TECHNICAL WHITE PAPER – **DECEMBER 2017**

VMWARE ALWAYSON DIGITAL WORKSPACE DESIGN GUIDE

Version 3.1

Table of Contents

About Design Guides

VMware design guides are created through architectural design development and review by subject matter experts. The guides provide overviews of solution architectures and high-level implementation guidance. As a reference asset, each document illustrates a design framework to support proof-of-concept, pilot, and full implementation. However, for each implementation, customer-specific design and associated documentation must be developed.

Design guides ensure the viability of logical designs or concepts in real-world real-world practice. This document does not replace or supersede product specifications and installation guidelines published for each VMware product. All detailed technical and functional product-level questions should be referred to the appropriate product documentation as well as assets developed by the services organization.

Introduction

This guide provides an overview of the VMware AlwaysOn Digital Workspace solution, its logical architecture, and validation of the capabilities. Based on products from VMware, this architecture represents the foundation on which customers and partners can build comprehensive solutions that require high availability for end-user-computing (EUC) services.

The term *workspace* refers to a software environment within which an end user accesses and interacts with one or more entitled applications. A workspace is comprised of a traditional Windows desktop, browser applications, and published applications.

The term *AlwaysOn* refers to a level of service availability where, in the event of a service disruption, the designed recovery time objective (RTO) is measured to be 15 minutes or less.

The solution described here is not exclusive to the third-party products referenced within the guide. Please consult with your VMware representative on how to implement this architecture using your preferred third-party vendors and supported products. This document will be updated as newer capabilities are incorporated in the AlwaysOn Digital Workspace solution.

Audience

This document is for enterprise architects, solution architects, sales engineers, field consultants, advanced services specialists, and customers who plan to design, configure, and deploy an AlwaysOn Digital Workspace solution.

Revisions and Additions

This version of the AlwaysOn Digital Workspace Design Guide includes the following additions and revisions as compared with version 3.0.

1. Revised terminology definition for AlwaysOn to reflect stateful components
2. Functionality based on the VMware Horizon® 7 version 7.3 platform (including the RDSH application delivery platform)
3. Incorporated *VMware Horizon 7 Enterprise Edition Multi-Site Reference Architecture* white paper
4. Incorporated *VMware Identity Manager Reference Architecture* white paper

Business Case

Organizations across all industries recognize that the Windows-based desktop service model based on physical PCs no longer meets their end users' increasing demands for mobility, device choice, security, and agility in application delivery and lifecycle management. Moreover, as end-user productivity grows due to the increasing adoption of technology, the net cost of downtime in EUC services continues to increase as a result. This calls for a solution that abstracts planned and unplanned desktop service disruptions (for example, patch updates, break/fix, OS upgrades, and more) from users' business productivity. Total Cost of Ownership (TCO) analyses routinely show that over 50 percent of current desktop TCO is due to end-user loss of productivity.

While virtual desktop technologies offer a highly desirable alternative for next-generation EUC services, the infrastructure delivering those services must be highly resilient. This is a necessary architectural requirement in order to counter-balance the inherent increase in fault domain. Resiliency requires a foundational architecture designed to support desired uptime and capacity SLA requirements.

What Is the AlwaysOn Digital Workspace?

VMware AlwaysOn Digital Workspace is a complete, end-to-end solution for a private (on-premises) cloud infrastructure for virtual applications and desktops based on the VMware Horizon platform. The solution offers critical capabilities in three areas:

- Availability
- Workspace mobility
- Security

Availability

End-to-end redundancy is a primary aspect of the AlwaysOn Digital Workspace solution design. By providing high resilience through redundancy, the solution eliminates potential single points of failure in delivering the digital workspace experience to end users.

By combining end-to-end redundancy and an active-active architecture, the solution provides multiple available paths to access virtual desktop cloud(s) running in the customer's data center(s). If a path becomes unavailable due to component outage or planned maintenance, the system intelligently routes around the unavailable component or path and maintains delivery of desktop services to end users.

Workspace Mobility

Workspace mobility is a capability that allows end users access to their digital workspaces from any connected device (on corporate network or Internet) without having to log out and log in as they move between devices. As end users move from device to device and across locations, the AlwaysOn Digital Workspace solution securely reconnects end users to their virtual desktops or published application instances without losing session information, even when they access the organization from a remote location through the firewall.

From an end user's viewpoint, this functionality is also known as a "follow-me desktop." This type of session persistence can yield significant productivity gains because it allows users to move across devices and between locations while keeping their desktops and applications in the same state.

Security

Maintaining endpoint security is an ever-increasing concern for all organizations. The AlwaysOn Digital Workspace solution offers several new capabilities that significantly reduce the overall risk factor.

