



**Hewlett Packard**  
Enterprise

# **HPE Reference Architecture for VMware Horizon on HPE SimpliVity 380 Gen10**

# Contents

Executive summary.....	3
Introduction.....	3
Simplifying client virtualization.....	3
Technology overview.....	4
VMware Horizon 7 technology overview.....	6
Solution overview.....	7
Management infrastructure.....	8
Desktop infrastructure.....	9
Solution components.....	10
Best practices and configuration guidance for the solution.....	10
Management infrastructure.....	10
Desktop infrastructure.....	14
Business continuity and disaster recovery.....	16
Scenario #1 – Active/Active Site Recovery and Failback with Linked Clones.....	17
Scenario #2 – Active/Active Site Failure – DR with a Tertiary Site.....	19
Scenario #3 – Active/Active Site Recovery and Failback (Full Clones).....	21
Desktop infrastructure capacity and sizing.....	23
Workload description.....	24
Analysis and recommendations.....	25
Summary.....	26
Appendix A: Bill of materials.....	26
Resources and additional links.....	28

## Executive summary

Client virtualization is a top initiative for many IT organizations; driven in part by the promise of a flexible, mobile computing experience for end users and consolidated management for IT. Organizations are looking to client virtualization solutions, like VMware® Horizon 7, to reduce application distribution and administration expenses, to minimize the expenses associated with the operation of their desktop environment, and to improve data security and ensure compliance.

Too often, client virtualization deployments are plagued by sluggish and unpredictable desktop performance and higher than expected up-front capital expenses. As a result, organizations who adopt client virtualization solutions need to compromise between the competing needs of strong solution performance, resiliency, and cost reduction.

The HPE SimpliVity 380 Gen10 is a market-leading hyperconverged infrastructure platform, ideally suited for addressing the challenges of client virtualization. It fulfills the primary requirement of many organizations, performance, without sacrificing economics or resilience. The HPE SimpliVity 380 Gen10 provides:

- Simplified deployment with hyperconverged building blocks.
- Ability to start small and scale out in affordable increments—from pilot to production.
- Highest density of desktops, per node, in the hyperconverged infrastructure category.
- Independently validated, unmatched client virtualization performance for a superb end user experience.
- Deployment of full-clone desktops with the same data efficiency as linked clones.
- Enterprise-class data protection and resiliency.

This Reference Architecture (RA) provides a roadmap for architecting these capabilities and showcases third-party validated Login VSI performance testing. It provides an enterprise-scale architecture for implementing VMware Horizon 7 on HPE SimpliVity 380 Gen10 hyperconverged infrastructure and describes the tests performed by Hewlett Packard Enterprise to validate the efficiency of the recommended solution.

The performance testing illustrates the ability of the HPE SimpliVity 380 Gen10 to deliver an excellent end user experience on VMware Horizon 7. Highlights include:

- Performance at scale: In Login VSI testing, consistently low latency was observed for both hosted desktops running Microsoft® Windows® 10 with Microsoft Office 2016, even as additional nodes are added to the solution.
- Best in class user density: 1500 Knowledge Worker users hosted on only 8x HPE SimpliVity 380 Gen10 nodes, including resilient N+1 design.

**Target audience:** This document is intended for customer IT architects, managers, and administrators, as well as channel partner engineers, professional services personnel, and other IT professionals who plan to deploy the HPE SimpliVity 380 Gen10 hyperconverged solution to support VMware Horizon 7.

**Document purpose:** The purpose of this document is to describe an enterprise-scale reference architecture, highlighting recognizable benefits to technical audiences.

## Introduction

### Simplifying client virtualization

Many businesses are constrained by legacy IT infrastructure that isn't well suited for client virtualization initiatives. Siloed data centers, composed of independent compute, storage, network and data protection platforms with distinct administrative interfaces are inherently inefficient, cumbersome, and costly. Each platform requires support, maintenance, licensing, and power and cooling—not to mention a set of dedicated resources capable of managing and maintaining each component. Rolling out a new complex solution, such client virtualization, becomes a time-consuming effort involving many different technology platforms, management interfaces, and operational teams. Expanding system capacity can take days or even weeks, due to cumbersome processes and administration. Troubleshooting problems and performing routine data backup, replication, and recovery tasks can be just as inefficient.

While grappling with this complexity, organizations also need to address challenges that are unique to virtualization, including:

- Difficulty sizing client virtualization workloads upfront, due to random and unpredictable user behaviors.
- Periodic spikes in demand, such as “login storms” and “boot storms”, that may significantly degrade performance, if not properly handled.
- High cost of downtime in the event of an outage, either unexpected or due to system maintenance.

The HPE SimpliVity 380 Gen10 addresses each of these challenges by providing a scalable, building block-style approach to deploying infrastructure for virtualization. This solution offers the enterprise predictable cost and delivers a high-performing desktop experience to end users.

### **Superior user experience through unmatched client virtualization performance**

The HPE SimpliVity 380 Gen10 enables high performance even at high user density. It eliminates the performance impact of client login storms, delivering 1,000 logins in 1,000 seconds – nearly 3x faster than the standard Login VSI benchmark client login speed and unparalleled in the hyperconverged infrastructure solution market.

### **Linear scalability from pilot to production with cost-effective client virtualization deployments**

The scale-out architecture of HPE SimpliVity minimizes initial capital expense and tightly aligns investments with business requirements. HPE SimpliVity building blocks are added incrementally, providing a massively-scalable pool of shared resources.

### **Enterprise-grade data protection and resiliency for client virtualization workloads**

The HPE SimpliVity 380 Gen10 provides built-in backup and disaster recovery capabilities for the entire client virtualization infrastructure, as well as for supporting management workloads. The solution ensures resilient, highly available, desktop operations and the ability to withstand node failures with no loss of desktops and minimal increase in latency.

## **Technology overview**

The HPE SimpliVity hyperconverged infrastructure solution is designed from the ground up to meet the increased performance, scalability and agility demands of today's data-intensive, highly virtualized IT environments. HPE SimpliVity technology transforms IT by virtualizing data and incorporating all IT infrastructure and services below the hypervisor into compact building blocks. With 3x total cost of ownership (TCO) reduction, HPE SimpliVity delivers the best of both worlds: the enterprise-class performance, protection and resiliency that today's organizations require, with the cloud economics businesses demand.

As a solution, HPE SimpliVity provides a single, shared resource pool across the entire IT stack, eliminating point products and inefficient siloed IT architectures. The solution is differentiated from other converged infrastructure solutions by three unique attributes: accelerated data efficiency, built-in data protection functionality and global unified management capabilities.

- **Accelerated data efficiency:** HPE SimpliVity performs inline data deduplication, compression and optimization on all data at inception across all phases of the data lifecycle, all handled with fine data granularity of just 4KB-8KB. On average, HPE SimpliVity customers achieve 40:1 data efficiency while simultaneously increasing application performance.
- **Built-in data protection:** HPE SimpliVity includes native data protection functionality, enabling business continuity and disaster recovery for critical applications and data, while eliminating the need for special-purpose backup and recovery hardware or software. The inherent data efficiencies of the HPE SimpliVity platform minimize I/O and WAN traffic, reducing backup and restore times from hours to minutes, while obviating the need for special-purpose WAN optimization products.
- **Global unified management:** The VM-centric approach of the HPE SimpliVity platform to management eliminates manually intensive, error-prone administrative tasks. System administrators are no longer required to manage LUNs and volumes; instead, they can manage all resources and workloads centrally, using familiar interfaces such as VMware vCenter™ Server.

An individual HPE SimpliVity 380 Gen10 node includes:

- **OVC – HPE OmniStack Virtual Controller** – A virtual machine is deployed and pinned to the host, used to expose HPE SimpliVity storage as NFS-based vSphere® Datastores. vSphere DirectPath I/O is used to pass through the local SCSI controller and the HPE OmniStack Accelerator Card to the OVC. Multiple OVCS in a vSphere Cluster present a unified namespace of storage across all HPE SimpliVity nodes within a vSphere Cluster.

- **OAC – OmniStack Accelerator Card** – Acknowledges writes, performs data efficiency operations, manages metadata and works with OVC to store metadata in the SSD pool. DRAM is used for transient data. Super capacitors are used to de-stage DRAM to SSD in the event of a power failure.
- **SSD Pool** – SSD drives (number and sizes vary based on HPE SimpliVity 380 Gen10 model) protected with RAID5 or RAID6 using a local SCSI controller – provides a single tier of storage for all system and user data requirements.

