HPE Reference Architecture for Citrix Virtual Apps and Desktops 7 on HPE Simplivity 380

CONTENTS

EXECUTIVE SUMMARY

Client virtualization is a key 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 look to client virtualization solutions, such as Citrix® Virtual Apps and Desktops to reduce distribution and administration expenses, to minimize operating expenses of their desktop environment, and to improve security and 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, adopting organizations often have to compromise between their needs for solution performance, resiliency, and cost reduction.

The HPE SimpliVity 380 is a market-leading hyperconverged infrastructure platform, providing the end-user experience that organizations require, without sacrificing economics or resilience. As such, it is ideally suited for addressing the multiple challenges of client virtualization including:

- Simplified deployment with hyperconverged building blocks
- Ability to start small and scale out in affordable increments, easily moving from pilot to production
- Independent validation of client virtualization performance
- Deployment of full-clone desktops with the equivalent data efficiency of linked clones
- Enterprise-class data protection and resiliency

This Reference Architecture provides a road map for architecting these capabilities and showcases the third-party validated Login VSI performance testing. It provides an enterprise-scale architecture for implementing Citrix Virtual Apps and Desktops 7 1808 hosted desktops on HPE SimpliVity 380 hyperconverged infrastructure. It also 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 to deliver excellent end-user experience in client virtualization deployments, as you scale your environment. The testing highlights include:

- **Performance at scale:** In Login VSI performance testing, a consistent low-latency was observed for hosted desktop implementations, even when additional nodes were added to the solution.
- **User density of 2150 user sessions:** Figure 15 and Figure 16 consists of two separate components. Firstly, 750 performance-critical desktops hosted on a 4&0 setup with four HPE SimpliVity 380 nodes. Secondly, 1400 desktops hosted on a 4&4 setup, combining four HPE ProLiant DL380 compute nodes together with another four HPE SimpliVity 380 nodes to provide a balance between performance and economy. This complete solution of eight HPE SimpliVity 380 nodes and four HPE ProLiant DL380 compute nodes still maintains a resilient N+1 design.

Target audience: This document is intended for customer IT architects, managers, and administrators together with channel partner engineers, professional services personnel, and other IT professionals who plan to deploy the HPE SimpliVity 380 hyperconverged solution to support Citrix Virtual Apps and Desktops 7 1808.

Document purpose: The purpose of this document is to describe an enterprise-scale design, 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 as 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. The challenges are:

- 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 deploy 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 enables high performance even at high user density. It eliminates the performance impact of client login storms, delivering 1000 logins in 1000 seconds – nearly 3x times 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 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 worlds 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. The 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.

HPE SimpliVity 380 node include:

- **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 shows architecture of HPE Simplivity.

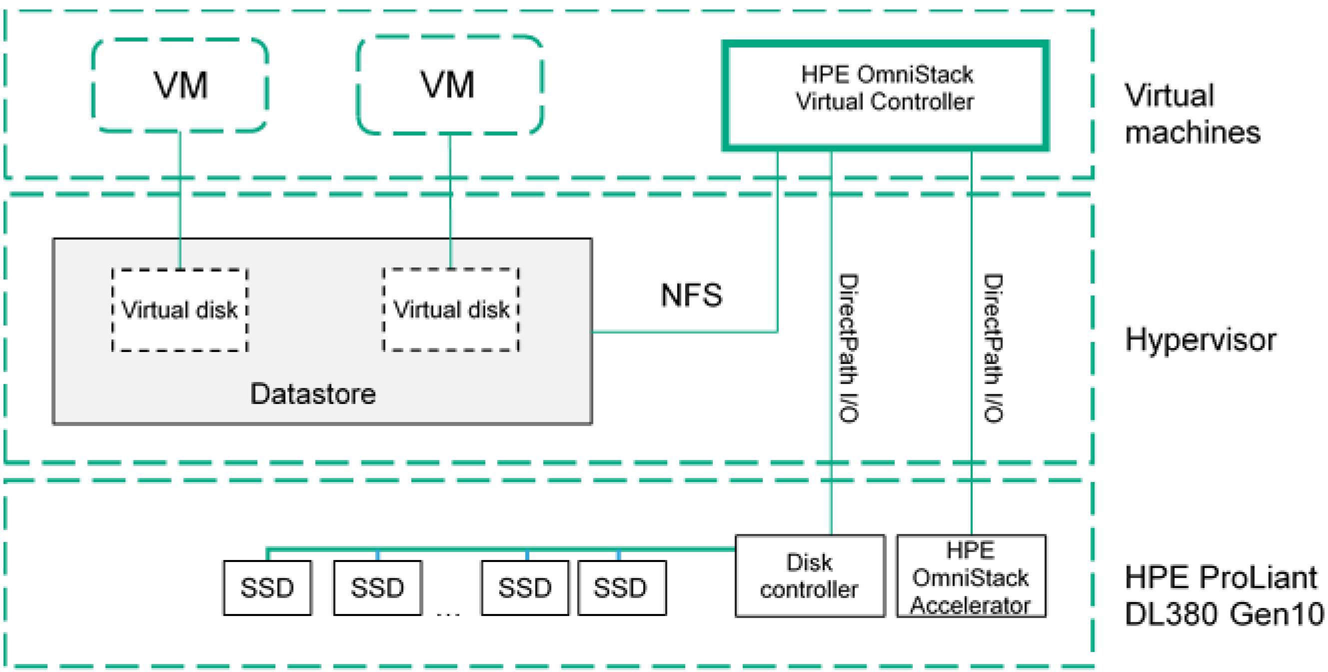

FIGURE 1. HPE SimpliVity architecture

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 shows the high level architecture of HPE Simplivity.