- **Endpoint data security** – Network communication between a client device and Horizon 7 virtual desktop infrastructure (VDI) is based on Blast Extreme or PCoIP protocols. Designed for real-time streaming of the graphical user interface (GUI), no data content is included in the communication stream to the user device. Therefore, traditional data protection measures, such as endpoint encryption, are not necessary. Similarly, loss of the end-user device has minimal security consequence because no data is locally stored or cached on the device.
- **User authentication** – The AlwaysOn Digital Workspace supports direct authentication via Active Directory as well as VMware Identity Manager™ (which is a component of the VMware Workspace™ ONE™ platform). Third-party proximity cards can also be added to streamline the authentication process.
- **Antivirus protection** – The AlwaysOn Digital Workspace solution is compatible with most of the top antivirus protection platforms such as Trend Micro, McAfee, Symantec, and Sophos. These platforms are capable of running their services within VMware vSphere® hypervisors. Offloading those services from the virtual desktops yields higher capacity and better performance.
- **Cybersecurity and compliance** – The Horizon 7 platform can significantly reduce risks caused by zero-day vulnerabilities. Security patches can be applied to centrally managed “gold” images and become immediately available to end users’ virtual desktop and published application sessions.

VMware Horizon 7 meets the following compliance standard requirements:

- PCoIP protocol is compliant with FIPS 140-2.
- VMware is SOC 2, Type I certified.
- Supports Criminal Justice Information Services (CJIS) 5.3.
- Supports PCI DSS version 3.0.

AlwaysOn Digital Workspace Architecture

As previously stated, at the core of AlwaysOn Digital Workspace architecture is the notion of multi-pathing. This concept ensures that end user requests for virtual or published applications or desktop sessions can be fulfilled via redundant resources.

The solution is comprised of two types of services required to deliver the desired resiliency:

- **Stateless services** – These are services that do not store, cache or persist data across user sessions; that is, each new session starts from a known gold image. Examples of stateless services are Horizon 7 published application and linked-clone virtual desktop sessions. All user-based data (that is, persona data) is stored separately in UEM folders stored in SMB file shares. The key aspect of stateless services in this architecture is that they can operate within an active-active infrastructure, where a session can be served up from any one of the available Horizon 7 instances because there is no user affinity at session startup time.
- **Stateful services** – This category is made up of services that do not work in an active-active mode (that is, all have operations that require dynamic data that must be persisted). They include VMware Identity Manager, VMware App Volumes™, and VMware User Environment Manager™.

The redundancy considerations associated with this category of service are described in the Horizon 7 Multi-Site Reference Architecture.

Although the overall uptime of this solution is driven by the availability of individual stateful components, explicit designs can be employed to ensure that the target recovery time objective (RTO) is kept below the target 15 minutes duration. Figure 1 shows this architecture at a high level.