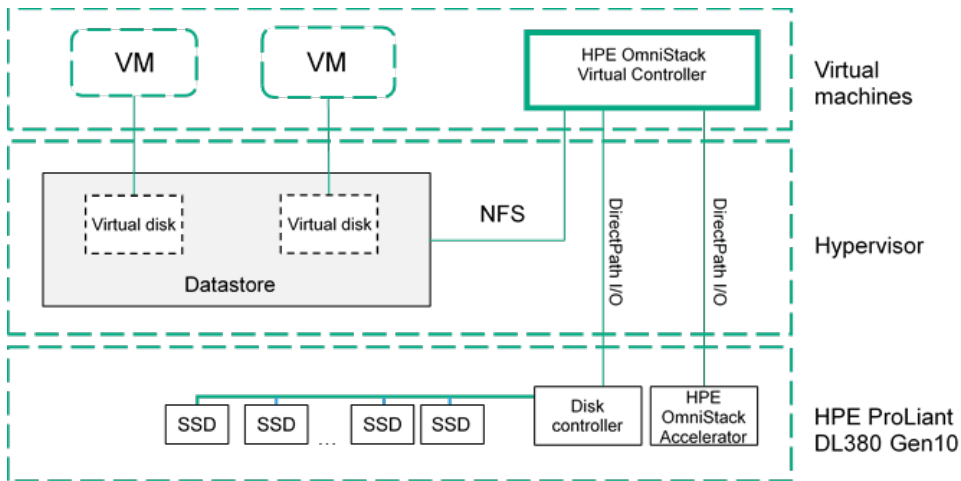


Figure 1. HPE SimpliVity architecture

HPE SimpliVity is a software-defined hyperconverged infrastructure solution. Clustering multiple HPE SimpliVity-powered hyperconverged infrastructure units forms a shared resource pool and delivers mobility, high availability, and efficient scaling of performance and capacity.

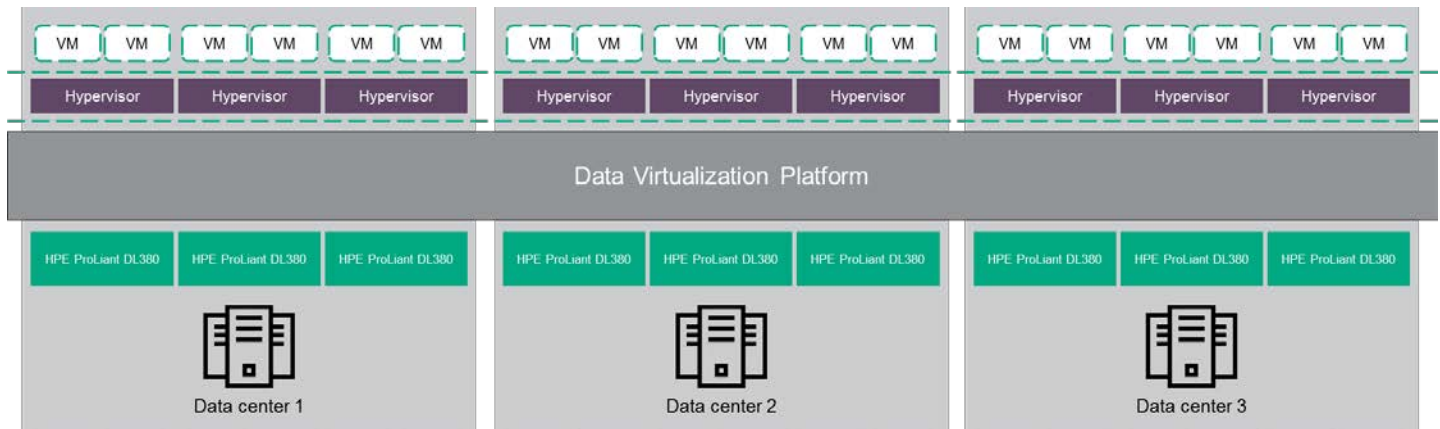


Figure 2. HPE SimpliVity high level architecture

## VMware Horizon 7 technology overview

VMware Horizon 7 is a virtual workspace and application delivery suite designed to allow the administrator to easily and effectively get the data and applications users need to them wherever they are and on any device, all while keeping the company's data securely inside their data center. This data center-centric approach to delivering desktops and applications allows IT staff to get these resources to end users in a way that minimizes the time to delivery and maximizes availability, all without sacrificing security in the process. The components of Horizon 7 focused on in this section are:

- Connection Server
- Security Server
- Universal Access Gateway
- vCenter Server
- View Composer

### Connection Server

The Connection Server is the backbone for the management of the Horizon 7 solution. It provides for management via web interface, as well as an Active Directory Lightweight Domain Services (AD LDS) database containing the objects within the Horizon 7 inventory. The AD LDS database is replicated between Connection Servers in a Horizon 7 pod to keep the state consistent across the pod.

The Connection Server also provides authentication and brokering of sessions to desktops and applications. Authentication is provided via Active Directory integration, as well as support for two-factor authentication. Two-factor authentication can be achieved via RSA SecurID integration or any RADIUS-compliant solution. Once a user has been authenticated, they are directed to their desktop or application by the Connection Server. This is referred to as session brokering.

### Security Server

The Security Server provides for secure access to virtual desktops and applications from outside a perimeter network firewall. It is typically placed inside a DMZ, where it acts as a reverse proxy for Teradici PC-over-IP, Microsoft RDP, or VMware Blast traffic. Security Servers are typically paired 1-to-1 with Connection Servers on an internal network.

### Universal Access Gateway

Universal Access Gateway is similar to the Security Server in the services it provides, as it allows for secure access to resources from outside a perimeter network firewall. It is, however, completely unique from a deployment and operation perspective. As a hardened, Linux-based virtual appliance, it is no longer need to put a Windows Server machine in a DMZ.

### vCenter Server

The Horizon 7 infrastructure is tied into vCenter Server to give access to vSphere resources. This provides the physical compute, storage, and network resources that the virtual desktops will use to function, as well as the operation of the virtual machines themselves.

### View Composer

View Composer is used to create, manage, and destroy linked clone virtual machines within Horizon 7, as well as keep inventory on all objects for which it has responsibility. As of Horizon 6.2, this functionality has been expanded from single-user hosted desktops (VDI) to support multi-user hosted shared desktops using Microsoft Remote Desktop Services (RDS), as well as Microsoft RDS-based hosted application servers.

View Composer is able to perform a number of maintenance tasks on desktop and RDS server pools, which provide for ease of administration of those resources. These include:

- **Recompose** – This allows the administrator to define a new snapshot for the desktops in a pool. This is typically done to update them en masse after patches have been applied to the parent image and a new snapshot taken.
- **Rebalance** – This takes a pool of desktops and redistributes them across datastores as defined in the desktop pool configuration.
- **Refresh** – This takes desktops and puts them back to their original post-deployment state. The delta disk attached to the virtual desktop is reset, thereby removing any changes made to the desktop.
- **Reset** – The administrator is given the ability to reset (warm boot) user desktops from the Horizon 7 administrative interface.

## Solution overview

The solution outlined in this document provides guidance for implementing HPE SimpliVity 380 Gen10 to enable a single client virtualization building block, supporting 1500 Knowledge Worker users. This architecture can be used to scale to tens of thousands of users, by duplicating the building blocks as outlined below.

This solution leverages HPE SimpliVity 380 Gen10 hyperconverged infrastructure as the fundamental element of the design. HPE SimpliVity nodes are combined, forming a pool of shared compute (CPU and memory), storage, and storage network resources. VMware vSphere and Horizon 7 provide a high-performance client virtualization environment that is highly available and highly scalable.

The building block includes:

- HPE SimpliVity 380 Gen10 nodes with Intel Xeon Scalable Gold 5120 CPUs and 315GB usable memory for management workloads
- HPE SimpliVity 380 Gen10 nodes with Intel Xeon Scalable Gold 6150 CPUs and 699GB usable memory for desktop workloads
- 2TB datastore per HPE SimpliVity node – applies to all workloads
- 10GbE networking
- Windows 10 Enterprise for virtual desktops
- Windows Server 2016 for management server workloads
- N+1 design for management workloads and infrastructure where possible

The testing performed for this Reference Architecture was designed to validate the functionality of the system as a 4&4 block. This is defined as four HPE SimpliVity 380 nodes tied together with four HPE ProLiant DL380 Gen9 compute nodes in a 1:1 configuration.

Login VSI 4.1.32.1 was used for testing all workloads, including Login VSI Standard Knowledge Worker for hosted desktops and Login VSI Standard Office Worker for hosted shared desktops. The intent was to find the maximum number of users that could effectively be run on this infrastructure configuration for each workload type. Proper sizing methodology is shown below to ensure that the configuration and load described in this document are production-ready for customer environments and resilient to N+1 design requirements.