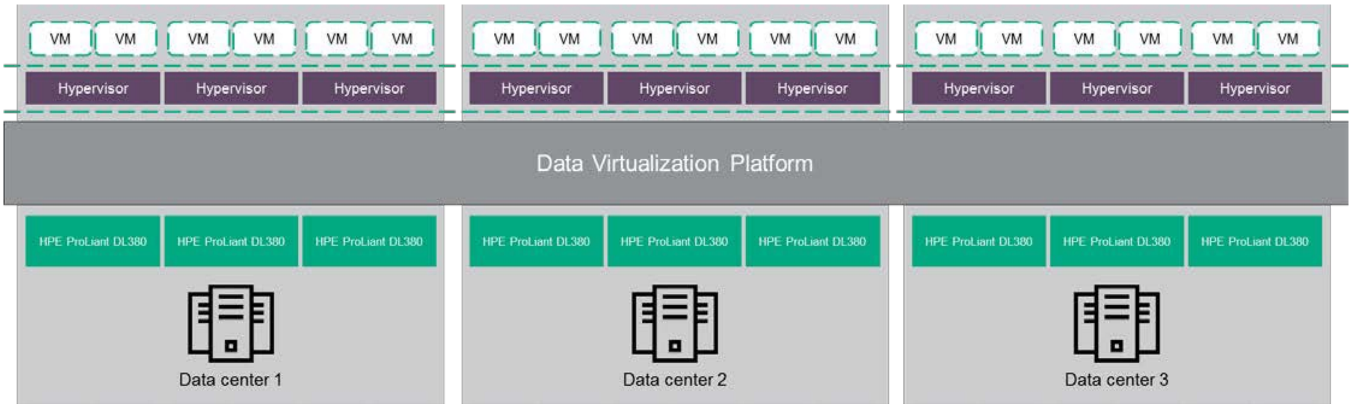

FIGURE 2. HPE SimpliVity high-level architecture

Citrix Virtual Apps and Desktops 7 technology overview

Citrix Virtual Apps and Desktops are applications and desktop virtualization solutions that control virtual machines, applications, licensing, and security while providing anywhere access from any device. Citrix FlexCast Management Architecture (FMA) is a unified architecture that integrates XenApp and Virtual Apps and Desktops to centralize management. Different types of workers across the enterprise have different requirements for their desktops. Some are satisfied with simplicity and standardization, while others need high-performance and personalization. Virtual Apps and Desktops can meet these requirements in a single solution using Citrix FlexCast delivery technology. With FlexCast, IT can deliver multiple variants of their virtual desktop, with each specifically tailored to meet the performance, security, and flexibility requirements of individual users.

With Citrix Virtual Apps and Desktops:

- End users can run applications and desktops, independent of the devices operating system and interfaces
- Administrators can manage the network and provide or restrict access from selected devices or from all devices
- Administrators can manage an entire network from a single data center

Citrix Virtual Apps and Desktops components

Figure 3 shows the key components of Citrix FlexCast Management Architecture (FMA).