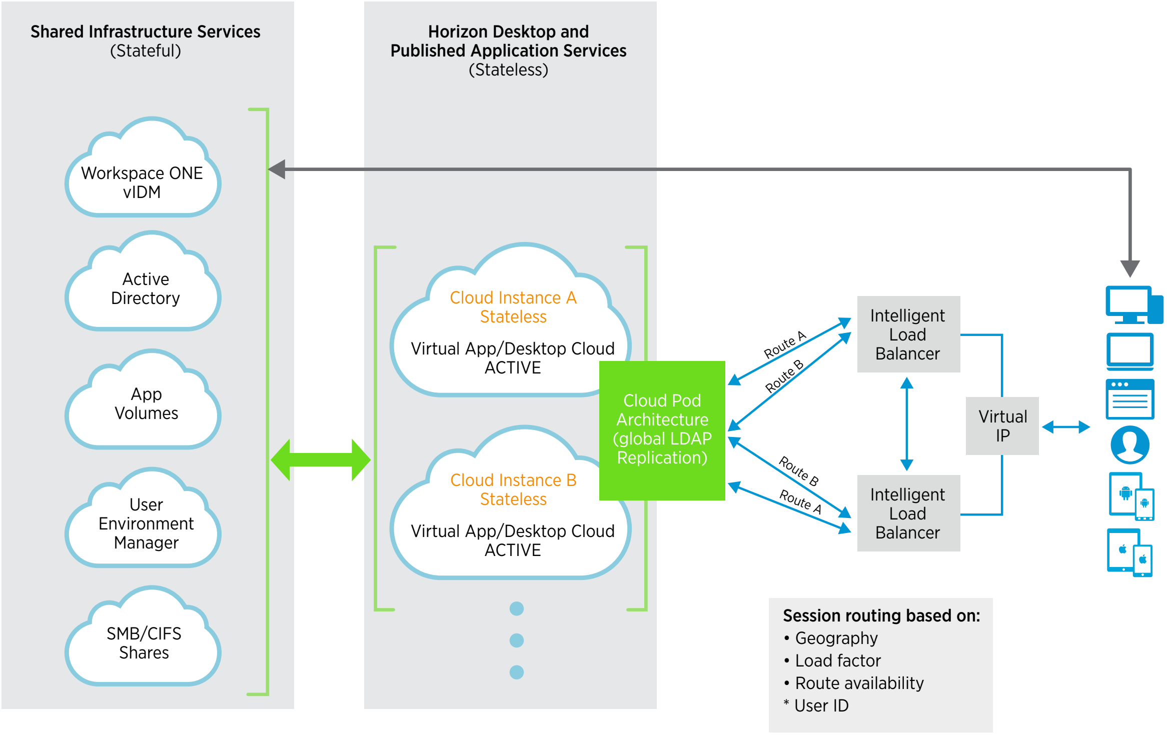

Figure 1: AlwaysOn Digital Workspace High-Level Architecture

Each instance of virtual app/desktop cloud is comprised of View desktops as well as RDSH hosts running published applications—in any combination based on capacity requirements. Horizon 7 instances are independent of each other and while they operate in an active-active mode, access to them is dependent on the availability of the VMware Identity Manager portal service for authentication. Once authenticated, incoming session requests are routed by the intelligent load balancer based on predefined routing logic as well as the availability of each Horizon 7 instance.

Active session states are shared across the instances through the Cloud Pod Architecture (CPA) functionality. This means that an end user with an existing active (logged-in) session is automatically routed to the Horizon 7 instance that is running that session.

The AlwaysOn Digital Workspace architecture can include as many instances of virtual app/desktop clouds as supported by CPA. Each instance is made up of one or more View pods. Please refer to the Horizon 7 technical documentation for the components within a pod.

AlwaysOn Digital Workspace Configuration Options

This section describes various configuration alternatives that accommodate different data center topologies as well as infrastructure cost variables. The final configuration design must be developed during the design phase.

Single Data Center Design

The architecture can contain two instances in one data center.

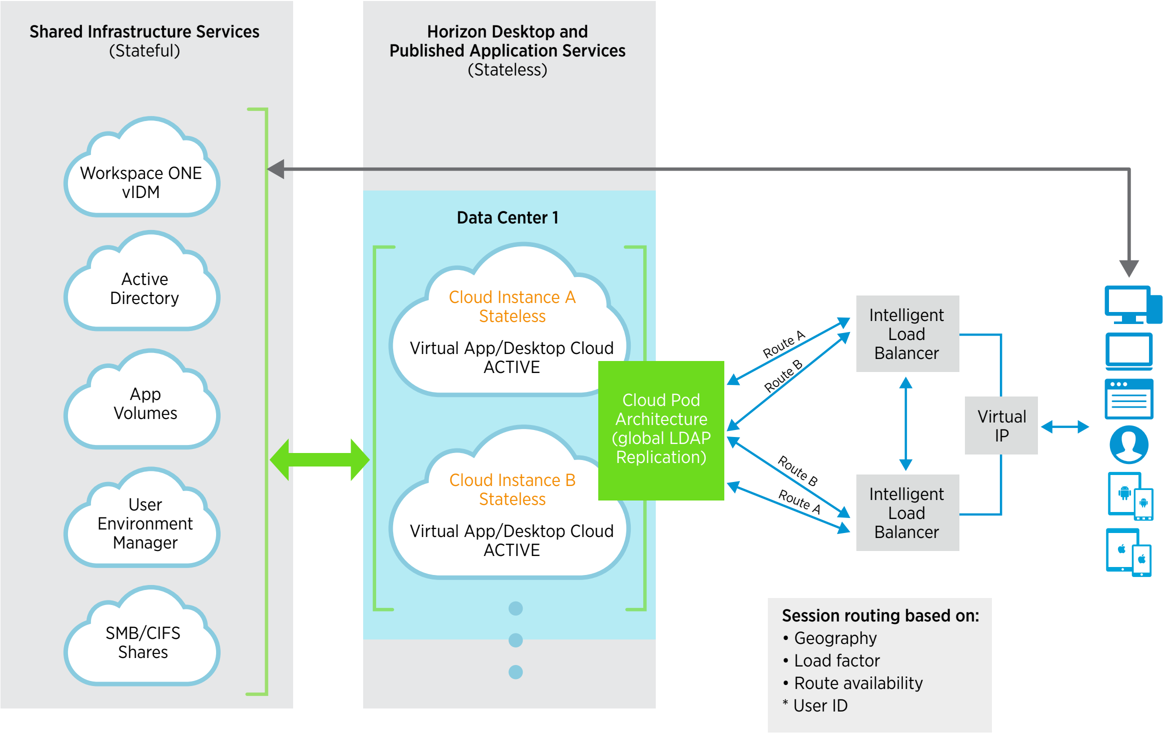
Shared Infrastructure Services
(Stateful)
Workspace ONE vIDM
Active Directory
App Volumes
User Environment Manager
SMB/CIFS Shares
Horizon Desktop and Published Application Services
(Stateless)
Data Center 1
Cloud Instance A
Stateless
Virtual App/Desktop Cloud ACTIVE
Cloud Instance B
Stateless
Virtual App/Desktop Cloud ACTIVE
Cloud Pod Architecture (global LDAP Replication)
Route A
Route B
Route B
Route A
Intelligent Load Balancer
Intelligent Load Balancer
Virtual IP
Session routing based on:
• Geography
• Load factor
• Route availability
* User ID

Figure 2: Single Data Center Design

Dual Data Center Design

The architecture can be set up with two instances split between two data centers.

