewall	Log
10	Desktop FW Rules	1036	View Clients 001 View Connection Servers...	VDI Desktops 001	Horizon View 6x Desktops	Allow	Distributed Firewall	Log
11	Finance Filer Access	1044	Finance Desktops	FIL01	Server Message Block (... SMB Server SMB Server UDP	Allow	Distributed Firewall	Log
12	Legal Project Mgmt Server	1046	Legal Desktops	PMP01	HTTP	Allow	Distributed Firewall	Log
13	Email Service	1045	Finance Desktops IT Desktop Legal Desktops	EML01	POP3	Allow	Distributed Firewall	Log
14	Domain Services	1047	Finance Desktops IT Desktop Legal Desktops	Domain Controllers	Active Directory Server Active Directory Server U... DHCP-Client DHCP-Server DNS	Allow	Distributed Firewall	Log

NSX abstracts the network infrastructure into a logical representation that is made up of a library of network services, that can be interconnected to create a virtual network overlay. This enables the tenant to place each virtual desktop in its own individual network container. Using the NSX DFW (Distributed Firewall) it is possible to deliver micro-segmentation, eliminating unauthorized cross-talk between virtual desktop or server workloads.

Tenant administrators can define policy rule-sets that automatically attach to a virtual desktop, or virtual machine groups so the policy is enforced at the time the virtual machine is created.

NSX also exposes a RESTful API, allowing cloud management platforms to automate the delivery of network services.

Service Levels

Service levels and infrastructure blueprints will determine five key service qualities for the solution; availability, performance, security, management, and recoverability.

The following example is a very basic (Gold, Silver, Bronze) service differentiator. Adapt the example to suit the service level agreement (SLA) offerings by a service provider.

Gold	●●●●●	●●●●●	●●●●●	●●●●●	●●●●●
Silver	●●●○○	●●●●○	●●●●○	●●●○○	●●●○○
Bronze	●●●○○	●●●○○	●●●○○	●●○○○	○○○○○

Many cloud service providers offer uptime guarantees as part of their SLA. While the usage of uptime percentages for cloud providers are common, such as 99.99% (approximately 52 minutes a year downtime), it is the ability to recover from an unexpected failure within a maximum RTO (Recovery Time Objective) and RPO (Recovery Point Objective) priority value that matters.

In order to meet very strict RTO priorities of < 1 hour, there must be no single points of failure in the physical infrastructure or software components. Cloud Pod Architecture (See page 23) takes advantage of two or more Pod implementations that can be geographically separated to solve disaster recovery challenges. This is reflected in the Gold service level described below.

Information security is often overlooked when SLAs are defined, however, many desktop environments are subject to strict security governance and compliance no matter how small or large the environment. Offering the digital workspace to enterprise tenants needs to be secure, both in connectivity and data storage requirements. Components such as NSX, SSL Certificates, and 2-Factor Authentication are mentioned here. However, other security areas that service providers may be subject to as part of the SLA include: Penetration and vulnerability testing, auditing, configuration and change management, forensics, and information impact levels.

9.1 SLA Example

9.1.1 Gold

9.1.1.1 Availability

This will provide the highest level of availability for the tenant. A dedicated management cluster is configured with High Availability (HA) in an N+2 configuration and all resource clusters are N+2. vSAN will be configured with a Failures to Tolerate (FTT) of 2.

9.1.1.2 Performance

Virtual machine memory reservation in management and resource clusters will be configured at 100% to guarantee memory is available, and eliminate virtual machine swap files. The ratio of virtual desktops per CPU core on each host will be sized to no more than 5:1. RDS hosts will be configured with a 1:1 ratio for vCPU to physical CPUs on each host, and session workloads sized with 1024MB average per user session.

9.1.1.3 Security

The service provider will manage and deploy trusted SSL certificates on all infrastructure components (vCenter Server, Horizon, vSphere, Identity Manager). The service provider will fully manage firewalls including NSX micro-segmentation. 2-Factor authentication will be used for all access to Identity Manager.

9.1.1.4 Management

The tenant will be provided with a dedicated management block providing tenant access (optional) to vCenter Server and all management components. vRealize Operations for Horizon will be implemented for monitoring the solution. vRealize Log Insight will be implemented and used as a syslog server to collect all logging information for ESXi hosts and Horizon server components.

9.1.1.5 Recoverability

Cloud Pod Architecture will be used to maintain Pods across two data centers (sites). User data and applications (ThinApp and App Volumes) are replicated between data centers using storage replication.

9.1.2 Silver

9.1.2.1 Availability

The management cluster is configured with High Availability (HA) in an N+1 configuration and all resource clusters are N+1. vSAN will be configured with a Failures to Tolerate (FTT) of 1.

9.1.2.2 Performance

Virtual machine memory reservation in management and resource clusters will be configured at 50%. The ratio of virtual desktops per CPU core on each host will be sized to no more than 10:1. Maximum host memory utilization will be 90%, and maximum CPU utilization will be 80%. RDS hosts will be configured with a 1:1 ratio for vCPU to physical CPUs on each host, and session workloads sized with 512MB average per user session.

9.1.2.3 Security

The service provider will manage and deploy trusted SSL certificates on all infrastructure components (vCenter Server, Horizon, vSphere, Identity Manager). NSX will provide micro-segmentation, managed by the tenant. 2-Factor authentication will be used for user access to Horizon.

9.1.2.4 Management

Tenant management resources will be hosted on a shared management block. vRealize Operations for Horizon will be implemented for monitoring the solution. vRealize Log Insight will be implemented and used as a syslog server to collect all logging information for ESXi hosts and Horizon server components.