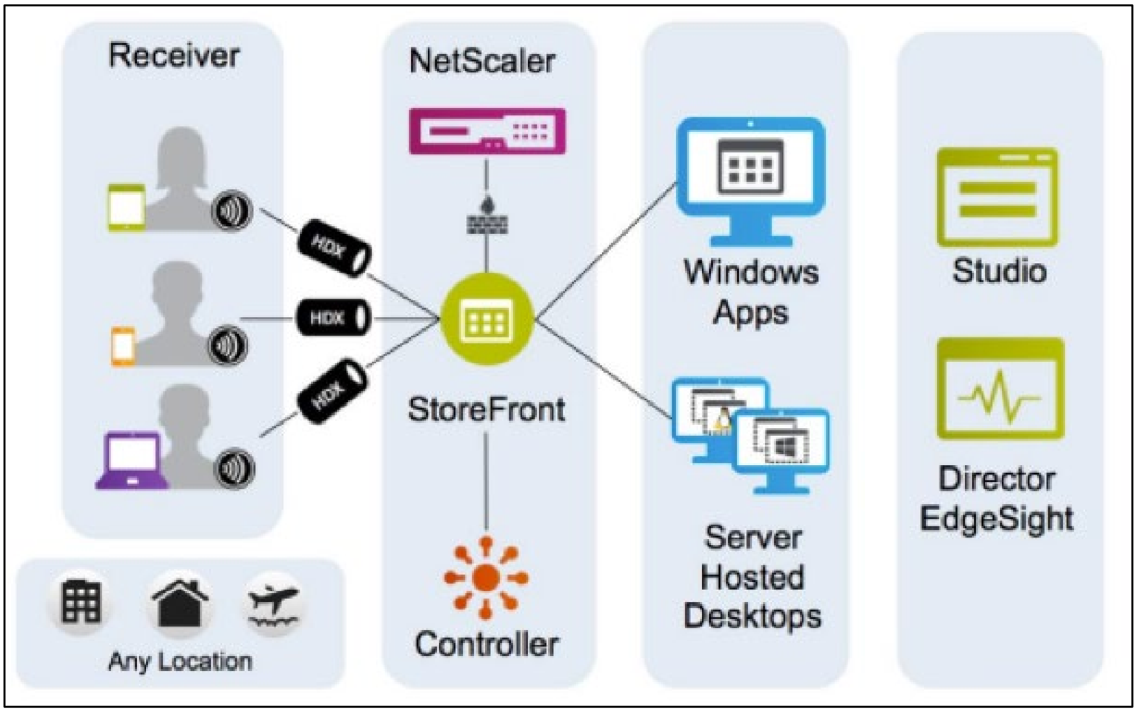

FIGURE 3. Citrix FlexCast Management Architecture

Delivery Controller

The Delivery Controller is the central management component of a XenApp or Virtual Apps and Desktops site. Each site has one or more Delivery Controllers, installed on at least one server in the data center. For site reliability and availability, controllers should be installed on more than one server. If your deployment includes virtual machines hosted on a hypervisor or cloud service, the controller services communicate with the hypervisor to:

- Distribute applications and desktops
- Authenticate and manage user access
- Broker connections between users and their virtual desktops and applications
- Optimize user connections
- Load-balance connections

The controllers broker service tracks which users are logged on and where, what session resources the users have, and if users need to reconnect to existing applications. The Broker Service executes Windows PowerShell cmdlets and communicates with a broker agent on the VDAs over TCP port 80. It does not have the option to use TCP port 443.

Data from the controller services is stored in the site database. The controller manages the state of desktops, starting and stopping them based on user demand and administrative configuration. In some editions, the controller allows you to install profile management so that you can manage user personalization settings in virtualized or physical Windows environment.

Database

At least one Microsoft® SQL Server database is required for every XenApp or Virtual Apps and Desktops site, to store configuration and session data collected and managed by the services that make up the Controller. Install the database within your data center, and ensure it has a persistent connection to the Controller. The site also uses a configuration logging database and a monitoring database. By default, these are installed in the same location as the site database, but you can change this.

Virtual Delivery Agent (VDA)

The Virtual Delivery Agent (VDA) is installed on each physical or virtual machine in your site that you make available to users. These machines can deliver either applications or desktops. The VDA enables the machine to register with the Controller, which in turn allows the machine and the resources it is hosting to be made available to users. VDAs establish and manage the connection between the machine and the user device, they verify that a Citrix license is available for the user or session, and they apply whatever policies have been configured for the session.

The VDA communicates session information to the Broker Service in the Controller through the broker agent included in the VDA. The broker agent hosts multiple plug-ins and collects real-time data. It communicates with the Controller over TCP port 80. It does not have the option to use TCP port 443. Often, the term VDA is used to refer both the agent itself and the machine the agent is installed on.

VDAs are available for Windows Server® and desktop operating systems. VDAs for Windows Server operating systems allow multiple users to connect to the server concurrently. VDAs for Windows desktop operating systems allow only one user to connect to the desktop at a time. A Linux® VDA is also available.

Citrix StoreFront

Citrix StoreFront authenticates users to sites hosting resources and manages the stores of desktops and applications that users access. It can host your enterprise application store, giving users self-service access to the desktops and applications that you provide them. It also keeps track of user's application subscriptions, shortcut names, and other data to ensure users have a consistent experience across multiple devices.

Citrix Receiver

Citrix Receiver is installed on user devices and other endpoints (such as virtual desktops). It provides users with quick, secure, self-service access to documents, applications, and desktops from any of the users devices, including smartphones, tablets, and PCs. Citrix Receiver provides on-demand access to Windows, web, and software-as-a-service (SaaS) applications. For devices that cannot install Citrix Receiver software, Citrix Receiver for HTML5 provides a connection through an HTML5-compatible web browser.

Citrix Studio

Citrix Studio is a management console that enables you to configure and manage your XenApp or Virtual Apps and Desktops deployment. This eliminates the need for separate management consoles for managing delivery of applications and desktops. Citrix Studio provides various wizards to guide you through the process of setting up your environment, creating your workloads to host applications, and desktops, and assigning applications and desktops to users. You can also use Studio to allocate and track Citrix licenses for your site. Studio gets the information it displays from the Broker Service in the Controller, communicating over TCP port 80.

Citrix Director

Citrix Director is a web-based tool that enables IT support and help desk teams to monitor an environment, to troubleshoot issues before they become system critical and to perform support tasks for end users. You can use a single Director deployment to connect to, and monitor, multiple XenApp or Virtual Apps and Desktops sites.

The Citrix Director user interface displays:

- Real-time session data from the Broker Service in the Controller, including data that the Broker Service gets from the broker agent in the VDA
- Historical site data from the Monitor Service in the Controller
- Data about HDX traffic (also known as ICA traffic) captured by HDX Insight from the NetScaler, if your deployment includes a NetScaler and your XenApp or Virtual Apps and Desktops edition includes HDX Insight

Citrix Director also allows you to view and interact with the user sessions, using Windows Remote Assistance.