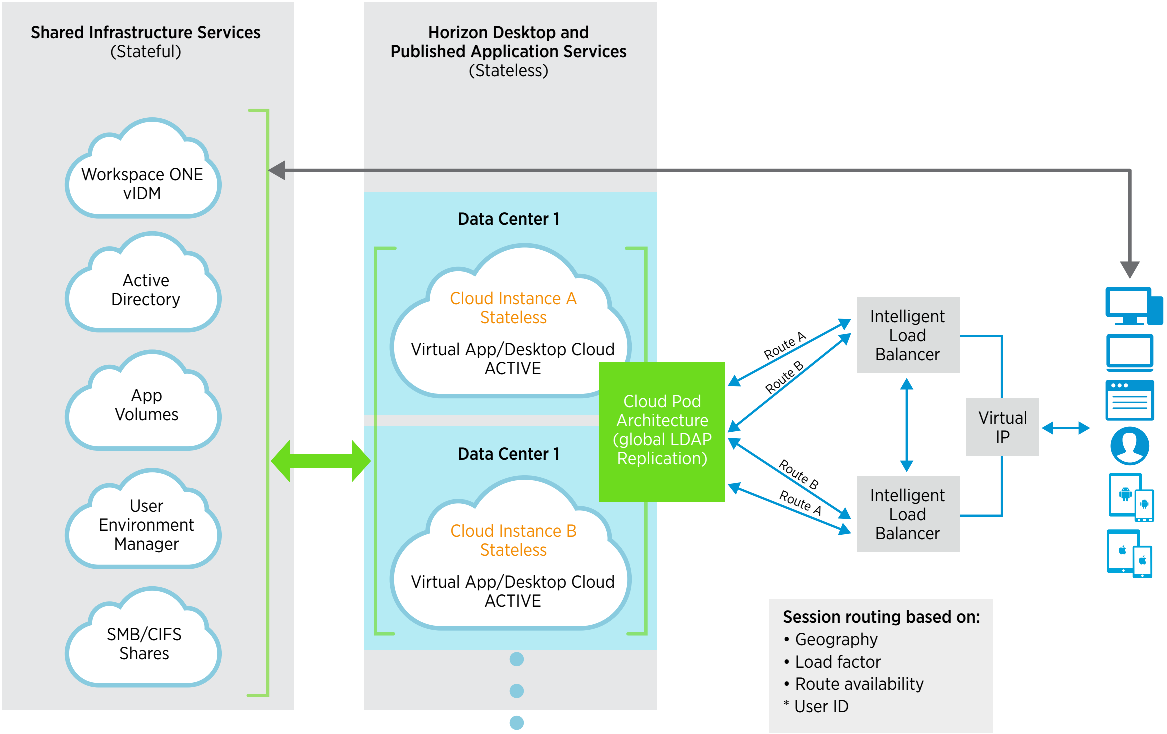

Figure 3: Dual Data Center Design

Single Data Center Design in N+1 Configuration

In this configuration, each virtual app/desktop cloud instance is sized for 50 percent of the total required session capacity. In the event that any one of the instances becomes unavailable, the remaining two instances continue delivering service at 100 percent required capacity—resulting in no service outage and no capacity impact.

The advantage of the N+1 configuration is in reduced hardware costs since the total hardware required for 100 percent session redundancy is only 1.5X (as compared to 2X).