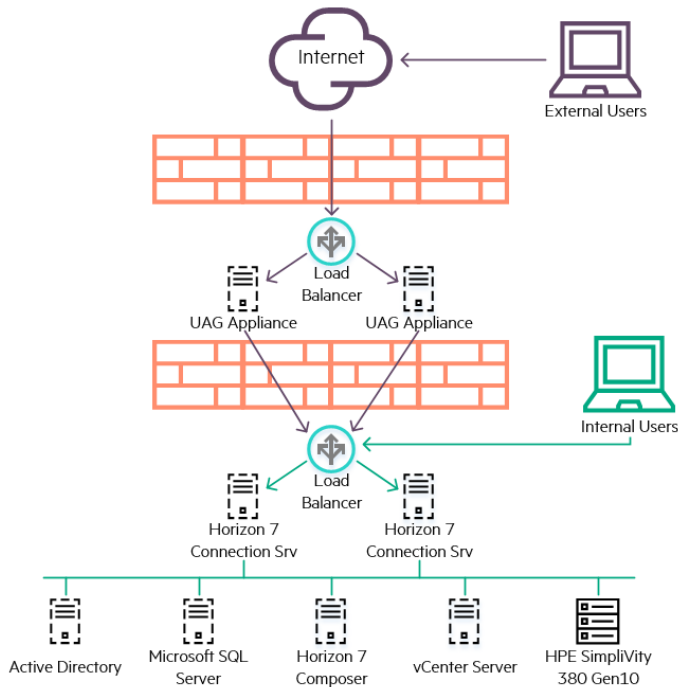


Figure 3. Solution Overview

## Management infrastructure

This section details the HPE SimpliVity environment dedicated to running the management workloads required to support 1500 user sessions, as described in the Desktop Infrastructure Design section. A separate, dedicated, HPE SimpliVity environment is also used for the VMware Horizon 7 virtual desktops, further detailed in the Desktop Infrastructure Design. The management workloads considered in this document are outlined in detail in the Best practices and configuration guidance for the solution section below.

These workloads are represented in Figure 4. Management infrastructure logical design.

- 2x HPE SimpliVity 380 Gen10 servers
- 2x Intel® Xeon® Gold 5120 CPUs (Intel Xeon Scalable 12-core) per node
- 315GB usable memory each
- 2x 2TB datastores
- 10GbE interconnect between systems (no 10GbE switch required, but may be used)



**Figure 4.** Management infrastructure logical design

## Desktop infrastructure

The desktop block for 1500 Knowledge Worker users consists of an eight-node HPE SimpliVity/vSphere Cluster. This configuration has been tested and validated to support the workload as defined, including N+1 design for compute and storage. Results of these tests are available in the Analysis and recommendations section of this document

The infrastructure shown in Figure 5. Desktop Infrastructure Logical Design was used to support these workloads, per cluster, was used to support these workloads, per cluster:

- 8x HPE SimpliVity 380 Gen10 6000 series small all flash nodes
- 2x Intel Xeon Gold 6150 CPUs (Intel Xeon Scalable 18-core) per node
- 699GB usable memory per HPE SimpliVity 380 Gen10 node
- 8x 2TB datastores
- 10GbE networking

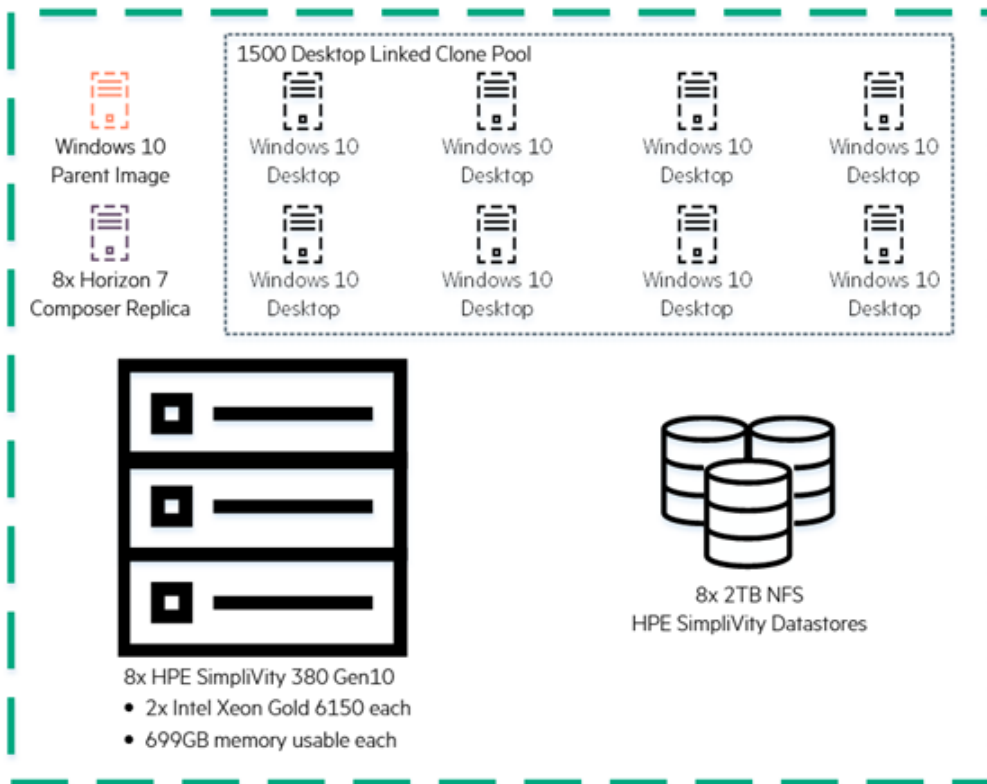


Figure 5. Desktop Infrastructure Logical Design

## Solution components

The following table provides an overview of the configuration for VMware Horizon 7, Login VSI 4.1.32.1, and the tested HPE SimpliVity 380 Gen10 building blocks.

**Table 1.** Solution components

Parameter	Setting
<b>Users per host/HPE SimpliVity 380 node</b>	
<ul style="list-style-type: none"> <li>Virtual desktops</li> </ul>	<ul style="list-style-type: none"> <li>187.5 per HPE SimpliVity 380 Gen10 node</li> </ul>
<b>VMware Horizon 7 Pools Deployed</b>	
<ul style="list-style-type: none"> <li>Virtual desktops</li> </ul>	<ul style="list-style-type: none"> <li>A Linked Clone pool with 1500x Windows 10 virtual desktops</li> </ul>
<b>Login VSI configuration</b>	
<ul style="list-style-type: none"> <li>Benchmark mode</li> <li>Virtual desktop workload</li> <li>Virtual desktop connection</li> <li>Number of launchers</li> </ul>	<ul style="list-style-type: none"> <li>Enabled</li> <li>Knowledge Worker</li> <li>Direct desktop connection (DDC)</li> <li>5 launchers</li> </ul>

## Best practices and configuration guidance for the solution

In this section, configuration guidance and best practices will be shown for deploying VMware Horizon 7 on HPE SimpliVity 380 infrastructure. Both management and desktop workload supporting infrastructure guidance are shown here.

### Management infrastructure

The following infrastructure was deployed to support the VMware Horizon 7 workloads run to produce this Reference Architecture.

**Table 2.** Management infrastructure workloads

Workload	Version	vCPUs	RAM	Disk	OS
vCenter Server Appliance – Desktop (Large)	6.5	16	32GB	640GB	VMware Virtual Appliance
vCenter Server Appliance – Mgmt (Small)	6.5	4	16GB	290GB	VMware Virtual Appliance
Arbiter – Desktop	3.6.2	2	4GB	40GB	Windows Server 2016
Arbiter – Mgmt (lives outside of cluster)	3.6.2	2	4GB	40GB	Windows Server 2016
Microsoft SQL Server x 2 (for Always On Availability Group)	2014 SP2	4	8GB	100GB	Windows Server 2016
File Share Witness (for Always On Availability Group)	N/A	2	4GB	40GB	Windows Server 2016
AD DC/DHCP/DNS x 2	N/A	2	4GB	40GB	Windows Server 2016
Horizon 7 Composer Server	7.4	4	6GB	100GB	Windows Server 2016
Horizon 7 Connection Servers x 2	7.4	4	10GB	100GB	Windows Server 2016
Horizon Unified Access Gateway Servers x 2	3.2.1	2	4GB	20GB	VMware Virtual Appliance

**Table 3.** Management infrastructure vSphere design

Attribute	Value	Rationale
Number of vCenter Servers	1	Only a single vCenter Server is required to support this workload. The desktop-supporting vCenter Server Appliance was deployed as a small.
Number of vSphere Clusters	1	Given the number of HPE SimpliVity systems required to support the given workload, there is no need to split out hosts into separate vSphere Clusters.
Number of vSphere Datacenters	1	A single vSphere Cluster is present and no Datacenter-level separation is necessary.