Citrix License Server

The License Server manages your Citrix product licenses. It communicates with the Controller to manage licensing for each user session and with Studio to allocate license files. You must create at least one license server to store and manage your license files.

NetScaler Gateway

When users connect from outside the corporate firewall, XenApp or Virtual Apps and Desktops can use Citrix NetScaler Gateway (formerly Access Gateway) technology to secure these connections with TLS. The NetScaler Gateway or NetScaler VPX virtual appliance is an SSL VPN appliance that is deployed in the demilitarized zone (DMZ) to provide a single secure point of access through the corporate firewall.

Citrix Provisioning Services

Citrix Provisioning Services (PVS) is an optional component of XenApp and Virtual Apps and Desktops available with some editions. It provides an alternative to Machine Creation Services (MCS) for provisioning virtual machines. Citrix Provisioning Services streams the master image to a user device, whereas MCS creates copies of a master image. PVS doesn't require a hypervisor to do this, so you can use it to host physical machines. When Citrix Provisioning Services is included in a site, it communicates with the controller to provide users with resources. For this Reference Architecture, we used Machine Creation Services exclusively and so Citrix Provisioning Services was not used for any of the results presented.

Citrix provisioning methods

Virtual Apps and Desktops 7 1808 has two integrated solutions to provide different benefits to business needs, namely, Provisioning Services and Machine Creation Services.

Citrix Provisioning Services 7 1808

Most enterprises struggle to keep up with the proliferation and management of computers in their environment. Each computer, whether it is a desktop PC, a server in a data center, or a kiosk-type device, must be managed as an individual entity. The benefits of distributed processing come at the cost of distributed management. It costs time and money to configure, update, support, and ultimately decommission each computer. The initial cost of the machine is often dwarfed by operational costs.

Citrix Provisioning Services takes a very different approach from traditional imaging solutions by fundamentally changing the relationship between hardware and the software that runs on it. By streaming a single shared disk image (vDisk) rather than copying images to individual machines, Citrix Provisioning Services enables organizations to reduce the number of disk images they manage, even as the number of machines continues to grow. This provides the efficiencies of a centralized management together with the benefits of distributed processing. Also, because machines are streaming disk data dynamically and in real time from a single shared image, machine image consistency is ensured. At the same time, large pools of machines can completely change their configuration, applications, and even OS in the time it takes them to reboot.

Using Citrix Provisioning Services, any vDisk can be configured in Standard Image mode. A vDisk in Standard Image mode allows many computers to boot from it simultaneously, greatly reducing the number of images that must be maintained and the amount of storage that is required. The vDisk is in read-only format and the image cannot be changed by target devices.

PVS vDisk can be assigned to a single target device in Private Image mode, or to multiple target devices in Standard-Image mode.

When provisioning, use the following best practice settings:

- **vDisk attributes**
  - Read only during streaming
  - Shared by many VMs, simplifying updates
- **Write cache attributes**
  - One VM, one write cache file
  - Write cache file is empty after VM reboots
  - Recommended storage protocol: NFS for write cache
  - The write cache file size is 4 GB to 10 GB for hosted desktops. If vDisk is used in the deployment, the write cache size will be smaller.

Provisioning server RAM cache

A RAM write cache option (cache on device RAM with overflow on hard disk) is available in provisioning server 7 1808. Write cache can seamlessly overflow to a differencing disk and RAM cache become full. At the Citrix Provisioning Services console, select vDisk properties and choose Cache type as Cache in device RAM with overflow on hard disk. This will use hypervisor RAM first and then use hard disk. Choose the RAM size based on OS type. Refer https://support.microsoft.com/en-us/help/4028142/windows-10-system-requirements to know the system requirements.

Figure 4 shows the Citrix Provisioning Services file layout.

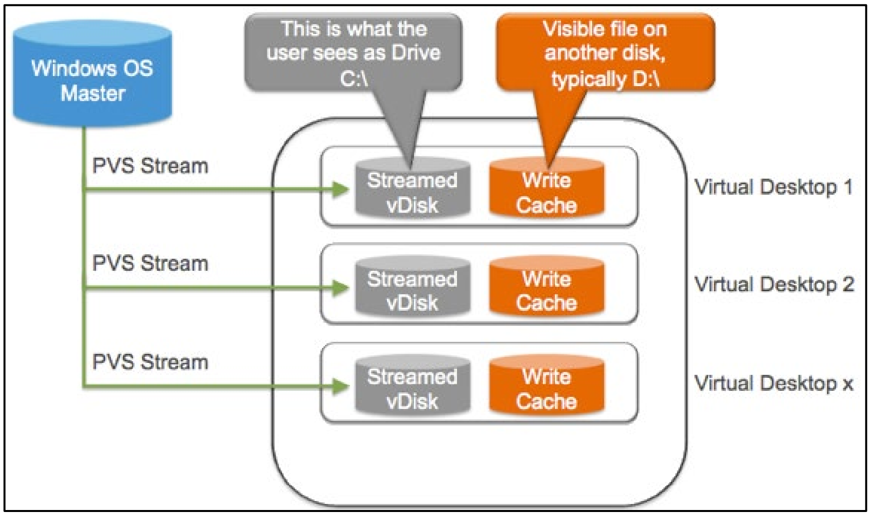

FIGURE 4. Citrix Provisioning Services file layout

Machine Creation Services

Machine Creation Services (MCS) is an image delivery technology, integrated within Virtual Apps and Desktops. Machine Creation Services utilizes hypervisor APIs to create a unique, read-only thin-provisioned clone of a master image where all writes are stored within a differencing disk. When you provision desktops using MCS, a master image is copied to each datastore. This master image copy uses the hypervisor snapshot clone. Within minutes of the master image copy process, Machine Creation Services creates a differencing disk and an identity disk for each VM. The size of the differencing disk is same as the master image in order to host the session data. The identity disk is normally 16 MB and is hidden by default. The identity disk has the machine identity information such as host name and password. Figure 5 shows the Machine Creation Services file layout.