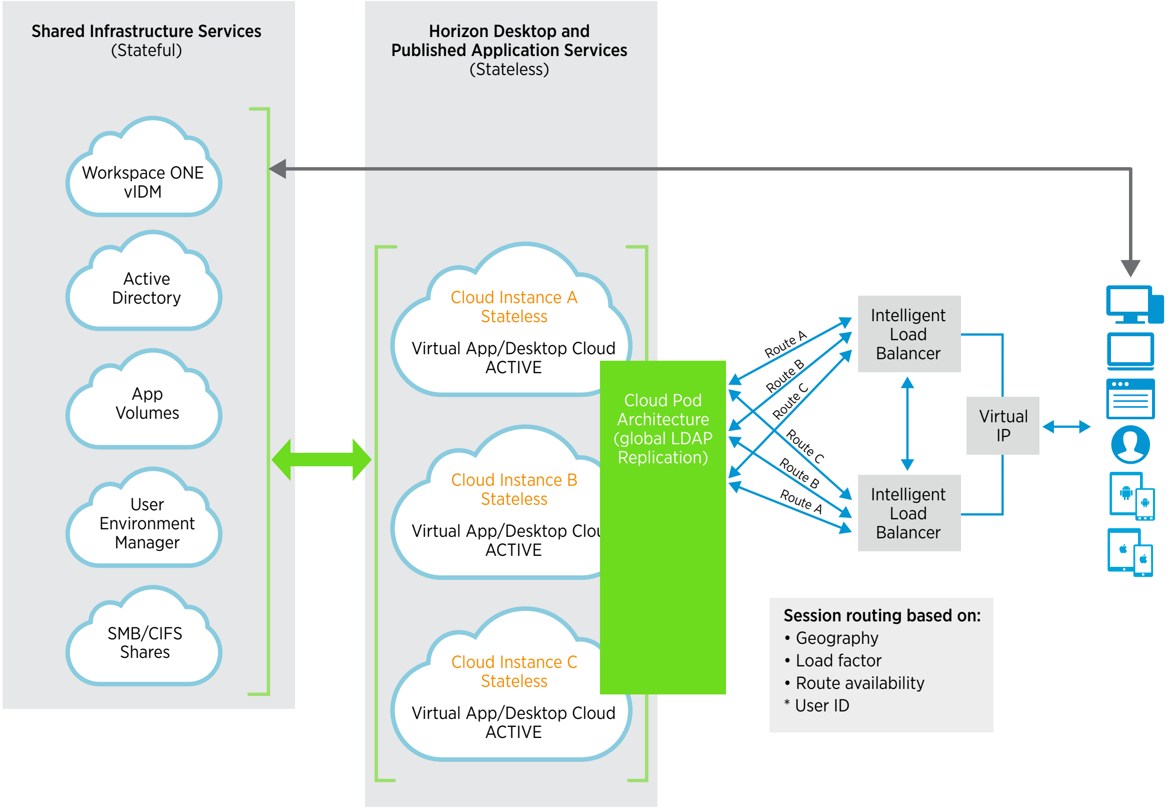

Figure 4: Single Data Center N+1 Configuration

Multiple Data Center Design in N+1 Configuration

This design is an adaptation of the single data center N+1 configuration where Horizon 7 instances are spread across multiple data centers. The advantage remains the same (that is, 1.5X hardware instead of 2X), however, this configuration also provides data center disaster recovery capability.

The same approach can be applied to the scenario where there are only two data centers: primary and secondary. In this case, the N+1 configuration would be comprised of two Horizon 7 instances in the primary data center and one in the secondary data center. The design remains exactly the same where all instances are in active state and deliver application and desktop sessions.