Attribute	Value	Rationale
vSphere HA Configuration	HA enabled	Enabled to restart VMs in the event of an ESXi host failure
	Admission Control enabled	Ensure VM resources will not become exhausted in the case of a host failure.
	% of cluster resources reserved – 50%	Set to the percentage of the cluster a single host represents.
	Isolation Response – Leave Powered On	Ensure a host isolation event does not needlessly power off desktops.
vSphere HA – Advanced Settings	das.vmmemoryminmb – 9728MB	Both are set to averages of the workloads in the cluster. This serves to set the percentage of cluster resources in HA calculation to that of an average VM.
	das.vmcpuminhz – 1000MHz	

**HPE SimpliVity Servers** – To support the management workloads outlined in this document, a two-host vSphere Cluster, comprised of a pair of HPE SimpliVity 380 nodes, is recommended. Unlike other HCI solutions, HPE SimpliVity fully supports a two-host cluster in its minimum configuration. Using HPE SimpliVity has the ability to start small, with only the infrastructure needed, and scale out as the client virtualization environment grows.

**vCenter Servers** – All roles were deployed on to a single vCenter Server Appliance (VCSA) instance. No CPU or memory pressure were observed during testing, so dedicating servers for each service was unnecessary. Integrated vPostgres databases were used for each VCSA deployed.

**Infrastructure Services (Domain Controllers/DNS/DHCP)** – These services were all co-located on the same virtual machines. No CPU or memory pressure was observed during testing. In-depth Active Directory design and recommendations are outside the scope of this document. Please see <https://msdn.microsoft.com/en-us/library/bb727085.aspx> for more information and best practices.

**Microsoft SQL Server** – All supporting databases for this reference design, were run on a pair of virtual machines running Microsoft SQL Server 2014 Update 1 with Always On Availability Groups. A small file server virtual machine was used as cluster witness. These databases are referenced in the table below.

**Table 4.** Required Microsoft SQL Databases

Database	Authentication	Size	Recovery Mode
Horizon 7 Composer DB	SQL Authentication	Default	Full
Horizon 7 Events DB	SQL Authentication	Default	Full

**Sizing** – Compute, Storage, and Network resources for each infrastructure VM were selected using VMware best practices for Horizon 7 as a baseline and modified based on their observed performance on the HPE SimpliVity 380 Gen10 systems.

**HPE SimpliVity arbiter placement** – The HPE SimpliVity arbiter should always be deployed outside of the HPE SimpliVity infrastructure it manages. The arbiter instance supporting the desktop VCSA will be deployed as a virtual machine in the management cluster. The arbiter instance supporting the management VCSA should be deployed outside of the management cluster. In this instance, it is deployed as a small virtual machine outside of this configuration.

**vStorage API for Array Integration (VAAI)** – VAAI is a vSphere API that allows storage vendors to offload some common storage tasks from ESXi to the storage itself. The VAAI plugin for HPE SimpliVity is installed during deployment, so no manual intervention is required.

**Datastores** – A single datastore per HPE SimpliVity node is required to ensure even storage distribution across cluster members. This is less important in a two node HPE SimpliVity server configuration; however, following this best practice guideline will ensure a smooth transition to a three node HPE SimpliVity environment, should the environment grow over time. This best practice has been proven to deliver better storage performance and is highly encouraged for management workloads. It should be noted that this is a requirement for desktop-supporting infrastructure.

**Table 5.** vSphere VMkernel port to OVC Storage IP Mapping

HPE SimpliVity node (VMKernel IP)	OVC Mapped (Storage IP)
HPE SimpliVity 380 Gen10 #1 (10.111.33.121/22)	OVC #1 (10.111.32.121/22)
HPE SimpliVity 380 Gen10 #2 (10.111.33.122/22)	OVC #2 (10.111.32.122/22)

**Networking** – The following best practices were utilized in the vSphere networking design:

- Segregate OVC networking from ESXi host and virtual machine network traffic
- Leverage 10GbE where possible for OVC and virtual machine network traffic

These best practices offer the highest network performance to VMs running on HPE SimpliVity. Taking this into consideration, a single vSphere Standard Switch is deployed for management traffic.

**Table 6.** Management Infrastructure vSphere Standard Switch Configuration

Parameter	Setting
Load balancing	Route based on Port ID
Failover detection	Link status only.
Notify switches	Enabled.
Failback	No.
Failover order	Active/Active
Security	Promiscuous Mode – Reject MAC Address Changes – Reject Forged Transmits – Reject
Traffic Shaping	Disabled
Maximum MTU	1500
Number of Ports	128
Number of Uplinks	2
Network Adapters	1GbE NICs on each host
VMkernel Adapters/VM Networks	vmk0 – ESXi Management – Active/Active – MTU 1500 VM – vCenter Server – Active/Active – MTU 1500

In addition to the vSphere Standard Switch deployed for management workloads, a single vSphere Distributed Switch is deployed for the remaining traffic, including:

- Virtual Machines
- HPE SimpliVity Federation
- HPE SimpliVity Storage
- vMotion

**Table 7.** Management Infrastructure vSphere Distributed Switch Configuration

Parameter	Setting
Load balancing	Route based on physical NIC load.
Failover detection	Link status only.
Notify switches	Enabled.
Failback	No.
Failover order	Active/Active
Security	Promiscuous Mode – Reject MAC Address Changes – Reject Forged Transmits – Reject
Traffic Shaping	Disabled
Maximum MTU	9000

Parameter	Setting
Number of Ports	4096
Number of Uplinks	2
Network Adapters	10GbE NICs on each host
Network I/O Control	Disabled
VMkernel ports/VM Networks	vmk1 – vMotion vmk2 – Storage vMotion – Active/Standby – MTU 9000 Federation – Standby/Active – MTU 9000 Storage – Standby/Active – MTU 9000 Management VMs – Active/Active – MTU 9000
Port Binding	Static

To illustrate the port-binding and network segmentation described above, please refer to the figure below.

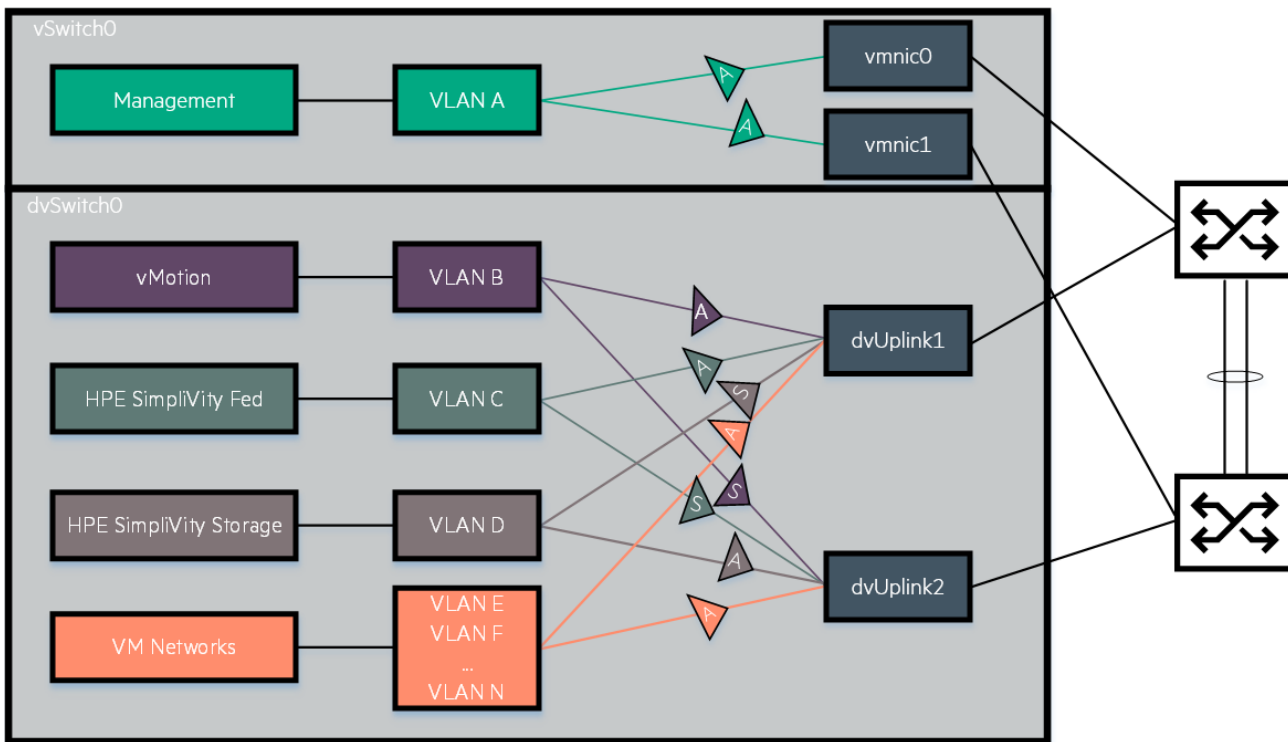


Figure 6. vSphere network logical design

## Desktop infrastructure

The following configurations were performed during our testing for this document.