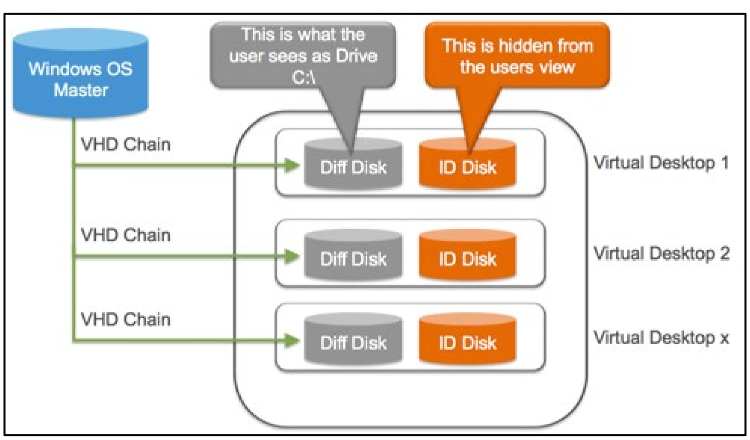

FIGURE 5. Machine Creation Services file layout

Machine Creation Services Storage Optimization (MCSIO) is a feature within Machine Creation Services provisioning that was introduced in XenApp and Virtual Apps and Desktops 7.9. This feature helps to reduce traffic to shared storage in XenApp and Virtual Apps and Desktops environments.

Machine Creation Services provisioning leverages the underlying hypervisor. However, the resulting input/output (I/O) profile of a Machine Creation Services provisioned environment is dictated by your chosen hypervisor and shared to network storage. Machine Creation Services Storage Optimization reduces I/O load through a 2-tier caching system. An in-memory cache, known as the temporary memory cache, is used as the first storage tier while the temporary disk cache provides a second tier by overflowing onto an additional disk attached to the provisioned machine. On the surface, this is comparable to Citrix Provisioning Services RAM cache with overflow. However, MCSIO is primarily intended for caching of write IOPS while it can read cached data and can impact read IOPS on shared storage.

The changes in the architecture introduce a subtle, but important, change to how MCSIO operates compared to standard MCS. In MCS, the delta disk caches all guest read/write operations and resides on the same storage as the master disk. Now, you can separate persistent and temporary data. You can store the temporary cache disk on different storage to the master disk, such as local hypervisor storage or any supported shared storage and then tailor your RAID configurations to the workload.

To achieve this, machines provisioned using MCSIO have an additional driver to intercept and manage I/O operations. When the operating system boots, it starts drivers, including the MCSIO driver, which are read from the base disk. Only after loading the MCSIO drive, it can write to temporary cache. Then, the MCSIO driver can intercept and cache the writes, which would have gone to the delta disk in its temporary memory and disk cache. This results in the temporary cache offloading the I/O operations that were previously handled by the delta disk. The delta disk is still present on MCSIO provisioned machines, but it is only used by MCSIO for maintenance tasks. The MCSIO file layout is in Figure 6.

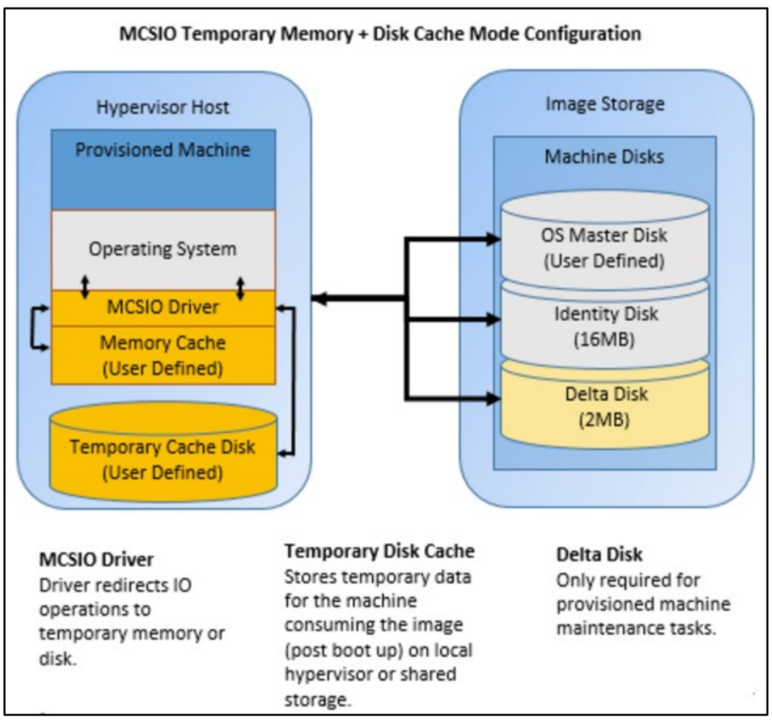

FIGURE 6. MCSIO file layout

Machine Creation Services full clones

The ability to deploy full clones from MCS was introduced in version 7.11 to support storage migration, VM migration, and backup solutions that don't support a delta (linked clone) structure. This allows for a full copy of the base disk to be deployed, instead of a delta, as is the case with MCS fast clones. With the HPE SimpliVity platform, these full copy base disks are stored as space efficiently as delta-based clones, so you can leverage MCS full clones without any negative impact on storage capacity. Figure 7 shows the MCS full clone disk layout.