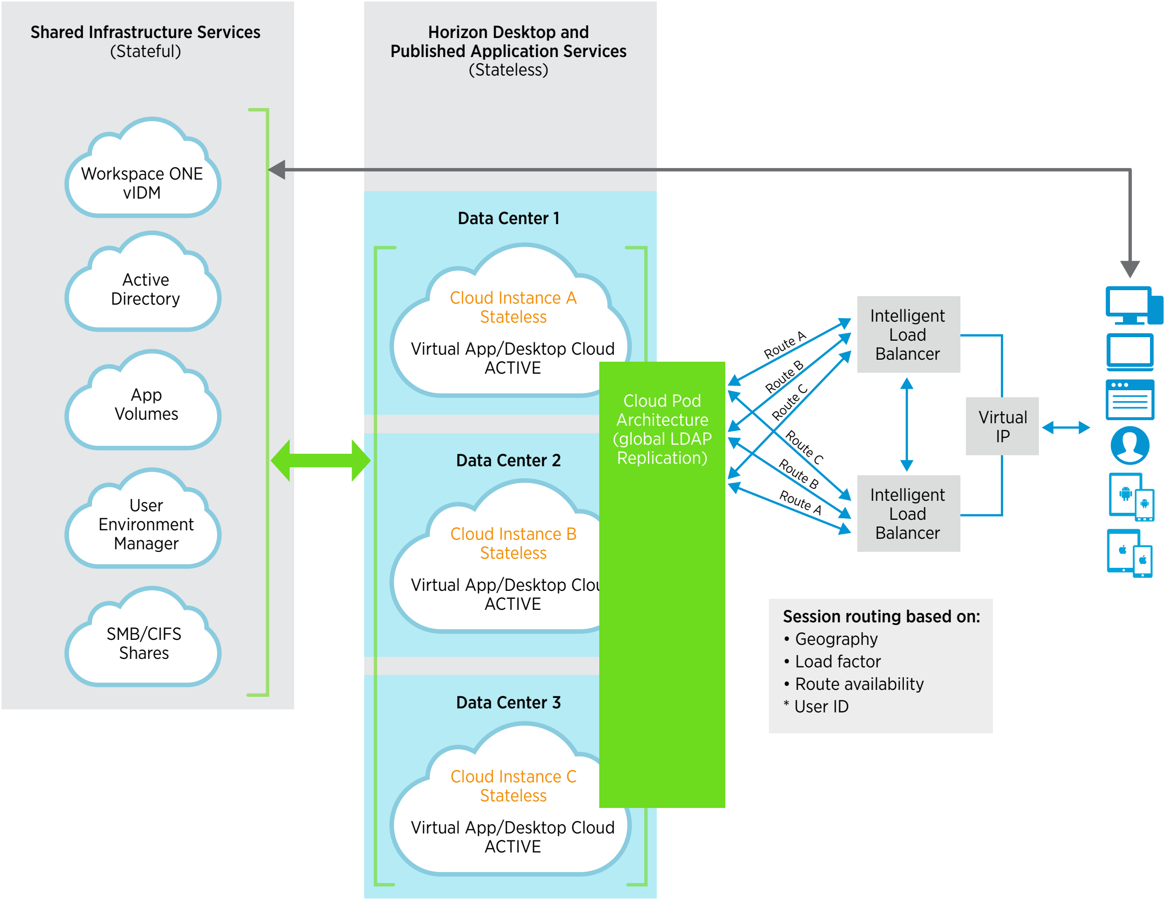

Session routing based on:
- Geography
- Load factor
- Route availability
- * User ID

Figure 5: Multiple Data Center N+1 Configuration

AlwaysOn Digital Workspace Supported Functions

The following functions and associated products are incorporated in the AlwaysOn Digital Workspace design guide.

COMPONENT	DETAILS
End-user authentication	Three modes are supported: • VMware Identity Manager (supports Active Directory and other identity provider platforms) • Via VMware Horizon Client (supports Active Directory) • Via HTML5 browser (supports Active Directory)
Proximity card access	Imprivata OneSign. Two modes are supported: • Active Directory authentication to the VMware Identity Manager portal. Supports SAML authentication to Horizon 7 or RDSH. • Active Directory authentication and launch Horizon Client. Single sign-on to entitled desktops or RDSH applications. Other third-party access products with similar functionality can be included.
Single sign-on (SSO)	Four modes are supported (in any combination): • SAML integration within the VMware Identity Manager portal. • HTML password vault accessed from the VMware Identity Manager portal. • TrueSSO for Windows desktops accessed from the VMware Identity Manager portal. • Imprivata OneSign Windows agent for Windows application sign-on.
Application portal/ catalog	Two modes are supported: • VMware Identity Manager – Shows all applications entitled to authenticated end user. • Direct through Horizon Client – Shows all applications entitled to the end user in Active Directory.
Virtual desktop platform (VDI)	View in Horizon 7
Application publishing platform	RDSH in Horizon 7
Global entitlement	VMware Cloud Pod Architecture (CPA): • Enables multiple View pods to act as a single View environment for entitlement purposes. • Single URL for user access – Unified namespace when used with global load balancing. - Simplified administration – Global entitlement (entitle users across pods and pools).
Application management and delivery	Multiple methods are supported (in any combination): • VMware App Volumes in Horizon 7 • Directly installed in gold image(s) • Via VMware Identity Manager as SaaS applications • Via VMware Identity Manager as XenApp published applications • As VMware ThinApp® packaged applications: - Via App Volumes in a View desktop - In a gold image in a View desktop - As an icon on the VMware Identity Manager portal

COMPONENT	DETAILS
Application isolation	VMware ThinApp (third-party products with similar capabilities are also supported).
User profile management	VMware User Environment Manager (third-party products with similar capabilities are also supported).
Storage platform	VMware vSAN™ (third-party storage platforms with similar capabilities are also supported).
Security and micro-segmentation	VMware NSX®

Table 1: AlwaysOn Digital Workspace Functional Stack

AlwaysOn Digital Workspace – Availability Analysis

The availability analysis is described across the two functional layers previously described in AlwaysOn Digital Workspace Architecture: stateless services (published application and desktop sessions), and stateful services (shared infrastructure services).

Stateless Virtual Application and Desktop Sessions

Independent Horizon 7 instances function in an active-active mode to service incoming requests for virtual desktop and RDSH application services. Using VMware Cloud Pod Architecture (CPA), these instances operate in unison and present a common (single namespace) interface to end users. Using the Global Entitlement construct, desktop and RDSH pools within each Horizon 7 instance can be exposed to incoming session requests.

Upon authentication via VMware Identity Manager, end users with already running desktop (View) or application (RDSH) sessions are automatically routed to their sessions. End users requesting new desktop (View) or application (RDSH) sessions are routed to the appropriate Horizon 7 instance based on logic/policies in the load balancer and CPA.

In the event that one of the Horizon 7 instances becomes non-responsive due to planned or unplanned outage, the intelligent load balancer adjusts its routing logic and new incoming session requests are sent to the other Horizon 7 instance(s). The Recovery Time Objective (RTO) in the event of a Horizon 7 instance failure is driven by the load balancer component, however, it is typically measured in seconds.