**Table 8.** Knowledge Worker virtual machine configuration

Attribute	Specification
Operating System	Windows 10 LTSB 64-bit
Virtual Hardware	VM virtual hardware version 13
VMware Tools	Latest
Number of vCPUs	2
Memory	2304MB
Virtual Disk – VMDK	40GB
NTFS Cluster Alignment	8KB
SCSI Controller	VMware Paravirtual
Virtual Floppy Drive	Removed
Virtual CD/DVD Drive	Removed
NIC vendor and model	VMXNET3
Number of ports/NIC x speed	1x 10 Gigabit Ethernet
OS Page file	1.5GB starting and max
Number deployed	1500

**vStorage API for Array Integration (VAAI)** – VAAI is a vSphere API that allows storage vendors to offload some common storage tasks from ESXi to the storage itself. The VAAI plugin for HPE SimpliVity is installed during deployment, so no manual intervention is required.

**Datastores** – An equal number of datastores to the number of HPE SimpliVity 380 Gen10 systems in each vSphere Cluster must be deployed. In this configuration, eight datastores were created for the vSphere Cluster. This is done to more evenly distribute storage load across the HPE SimpliVity systems in the vSphere Cluster, as well as increase the likelihood any given desktop has locality with its VMDK disk.

Each datastore contains a Horizon Composer Replica. The write cache file contains all disk writes of a target device when using a write-protected vDisk (Standard Image).

**Table 9.** vSphere VMkernel port to OVC Storage IP Mapping

HPE SimpliVity node (VMKernel IP)	OVC Mapped (Storage IP)
HPE SimpliVity 380 Gen10 #1 (10.111.33.111/22)	OVC #1 (10.111.32.111/22)
HPE SimpliVity 380 Gen10 #2 (10.111.33.112/22)	OVC #2 (10.111.32.112/22)
HPE SimpliVity 380 Gen10 #3 (10.111.33.113/22)	OVC #3 (10.111.32.113/22)
HPE SimpliVity 380 Gen10 #4 (10.111.33.114/22)	OVC #4 (10.111.32.114/22)
HPE SimpliVity 380 Gen10 #5 (10.111.33.115/22)	OVC #1 (10.111.32.115/22)
HPE SimpliVity 380 Gen10 #6 (10.111.33.116/22)	OVC #2 (10.111.32.116/22)
HPE SimpliVity 380 Gen10 #7 (10.111.33.117/22)	OVC #3 (10.111.32.117/22)
HPE SimpliVity 380 Gen10 #8 (10.111.33.118/22)	OVC #4 (10.111.32.118/22)

**Networking** – The following best practices were utilized in the vSphere networking design:

- Segregate OVC networking from ESXi host and virtual machine network traffic
- Leverage 10GbE where possible for OVC and virtual machine network traffic

These best practices offer the highest network performance to VMs running on HPE SimpliVity. Taking this into consideration, a single vSphere Standard Switch is deployed for management traffic.

**Table 10.** Management Infrastructure vSphere Standard Switch Configuration

Parameter	Setting
Load balancing	Route based on Port ID
Failover detection	Link status only.
Notify switches	Enabled.
Failback	No.
Failover order	Active/Active
Security	Promiscuous Mode – Reject MAC Address Changes – Reject Forged Transmits – Reject
Traffic Shaping	Disabled
Maximum MTU	1500
Number of Ports	128
Number of Uplinks	2
Network Adapters	1GbE NICs on each host
VMkernel Adapters/VM Networks	vmk0 – ESXi Management – Active/Active – MTU 1500 VM – vCenter Server – Active/Active – MTU 1500

In addition to the vSphere Standard Switch deployed for management workloads, a single vSphere Distributed Switch is deployed for the remaining traffic, including:

- Virtual Machines
- HPE SimpliVity Federation
- HPE SimpliVity Storage
- vMotion

**Table 11.** Management Infrastructure vSphere Distributed Switch Configuration

Parameter	Setting
Load balancing	Route based on physical NIC load.
Failover detection	Link status only.
Notify switches	Enabled.
Failback	No.
Failover order	Active/Active
Security	Promiscuous Mode – Reject MAC Address Changes – Reject Forged Transmits – Reject
Traffic Shaping	Disabled
Maximum MTU	9000

Parameter	Setting
Number of Ports	4096
Number of Uplinks	2
Network Adapters	10GbE NICs on each host
Network I/O Control	Disabled
VMkernel ports/VM Networks	vmk1 – vMotion vmk2 – Storage vMotion – Active/Standby – MTU 9000 Federation – Standby/Active – MTU 9000 Storage – Standby/Active – MTU 9000 Management VMs – Active/Active – MTU 9000
Port Binding	Static

## Business continuity and disaster recovery

Planning for the ongoing support of end users in the face of a disaster that takes out a data center is critical when considering the design for client virtualization architecture. The inherent risk introduced by decoupling a user's working environment from their client device is huge if restoration of that decoupled environment isn't planned to the letter. Without planning, there exists a risk removing a user's ability to work. If there is no plan for the restoration of lost data, the entire environment is put at risk. In some cases, that can provide an existential risk to the business.

In this section, three different scenarios are laid out. These are:

- An active/active configuration across multiple sites leveraging linked clones;
- An active/active configuration across multiple sites with a tertiary site for regional DR;
- An active/active configuration across multiple sites leveraging full clone desktops.

These scenarios cover a large percentage of typical customer requirements for highly available desktop services, as well as those where recoverability of the environment is paramount. That being said, this is not an all-inclusive list. These scenarios and the designs detailed therein are flexible and can be adjusted to meet most every failure and recovery scenario.

## Scenario #1 – Active/Active Site Recovery and Failback with Linked Clones

In this first scenario, there are two sites, A & B, with fully configured and independent Horizon 7 pods. These include, per site:

- Desktop vCenter Server
- Connection Servers
- Composer Server
- File Server (for profile data)
- SQL Server - all co-located with the application server

Supporting this infrastructure, there is a management pod per site, with vCenter Server instances in Linked Mode across sites.

Each pod, management and desktop, is run on unique HPE SimpliVity infrastructure. There are two backup policies in place, both set to backup with 10 minute RPO to the opposite data center.

Horizon 7 Linked Clone Pools are configured with one A and one B pool per site - A is primary on Site A, B is primary on Site B. The others are for failover only and are powered off via pool configuration.

For View Persona Manager or VMware User Environment Manager (UEM), there 2 OUs created in AD – one for Site A, one for site B. Each is configured to redirect the profiles/folders for users logging in to those desktops.

### Failure scenario

Site A has users on it. Desktops are completely stateless, with profile data being saved to the file server in each site, depending on which site a user logs into (based on home sites).

1. Site A fails
2. Recovery of file server containing user data for Site A at Site B
3. Site A secondary desktops powered on at Site B
4. Site A users are able to work until Site A comes back online
5. Site A back online
6. Fail back Site A file server from Site B to Site A
7. Site A users fail back to Site A, normal operation proceeds.

### Best Practices

To fail over the file server from one physical site to another, one of a number of things must happen:

- Stretched L2 networking between sites so the file server does not need to change IP address
- Manual intervention to change the IP address of the file server when it moves between sites
- Scripted/automated intervention when moving the file server between sites

In the cases where the DNS A and PTR records will be incorrect for the file server when moving between sites, e.g., when the IP address is changed, the desktops on the recovery side will need to have their DNS caches flushed. The easiest way to make sure this happens is to ensure desktops on the recovery side are powered off via pool setting until after the recovery of the file server and re-IP is complete. This was validated in the scenario above.

To satisfy minimal RTO in an active/active configuration, both sites must be able to support 100% of the desktop workload of both sites. This means double the HPE SimpliVity infrastructure in each site. When this scales to 3+ sites, it is typical to protect for N+1 versus N+N. The Horizon 7 pods themselves (vCenter Server + Composer + Security Server + Connection Server) should be wholly contained within a site. Given the pod and block architecture, there is no reason to fail over these components between sites. Backup for disaster recovery of a site can be done across sites, but each site should remain self-sufficient from a Horizon 7 pod perspective.

The architecture described here will scale up to 1000 users. For every additional 1000 users, another file server will be required. That will, in turn, increase the RTO if manual intervention is required. In opportunities above 2000 users, it is recommended to leverage HPE SimpliVity RapidDR to automate failover process if there is an RTO of less than 30 minutes.