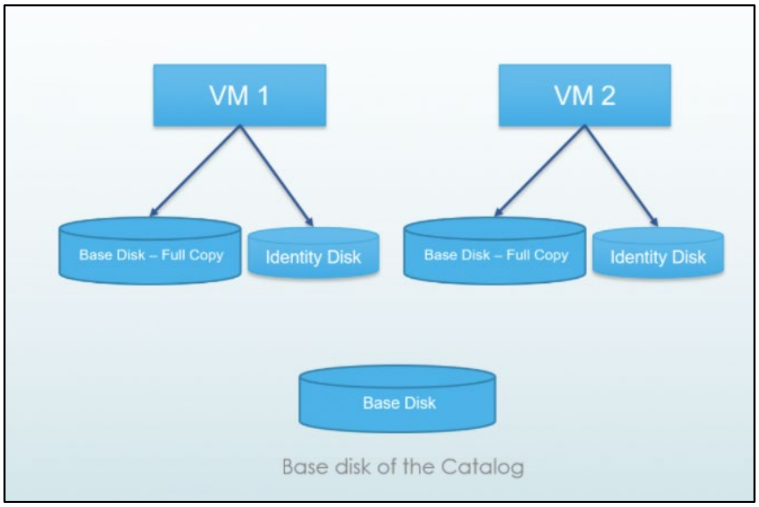

FIGURE 7. MCS full clone disk layout

Desktop types

The Desktop types are:

- **Hosted Shared Desktops (HSD):** In this scenario, many users are supported using a shared, server-based environment. Users get a desktop interface that can look like Windows 7 or Windows 10 but is in fact a published desktop on a Virtual Apps and Desktops server. These are inexpensive, locked-down Windows virtual desktops hosted from Windows Server operating systems. They are well suited for users, such as call center employees, who perform a standard set of tasks.
- **Hosted Virtual Desktops:** A Windows 7 or Windows 10 desktop runs as a virtual machine with a single user connecting remotely. As a result, a single user's desktop configuration has not impacted other users, so users can fully personalize their setup. There are many flavors for the hosted virtual desktop model (for example, existing, installed, pooled, dedicated and streamed), but they are all located within the data center. Each of these virtual desktops run on a Microsoft Windows desktop operating system rather than running in a shared, server-based environment.

For more information, see Citrix Virtual Apps and Desktops Release 7 1808.

Citrix Ready converged infrastructure

Digital transformation is quickly impacting every industry making it increasingly difficult to keep pace with an evolving digital workplace. Many businesses, especially those with smaller IT teams and lower budgets, struggle in implementing new solutions and delivering enterprise-level services like desktop virtualization and BYOD.

To overcome these challenges, Hewlett Packard Enterprise and Citrix partnered to create Citrix Ready converged appliance with HPE SimpliVity and Citrix Cloud. This integration makes it easy to virtualize your desktop apps and data on a central platform. For mid-sized businesses and enterprise remote office/branch office facilities, Citrix and Hewlett Packard Enterprise help you gain control over your desktop environment, quickly adapt to business changes, and run your desktop operations more efficiently.

Citrix Ready appliance components

Figure 8 shows the key components of Citrix Ready appliance.