9.1.2.5 Recoverability

The service provider will manage backups of all management and resource virtual machines.

9.1.3 Bronze

9.1.3.1 Availability

The management cluster is configured with High Availability (HA) in an N+1 configuration and all resource clusters are N+1. vSAN will be configured with a Failures to Tolerate (FTT) of 1.

9.1.3.2 Performance

The ratio of virtual desktops per CPU core on each host will be sized to no more than 15:1. RDS hosts will be configured with a 1:1 ratio for vCPU to physical CPUs on each host, and session workloads sized with 340MB average per user session.

9.1.3.3 Security

Trusted SSL certificates must be deployed and managed by the tenant.

9.1.3.4 Management

Basic environment monitoring will be provided by the service provider for outages or component failures.

9.1.3.5 Recoverability

No recovery of the solution. The tenant will maintain their own backups and copies of desktop images and applications.

Technical Reviewers

Simon Greaves is a Senior Consultant for VMware Professional Services Organization in EMEA. A virtualization and cloud specialist with over 12 years of experience working in highly technical environments for large scale enterprises with a focus on the VMware SDDC (Software Defined Datacenter) and Cloud Service Providers with VMware vCloud Director®.

Simon Long (Double VCDX #105) is an EUC Architect working for the VMware OneCloud team as a specialist in End User Computing and Cloud Architecture. He has been working for VMware for over six years in which he has helped some of the largest VMware European customers design and deploy VMware Horizon into their businesses. Prior to VMware, Simon had worked in a wide range of IT environments ranging from Cloud providers, education institutes to financial shops in a career spanning 16 years.

Travis Wood (Double VCDX #97) is a Principal Architect in the Global Services Engineering team. Travis develops the design methodology that VMware Professional Services use to design and implement EUC solutions around the world and develops the solution architecture for large or complex enterprises. In addition, he is a regular speaker at VMware events such as VMworld and vForum.

References

Additional information pertinent to this document and its topics:

Table 17. References

Document Title	Link or URL
vCloud Architecture Toolkit (vCAT) Blog	https://blogs.vmware.com/vcat
vSAN Ready Node Configurator	http://vsanreadynode.vmware.com/RN/RN
vSAN 6.2 Design and Sizing Guide	http://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/products/vsan/virtual-san-6.2-design-and-sizing-guide.pdf
VMware Horizon Sizing Estimator	http://vmware.com/go/horizoncalc
Horizon 7 Network Ports	http://www.vmware.com/techpapers/2015/vmware-horizon-7-network-ports-diagram-10492.html
VMware Horizon RDSH maximums	https://pubs.vmware.com/horizon-7-view/index.jsp?topic=%2Fcom.vmware.horizon-view.desktops.doc%2FGUID-C3F2BAB8-BA06-440D-BF2B-4624E71E9AED.html
View Connection Server Maximums and Virtual Machine Configuration	https://pubs.vmware.com/horizon-7-view/index.jsp#com.vmware.horizon-view.planning.doc/GUID-1AC83D7C-BDC6-4E32-A35C-652107BDB6D0.html
Services on a View Connection Server Host	https://pubs.vmware.com/horizon-7-view/index.jsp?topic=%2Fcom.vmware.horizon-view.security.doc%2FGUID-B226ED1C-2D38-462D-BABD-1D9A1C00EC3D.html
NVIDIA Machine Learning	http://www.nvidia.com/object/machine-learning.html
NVIDIA GRID Technology	http://www.nvidia.com/object/grid-technology.html
NVIDIA GRID User Guide	http://images.nvidia.com/content/pdf/grid/guides/GRID-vGPU-User-Guide.pdf
VMware Horizon 7 Instant Clone Desktops (white paper)	https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/techpaper/vmware-horizon-view-instant-clone-technology.pdf
Access Point Deployment Documentation	https://www.vmware.com/support/pubs/access-point-pubs.html

Document Title	Link or URL
Access Point Deployment with PowerShell	https://communities.vmware.com/docs/DOC-30835
Recommended Number of Nodes in VMware Identity Manager Cluster	http://pubs.vmware.com/identity-manager-27/index.jsp?topic=%2Fcom.vmware.wsp-install_27%2FGUID-3BFB1D4D-D5C2-480D-94E0-31ED6B0CAA63.html
VMware Identity Manager 2.8 Documentation	https://www.vmware.com/support/pubs/identitymanager-pubs.html
Sizing the vRealize Operations Manager Cluster	https://pubs.vmware.com/vrealizeoperationsmanager-63/index.jsp#com.vmware.vcom.core.doc/GUID-6108A393-2950-47E0-AC74-1C4FC532B292.html
Department of Defense Cloud Computing Security Requirements Guide	http://iasecontent.disa.mil/cloud/Downloads/Cloud_Computing_SRG_v1r2.pdf
Load Balancing with VMware Access Point	https://communities.vmware.com/docs/DOC-32792

10.1 Software Versions

The following software versions were referenced in this document:

Table 18. Software Versions

Product	Version
Access Point	2.7.2
VMware NSX	6.2.4
VMware App Volumes	2.11
VMware Horizon	7.0.2
VMware Identity Manager	2.8
VMware User Environment Manager	9.1.0
VMware vCenter Server	6.5
VMware vRealize Log Insight for vCenter™ and NSX	4.0.0
VMware vRealize Operations for Horizon	6.3.0
VMware vSphere (including vSAN)	6.5