Additionally, each Horizon 7 instance itself has built-in redundancies to handle various intra-instance component failures such as failures of Connection Servers, VMware vCenter Server®, physical hosts or clusters, and others.

Stateful Infrastructure Services

As described in the Architecture section, AlwaysOn Digital Workspace architecture uses several stateful infrastructure services as part of its operation. The availability and recovery time objectives associated with these services can directly impact the end-to-end availability of this solution.

These services are covered in the Horizon 7 Multi-Site Reference Architecture document. It is recommended that this document be used as a guide for implementing the referenced services.

AlwaysOn Digital Workspace Service Redundancy Analysis

The AlwaysOn Digital Workspace design has three primary tenets.

- **Eliminate any single point of failure that can cause an outage in the service.** This design objective is accomplished by ensuring that every layer of the stack is configured with built-in redundancy or high availability so that the failure of one component does not affect the overall availability of the desktop service.
- **Configure virtual desktop pools to be nonpersistent (using Instant Clone Technology).** The configuration allows the desktop service cloud to be managed as pools of homogenous virtual desktops without the complexity of managing user profiles or personas for every desktop. Any user can access any available virtual desktop pool based on entitlements and access policies.
- **Leverage the customer's existing enterprise storage (NAS or SAN) environment** for storing persistent user data, such as profiles, data files, and Outlook cache, as well as enterprise desktop gold images. This data must be accessible from any of the Horizon 7 instances.

The following sections examine and validate that the above design requirements are satisfied in the AlwaysOn Digital Workspace solution.

AlwaysOn Virtual App/Desktop Cluster

Redundancy measure: Multiple vSphere hosts are contained within each Horizon 7 cluster.

In the event that a physical host goes down, other hosts in the cluster (also called a block) continue uninterrupted. The only impact of such an outage is a reduction in the concurrent session capacity of the cluster, measured by the virtual machine density of each host. Users with aborted sessions running on the malfunctioning host log back into the service to receive a new session. No reconfiguration is necessary, other than replacing the defective host.

AlwaysOn Digital Workspace Pod

Redundancy measure: Each pod includes at least two Horizon 7 clusters (for both apps and desktops) and up to seven Connection Servers.

In the event that a Horizon 7 cluster becomes non-operational, the remaining cluster(s) continues to deliver full service. The Connection Servers in the pod work around the out-of-service cluster.

In the event that a Connection Server becomes unavailable, the other Connection Servers continue uninterrupted. Up to two (out of seven) Connection Servers can become unavailable before there is any impact to session request routing or brokering.

AlwaysOn Digital Workspace Private Cloud Instance

Redundancy measure: Many large implementations will have more than a single View pod within an instance of AlwaysOn Digital Workspace. In such cases, service is still available in that instance even if an entire pod becomes unavailable. The intelligent load balancer will be able to route new session requests to the other pod(s) in the instance

AlwaysOn Digital Workspace Service

Redundancy measure: AlwaysOn Digital Workspace Service includes at least two private cloud instances.

In the event that one of the private cloud instances becomes non-operational or requires a planned outage, the remaining private cloud instance(s) continue to deliver full service. The global intelligent load balancer sends new requests to the functioning private cloud instance(s).

The AlwaysOn Digital Workspace solution can be designed so that private cloud instances can operate in either active-active or active-passive mode.

In active-active mode, loss of a private cloud instance in its entirety does not impact service availability, because the functioning private cloud instance(s) continues to operate independently.

In active-passive mode, loss of an active private cloud instance requires that the passive instance be promoted to active status, typically through a DNS update.

Storage Infrastructure

Redundancy measure: Using VMware vSAN technology enables the storage architecture for AlwaysOn Digital Workspace to be highly modular. This is ideally suited to the nonpersistent nature of virtual desktops in this solution.

Each vSAN datastore is limited to a single cluster of vSphere hosts. This means that the largest failure domain for a datastore is a single cluster. Additionally, since a vSAN datastore supports RAID configuration, at least two flash storage arrays must become defective before there is any impact on the cluster's operation.

For implementation guidelines and best practices, see the vSAN documentation.

View Connection Server

Redundancy measure: Each View pod supports up to seven Connection Servers that function as a single logical entity. Loss of a Connection Server does not impact the availability of the View pod.

For implementation guidelines and best practices, see the Horizon 7 documentation.

Local Load Balancer

Redundancy measure: Each data center with an AlwaysOn Digital Workspace virtual cloud instance may include a redundant pair of local load balancers. In the event that one load balancer goes down, the other one continues to provide full service.

The AlwaysOn Digital Workspace is agnostic to the load balancer platform. For implementation guidelines and best practices, see the vendor's product documentation.