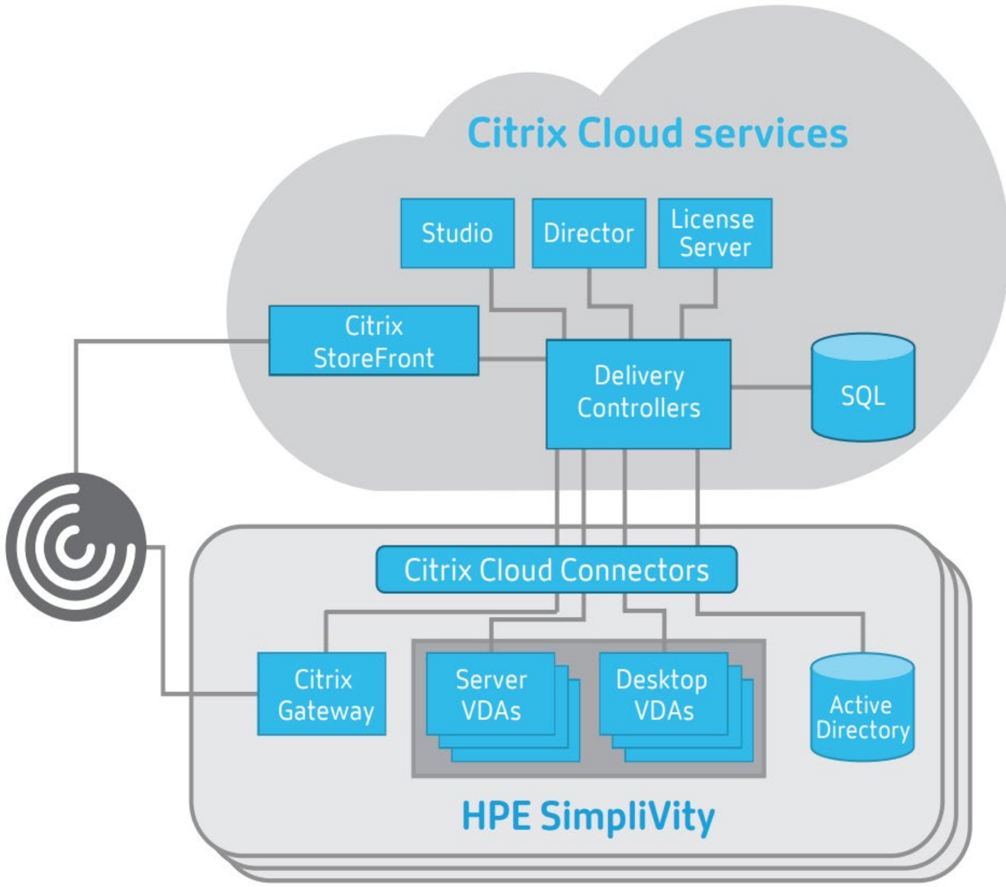

FIGURE 8. Citrix ready converged infrastructure appliance

Citrix Cloud

Citrix Cloud provides the HPE SimpliVity HCI appliance with the Citrix portfolio of workspace technologies as a service through a single management plane. Citrix Cloud keeps all the Citrix infrastructure components in the cloud up-to-date. You don't have to worry about performing complex upgrades to the underlying hardware and software components. Infrastructure software updates are quickly and seamlessly applied across the entire infrastructure stack simultaneously.

Citrix Cloud Connector

Citrix Cloud Connector acts as an interface between the Citrix Cloud and resource location. This helps in managing the hosts and virtual desktops from the cloud by removing the need for complex solutions and networking. Each resource location is recommended to have more than one cloud connector to provide high availability (HA). However, it is laborious to configure or deconfigure Citrix Cloud Connector whenever a resource location is created or deleted.

Figure 9 shows Citrix Cloud plug-in for HPE SimpliVity.

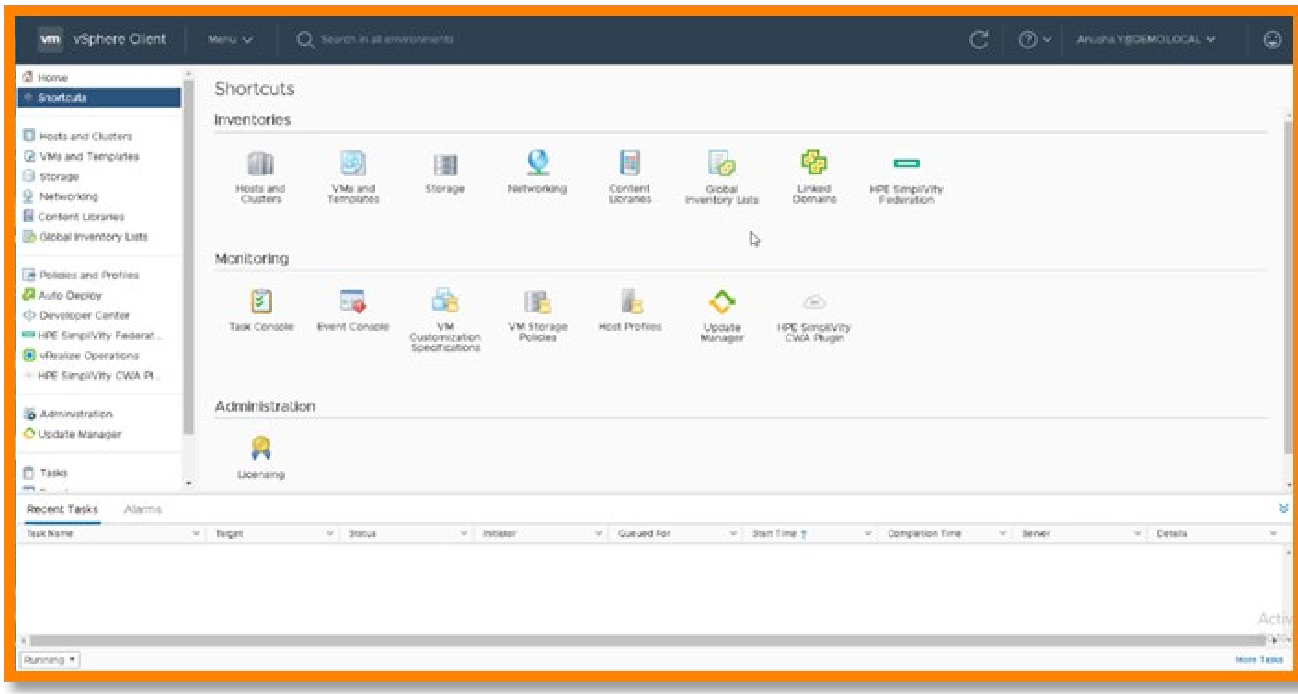

FIGURE 9. HPE SimpliVity Citrix Cloud plug-in

The plug-in can perform different operations to configure and deconfigure Citrix Cloud Connector virtual machines. The Citrix plug-in is supported for VMware vSphere hypervisor. Once the plug-in is installed, it can be accessed by the hypervisor management interface (vCenter).