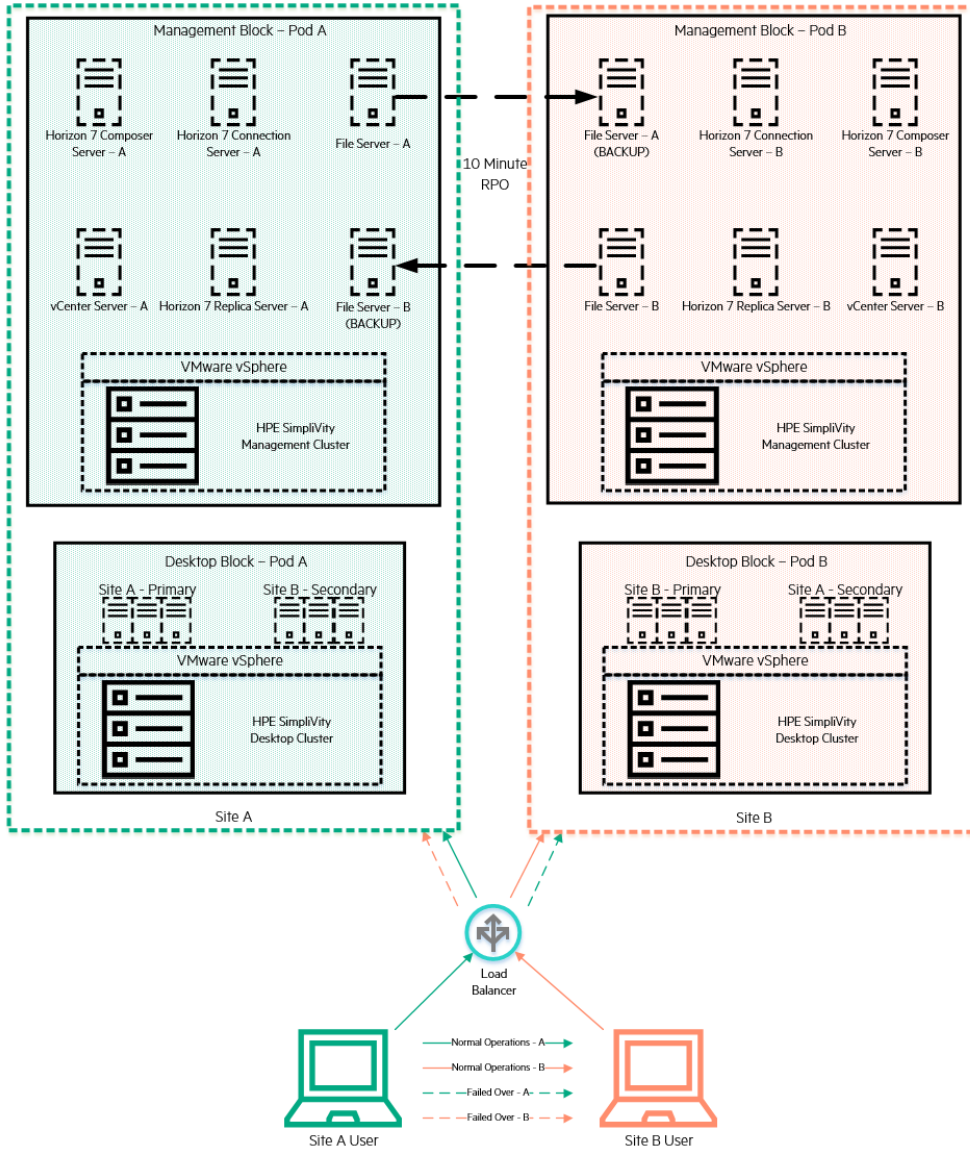


Figure 7. Active/Active Site Recovery with Linked Clones

**Scenario #2 – Active/Active Site Failure – DR with a Tertiary Site**

In this second scenario, there are two sites, A & B, with fully configured and independent Horizon 7 pods. These include, per site:

- Desktop vCenter Server
- Horizon 7 Connection/Replica Servers
- Horizon 7 Composer Server
- File Server (for profile data)
- SQL Server - all co-located with the application server

In addition to these two active sites, there is a disaster recovery site outside the geographical area of the sites serving up desktops for users. This site is in place only for disaster recovery.

Supporting this infrastructure, there is a management pod per site, with vCenter Server instances in Linked Mode across sites. Each pod, management and desktop, is run on unique HPE SimpliVity infrastructure.

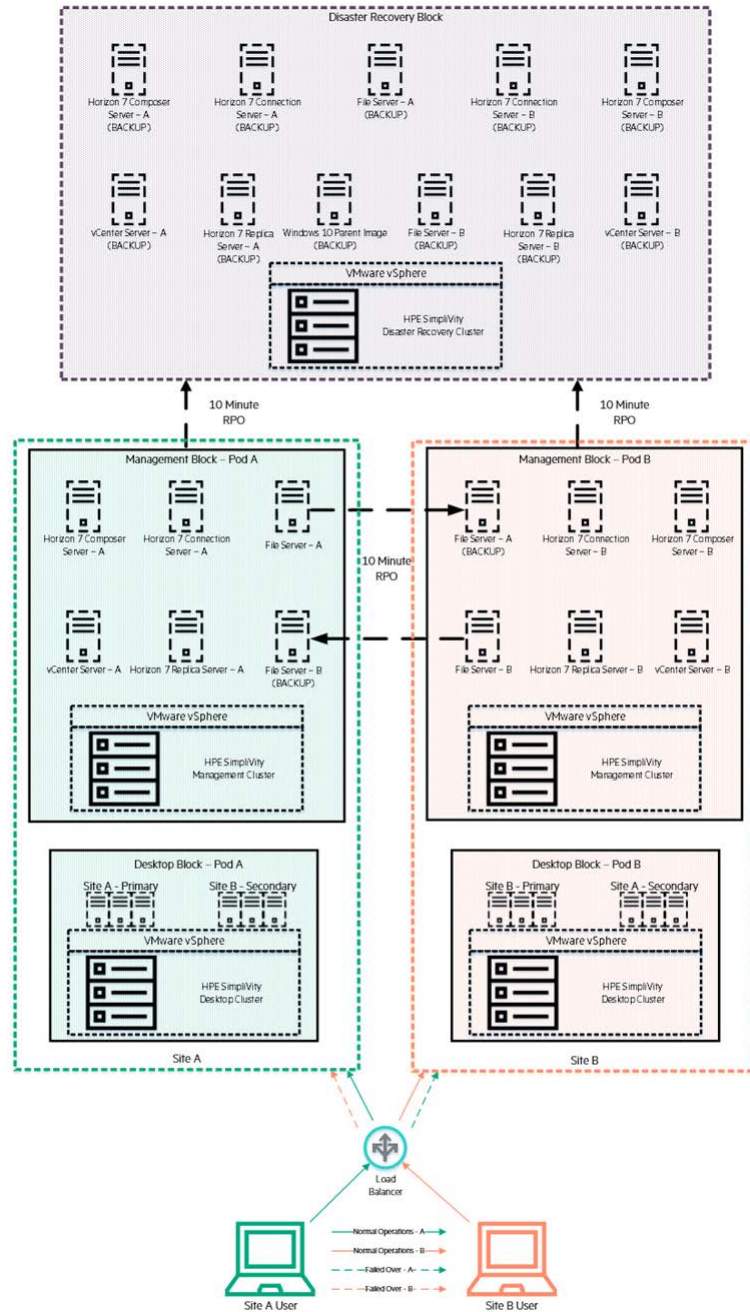
There is a backup policy in place, set to backup with 10 minute RPO to the DR data center, as well as a policy in place to backup user data between pods for easy recovery (see scenario #1 for further details).

Horizon 7 Linked Clone Pools are configured with one A and one B pool per site - A is primary on Site A, B is primary on Site B. The others are for failover only and are powered off via pool configuration.

**Failure Scenario**

Site A and B have users on them. Desktops are completely stateless, with profile data being saved to the file server in each site, depending on which site a user logs into (based on home sites).

- Regional disaster takes out both sites.
- Recovery of each site.
- Pools are rebuilt.
- All sites are back online, normal operation proceeds



**Figure 8.** Active/Active Site Failure – DR with a Tertiary Site

### Best Practices

The requirements to fulfill this scenario for DR to and from a tertiary site are typically minimal given the value it provides. Generally speaking, this can be accomplished with a small amount of HPE SimpliVity infrastructure given the use case. The tertiary site is only used as a backup repository for regional disasters.

There are no stretched networking requirements here like the previous scenario. The DR site should be contained and in a separate geographic (500+ miles away) location.

Backup policies to the DR site will typically be more lax on RPO, so the 10 minute policy outlined in the previous scenario is pretty atypical. In a large environment, this might mean a large amount of data passing over the wire if a tight RPO is defined or change rates are high.

### Scenario #3 – Active/Active Site Recovery and Failback (Full Clones)

In this scenario, there are two sites, A & B, with fully configured Horizon 7 pods. Unlike the previous scenarios, compute and storage resources from both sites available in each management domain. These include, per site:

- Desktop vCenter Server
- Horizon 7 Connection/Replica Servers
- Horizon 7 Composer Server - not used in this scenario
- File Server – not used in this scenario
- SQL Server – all co-located with the application server

Supporting this infrastructure, there is a management pod per site, with vCenter Server instances in Linked Mode across sites.