Global Load Balancer

Redundancy measure: The AlwaysOn Digital Workspace solution requires a redundant pair of global load balancers that provide a global namespace for all incoming desktop session requests. In the event that one load balancer is out of service, the other one continues to provide full service.

For implementation guidelines and best practices, see the vendor's product documentation.

VMware NSX for Horizon

Redundancy measure: Each NSX instance is bound to a single instance of VMware vCenter Server. Because the AlwaysOn Digital Workspace has redundant vCenter Server instances, NSX is also redundant.

NSX is a layer of network security software in the vSphere platform for AlwaysOn Digital Workspace. It provides the capability to segregate desktop pools using security policies (also called micro-segmentation functionality).

For implementation guidelines and best practices, see the NSX documentation.

AlwaysOn Digital Workspace Failure Scenario Analysis

Table 2 lists possible failure scenarios and the associated impact.

SCENARIO NUMBER	FAILURE SCENARIO	IMPACT ON DESKTOP SERVICE AVAILABILITY	DETAILS
1	A vSphere host goes out of service	None	View sessions that are running on the malfunctioning host are lost. However, the affected end users get a fresh View session upon logging in again. The operation of the View cluster to which the vSphere host belongs continues as normal.
2	A VMware vCenter Server instance goes out of service	None	There are two ways in which user impact is mitigated in this scenario. First, vCenter HA configuration protects against outage due to loss of a single server. Secondly, existing View sessions continue as normal. Logged-out sessions are not refreshed, and over time the capacity of the associated pool is reduced to zero. This draw-down effect is transparent to users because new login requests bypass the affected pool.
3	A Connection Server goes out of service	None	Each View pod includes up to seven Connection Servers. In this failure scenario, the remaining Connection Servers take over the load of the failed unit. Moreover, existing sessions are direct between Horizon Client in endpoints and the Horizon Agent in the virtual machines, and are unaffected by Connection Server outage.
4	A local drive in a vSAN cluster goes out of service	None	The vSAN data store is configured as a RAID array with built-in protection against failure of individual drives.
5	An entire View cluster goes out of service	None	AlwaysOn Digital Workspace architecture requires at least two clusters in each pod. In this scenario, the remaining cluster(s) continue normal operation. Total available session capacity is reduced by the size of the malfunctioning cluster.
6	A View private cloud instance is unavailable	None	AlwaysOn Digital Workspace architecture requires at least two View cloud instances. In this scenario, the remaining View private cloud instance(s) continue normal operation. Total available session capacity can remain unaffected in a 2N or N+1 configuration. Services such as App Volumes, User Environment Manager, and VMware Identity Manager, are available from any of the Horizon 7 instances.
7	Component outage in the App Volumes SQL database	SQL AlwaysOn RTO is measured in seconds. Requests during that window could fail.	App Volumes supports Microsoft SQL Server AlwaysOn. In the event that one of the SQL instances is down, service fails over to the other instance.
8	App Volumes SQL Server goes out of service	SQL AlwaysOn RTO is measured in seconds. Requests during that window could fail.	App Volumes supports Microsoft SQL Server AlwaysOn. In the event that one of the SQL instances is down, service fails over to the other instance.

SCENARIO NUMBER	FAILURE SCENARIO	IMPACT ON DESKTOP SERVICE AVAILABILITY	DETAILS
9	An App Volumes Manager goes out of service	None	App Volumes infrastructure service includes multiple App Volumes Managers operating in redundant mode.
10	VMware Identity Manager database is unavailable	SQL Always-On RTO is measured in seconds. Requests during that window could fail.	VMware Identity Manager service is configured with SQL Server AlwaysOn where the database service automatically fails over.
11	Optional component – Imprivata OneSign virtual appliance goes out of service	None	Imprivata OneSign platform includes multiple (up to 10) virtual appliances for authentication services. The OneSign agent maintains a list of available appliances and rolls over to the next available appliance in this scenario.
12	A vSphere host running the View admin services goes out of service	None	vSphere HA will restart any failed admin components such as a Connection Server. These server components will be distributed between multiple hosts so the loss of any one vSphere host does not impact service.

Table 2: AlwaysOn Desktop Solution Failover Scenarios

About the Author and Contributor

Farid Agahi is a Senior End-User-Computing Technical Strategist at VMware and a member of the Office of the CTO Team. Farid has over 30 years of experience in large-scale end-user-computing, solution engineering, and IT transformation. For the past seven years, Farid has been focused on working with VMware enterprise customers to design and execute their IT strategies to transform end-user-computing services. This work includes engineering specific solutions based on virtualization technologies and cloud computing platforms.

Tisa Murdock, Director of Industry Solutions, VMware, contributed content and review to this document.

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com

Copyright © 2017 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at http://www.vmware.com/go/patents. VMware is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies. Item No: VMW-TWP-ALWAYSONDIGWKSP-USLTR-20171220-WEB
12/17