Each pod has unique management infrastructure, which will be backed up to the opposite management-supporting HPE SimpliVity infrastructure to allow for failover of those management workloads. As mentioned previously, each pod has a desktop block capable of supporting 100% of the full clone workload within its management domain but in the opposite data center. This allows for failover of full clone desktops to the opposite data center while maintaining the relationship between full clone desktop and its vCenter Server and Horizon 7 instance.

There are four backup policies in place, each set to backup with 10 minute RPO to the opposite data center, one pair for desktops and one pair for servers. Horizon 7 full clone pools are configured per site. Both pools are configured to have desktops powered on.

### Failure Scenario

Site A has users on it. Desktops are stateful full clones. User data and applications are unique per desktop.

1. Desktops and management servers backed up across sites.
2. Site A fails.
3. Recovery of management workloads at site B.
4. Recovery of full clone desktops at site B.
5. Site A users are able to work until Site A comes back online.
6. Site A back online.
7. Fail back management workloads from Site B to Site A using SimpliVity Move or backup/restore.

8. Fail back full clone desktops from Site B to Site A using SimpliVity Move or backup/restore.
9. Site A users fail back to Site A, normal operation proceeds.

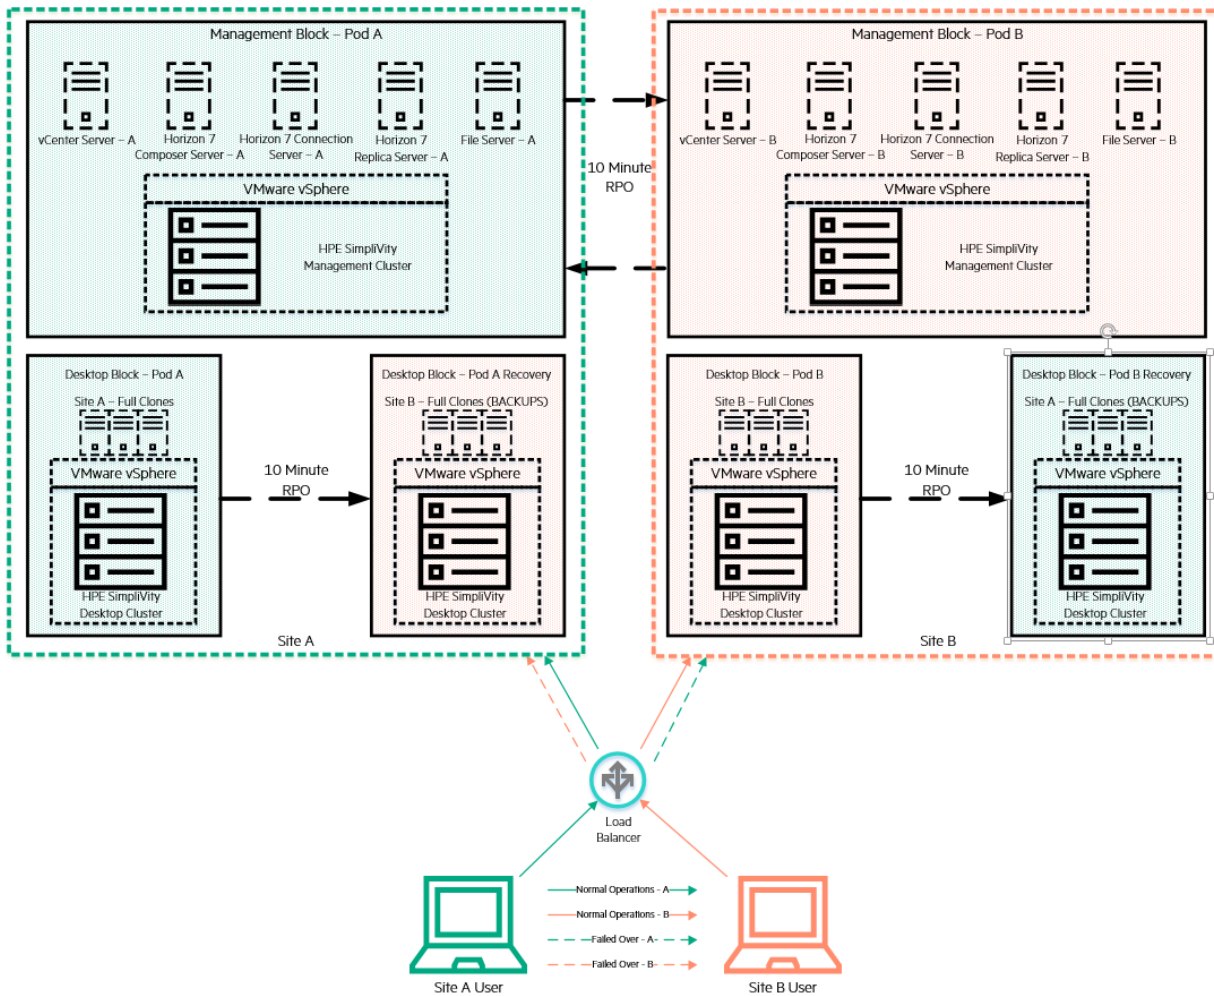


Figure 9. Active/Active Site Recovery and Failback (Full Clones)

**Best Practices**

To fail over the management workloads from one physical site to another, one of a number of things must happen:

- Stretched L2 networking between sites so the management workloads do not need to change IP addresses, or
- Manual intervention to change the IP address of the management when it moves between sites, or
- Scripted/automated intervention when moving the management workloads between sites.

Full clone desktops will be using DHCP to get their IP addresses, so no further intervention should be necessary. It should be noted that the desktops come up after the management workloads in this scenario. Given that the full clone desktops stay within their respective management domains (vCenter Server + Horizon 7 pod), no re-registration of desktops to another pod or failing back to the original pod is necessary, significantly cutting down on the complexity of operations during failover and failback.

To satisfy minimal RTO in an active/active configuration, both sites must be able to support 100% of the workload of both sites. This means double the HPE SimpliVity infrastructure in each site. When this scales to 3+ sites, it is typical to provide N+1 redundancy versus 2N.

## Desktop infrastructure capacity and sizing

**HPE SimpliVity federation and vSphere cluster sizing** –The solution is architected such that workloads are all run in a single vSphere Cluster with the 1500 virtual desktop workloads in a single vCenter Server.

### Note

This solution architecture was designed based on the Login VSI workload size standards. When sizing a production environment, proper assessment and use case definition should be done to accurately size the environment.

**Table 12.** Desktop infrastructure vSphere design

Attribute	Value	Rationale
Number of vCenter Servers	1	Only a single vCenter Server is required to support this workload. The desktop-supporting vCenter Server Appliance was deployed as a large.
Number of vSphere Clusters	1	A single vSphere Cluster is deployed.
Number of vSphere Datacenters	1	A single vSphere Datacenter object is needed for a single vSphere Cluster.
vSphere HA configuration	HA enabled Admission Control enabled % of cluster resources reserved – 12% Isolation Response – Leave Powered On	Enabled to restart VMs in the event of an ESXi host failure Ensure VM resources will not become exhausted in the case of a host failure. Set to the percentage of the cluster a single host represents. Ensure a host isolation event does not needlessly power off desktops.
vSphere HA – advanced settings	das.vmmemoryminmb – 2304MB das.vmcpuminhz – 187MHz	Both are set to averages of the workloads in the cluster. This serves to set the percentage of cluster resources in HA calculation to that of an average VM.
Reservations and limits	Full memory reservation for all desktop workloads	Ensures all desktop workloads have access to memory resources. Also avoids creation of VMkernel swap files on storage.

**HPE SimpliVity Servers** – A single 8-node vSphere Cluster comprised of 8x HPE SimpliVity 380 Gen10 nodes were deployed to support the Knowledge Worker virtual desktop workloads. The following design patterns were observed:

- **Limit physical CPU to virtual CPU oversubscription** – In this configuration, each HPE SimpliVity 380 Gen10 system has 36 physical cores. The OVC takes 4 physical cores, leaving 32 per system for desktop workloads to use. Each virtual desktop VM has two vCPUs.
- **Do not overcommit memory** – In this configuration, each HPE SimpliVity 380 Gen10 system has 768GB of available physical memory. 69GB of memory is reserved for the OVC on each system. That leaves 699GB available per HPE SimpliVity 380 Gen10 system for desktop workloads. Each hosted desktop VM has 2.25GB of memory.

**Table 13.** Desktop Infrastructure Resource Requirements

	HPE SimpliVity nodes	Total Usable	Required	Ratio/Overage
Hosted Desktop CPU	(36 pCPU – 4) x 8 hosts	256 pCPUs	3000 vCPUs	11.7:1
Hosted Desktop RAM	(768GB – 69GB) x 8	5592GB	3375GB	2.2TB spare

## Workload description

Login VSI has standard workloads for use in benchmark testing, including the Knowledge Worker workload which was used in this set of tests. This standard workload is defined as follows by Login VSI, per user.

**Table 14.** Knowledge Worker

Parameter	Setting
Apps Open	5 – 9
CPU Usage	100%
Disk Reads	100%
Disk Writes	100%
IOPS	8,5
Memory	1.5GB
CPU	2 vCPU

As a part of the Login VSI testing, a number of different applications are used to simulate specific types of work in an end user's virtual desktop. Microsoft Office 2016 was used to keep the benchmark as relevant as possible given its popularity over earlier versions. The rest of the applications used in the performance testing are stipulated either by the version of Windows desktop operating system used or the Login VSI setup.

**Table 15.** Applications Used

Applications Used
Microsoft Word 2016
Microsoft Excel 2016
Microsoft Outlook 2016
Microsoft PowerPoint 2016
Internet Explorer 11
MindMap
Flash Player
DoroPDF Printer
Photo Viewer

All performance testing documented utilized the Login VSI (<http://www.loginvsi.com>) benchmarking tool. Login VSI is the industry-standard load testing solution for centralized virtualized desktop environments. When used for benchmarking, the software measures the total response time of several specific user operations being performed within a desktop workload in a scripted loop. The baseline is the measurement of the response time of specific operations performed in the desktop workload, which is measured in milliseconds (ms).

There are two values in particular that are important to note: **VSIBase** and **VSImax**.

- **VSIBase:** A score reflecting the response time of specific operations performed in the desktop workload when there is little or no stress on the system. A low baseline indicates a better user experience, resulting in applications responding faster in the environment.
- **VSImax:** The maximum number of desktop sessions attainable on the host before experiencing degradation in host and desktop performance. The VSImax threshold is set at the VSIBase value plus 1000ms.

Hewlett Packard Enterprise used Login VSI 4.1.32.1 to perform the tests. The VMs were balanced across each of the servers, maintaining a consistent number of VMs on each node. All virtual machines were powered on, registered, and idle prior to starting the actual test sessions.

## Analysis and recommendations


The following are a summarization of the results from the tests described below:

**Table 16.** Summary of Login VSI test results

Test Case	VSIBase	VSImax Average	VSImax reached?
1500 Knowledge Workers	535ms	1055ms	No

### 1500 Knowledge Worker Users

VSIBase for the combined runs was 535ms, and VSImax was not reached in any run. VSImax average for the combined runs was 1055ms, and VSImax threshold was 1535ms. This is indicative of excellent performance in each run, as the point at which the VSImax average (blue line) crosses the VSImax threshold (horizontal red line) is where performance is deemed to be no longer acceptable by end users. As shown in Figure 7, average response time, represented by VSImax average, was consistently low. This is indicative of how well the HPE SimpliVity infrastructure handles the 1500 Knowledge Worker users. Maximum response time, represented in the chart by the plotted red line, exceeds the VSImax threshold toward the end of the test. This is acceptable behavior and expected as the CPU resources in the cluster are exhausted. The focus of the test is to ensure average response times stay within the VSImax threshold, which is nowhere near being breached.



**Figure 10.** 1500 Knowledge Worker users results

## Summary

This Reference Architecture provides guidance to organizations implementing VMware Horizon 7 on HPE SimpliVity 380 Gen10 hyperconverged infrastructure, and describes tests performed by HPE to validate and measure the operation and performance of the recommended solution. This includes third-party validated performance testing from Login VSI, the industry standard benchmarking tool for virtualized workloads.

Organizations are looking to client virtualization solutions like VMware Horizon 7 to reduce software licensing, distribution and administration expenses, and to improve security and compliance. The HPE SimpliVity 380 Gen10 is a market-leading hyperconverged infrastructure platform which helps to deliver the promised benefits of client virtualization, while overcoming many common challenges.

HPE SimpliVity 380 Gen10 for client virtualization provides:

- Simplified deployment with hyperconverged, building blocks.
- Ability start small and scale out in affordable increments—from pilot to production.
- Highest density of desktops per node in the hyperconverged infrastructure category.
- Independently validated, unmatched client virtualization performance for a superb end user experience.
- Deployment of full-clone desktops with the same data efficiency as linked clones.
- Enterprise-class data protection and resiliency.

This Reference Architecture describes solution testing performed in April 2018.

## Appendix A: Bill of materials

The following BOMs contain electronic license to use (E-LTU) parts. Electronic software license delivery is now available in most countries. HPE recommends purchasing electronic products over physical products (when available) for faster delivery and for the convenience of not tracking and managing confidential paper licenses. For more information, please contact your reseller or an HPE representative.

---

### Note

Part numbers are at time of publication/testing and subject to change. The bill of materials does not include complete support options or other rack and power requirements. If you have questions regarding ordering, please consult with your HPE Reseller or HPE Sales Representative for more details. [hpe.com/us/en/services/consulting.html](http://hpe.com/us/en/services/consulting.html)

---

**Table 17.** Bill of Materials (per HPE SimpliVity 380 Gen10 host)

Qty	Product #	Product Description
1	Q8D81A	HPE SimpliVity 380 Gen10 Node
1	Q8D81A 001	HPE SimpliVity 380 Gen10 VMware Solution
1	826884-L21	HPE DL380 Gen10 Intel Xeon-Gold 6150 (2.7GHz/18-core/165W) FIO Processor Kit
1	826884-B21	HPE DL380 Gen10 Intel Xeon-Gold 6150 (2.7GHz/18-core/165W) Processor Kit
1	826884-B21 OD1	Factory Integrated
2	Q8D88A	HPE SimpliVity 384G 6 DIMM FIO Kit
1	Q5V86A	HPE SimpliVity 380 for 6000 Series Small Storage Kit
1	873209-B21	HPE DL38X Gen10 x8/x16/x8 PCIe NEBS Riser Kit
1	873209-B21 OD1	Factory Integrated
1	P01366-B21	HPE 96W Smart Storage Battery (up to 20 Devices) with 145mm Cable Kit
1	P01366-B21 OD1	Factory Integrated
1	804331-B21	HPE Smart Array P408i-a SR Gen10 (8 Internal Lanes/2GB Cache) 12G SAS Modular Controller

Qty	Product #	Product Description
1	804331-B21 OD1	Factory Integrated
1	700751-B21	HPE FlexFabric 10Gb 2-port 534FLR-SFP+ Adapter
1	700751-B21 OD1	Factory Integrated
2	830272-B21	HPE 1600W Flex Slot Platinum Hot Plug Low Halogen Power Supply Kit
2	830272-B21 OD1	Factory Integrated
1	BD505A	HPE iLO Advanced 1-server License with 3yr Support on iLO Licensed Features
1	BD505A OD1	Factory Integrated
1	Q8A68A	HPE OmniStack 16-22c 2P Small SW
1	733664-B21	HPE 2U Cable Management Arm for Easy Install Rail Kit
1	733664-B21 OD1	Factory Integrated
1	867809-B21	HPE Gen10 2U Bezel Kit
1	867809-B21 OD1	Factory Integrated
1	826703-B21	HPE DL380 Gen10 SFF Systems Insight Display Kit
1	826703-B21 OD1	Factory Integrated
1	733660-B21	HPE 2U Small Form Factor Easy Install Rail Kit
1	733660-B21 OD1	Factory Integrated
1	H1K92A3	HPE 3Y Proactive Care 24x7 SVC
1	H1K92A3 R2M	HPE iLO Advanced Non Blade - 3yr Support
1	H1K92A3 Z9X	HPE SVT 380 Gen10 Node (1 Node) Support
1	H1K92A3 ZC0	HPE OmniStack 16-22c 2P Small Support
1	HA114A1	HPE Installation and Startup Service
1	HA114A1 5LY	HPE SimpliVity 380 HW Startup SVC
1	HA124A1	HPE Technical Installation Startup SVC
1	HA124A1 5LZ	HPE SVT 380 for VMware Remote SW St SVC

## Resources and additional links

HPE Reference Architectures, [hpe.com/info/ra](http://hpe.com/info/ra)

HPE Servers, [hpe.com/servers](http://hpe.com/servers)

HPE Storage, [hpe.com/storage](http://hpe.com/storage)

HPE Networking, [hpe.com/networking](http://hpe.com/networking)

HPE Technology Consulting Services, [hpe.com/us/en/services/consulting.html](http://hpe.com/us/en/services/consulting.html)

To help us improve our documents, please provide feedback at [hpe.com/contact/feedback](http://hpe.com/contact/feedback).

## Learn more at

<https://www.hpe.com/us/en/integrated-systems/simplivity.html>



---

**Sign up for updates**

---



---

© Copyright 2018 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Microsoft, Windows, and Windows Server are registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. Intel, Xeon, and Intel Xeon, are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries. VMware and vSphere are registered trademarks of VMware, Inc. in the United States and/or other jurisdictions. vCenter is a trademark of VMware, Inc. in the United States and/or other jurisdictions.

a00053300enw, August 2018