SOLUTION OVERVIEW

The solution outlined in this document provides guidance for deploying HPE SimpliVity 380 to enable a single client virtualization building block that can support up to 3000 users of mixed workload types from management infrastructure perspective. Through duplicating the building blocks as outlined in the following, the architecture can be used to scale tens of thousands of users.

This solution leverages HPE SimpliVity 380 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 Citrix Virtual Apps and Desktops provide a high-performance client virtualization environment that is highly available and highly scalable.

The building blocks include:

- HPE SimpliVity 380 nodes with Broadwell-based Intel® Xeon® Gold 6150 CPUs and 699 GB usable memory for desktop workloads
- HPE ProLiant DL380 Gen10 compute nodes with Broadwell-based Intel Xeon Gold 6150 CPUs and 768 GB memory for desktop workloads
- HPE SimpliVity 380 nodes with Broadwell-based Intel Xeon Gold 6150 CPUs and 699 GB usable memory for management workloads
- 2 TB datastore per HPE SimpliVity 380 node—this applies to all workload
- 10GbE networking

- Windows 10 Enterprise Long Term Servicing Branch (LTSB) for hosted desktops
- N+1 design for management workloads and infrastructure wherever possible

The tests used in this document was designed to validate the functionality of the system using both a 4&0 block and a 4&4 block. A 4&0 block is defined as four HPE SimpliVity 380 nodes in a single cluster. A 4&4 block is defined as four HPE SimpliVity 380 nodes tied together with four HPE ProLiant DL380 Gen10 compute nodes in a 1:1 configuration, as shown in Figure 10 in a single cluster.

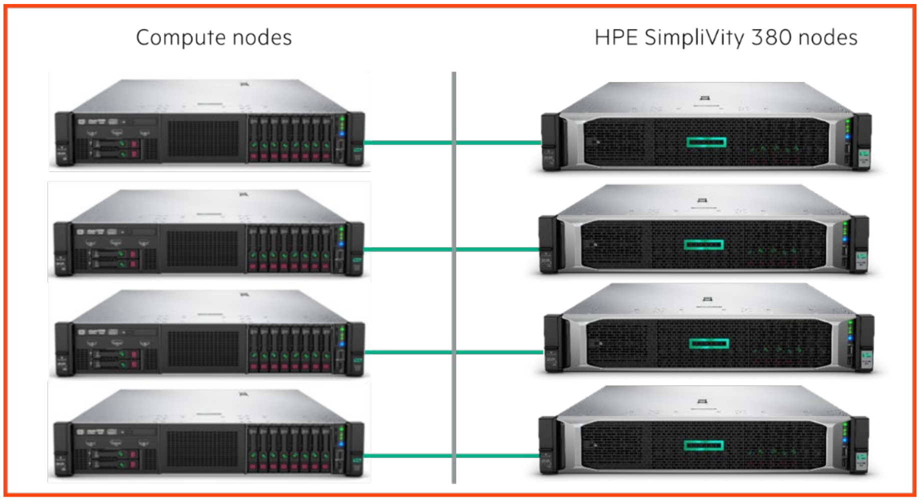

FIGURE 10. Compute nodes combined with HPE SimpliVity nodes in a 4+4 block

Login VSI 4.1.32.1 which was used for testing all workloads, including Login VSI Standard Knowledge Worker for hosted desktops. Proper sizing methodology is outlined in the following section to ensure that the configuration and load described in this document are production ready for customer environment and resilient to N+1 design requirements.

Management infrastructure

This section details the HPE SimpliVity environment dedicated to running the management workloads required to support 3000 user sessions.

- 2x HPE SimpliVity 380 servers
- Intel Xeon Gold 6150 (2.70 GHz, 18-core, 2 sockets per server)
- 699 GB usable memory each server
- 2 x 2 TB datastores
- 10GbE interconnect between systems (no 10GbE switch required, but may be used)

A separate, dedicated, HPE SimpliVity environment is also used for the Virtual Apps and Desktops hosted desktops, further detailed in the Desktop infrastructure section in this document. The management workloads considered in this document are outlined in detail in the Best practices and configuration guidance for the solution section in this document.

Desktop infrastructure

The 4+0 desktop block for 750 Knowledge Worker on MCS-deployed hosted desktops is a 4-node HPE SimpliVity 380, making a 4-node vSphere cluster. A separate 4+4 block (four HPE SimpliVity 380 nodes with four compute nodes) supports 1400 Knowledge Worker sessions on MCS-deployed hosted desktops. 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 in this document.

The following desktop infrastructure was used to support these workload:

- 4x HPE SimpliVity 380 Gen10 Small Enterprise All-Flash servers
  - Intel Xeon Gold 6150 (2.70 GHz, 18-core, 2 sockets per server)
  - 768 GB memory
- 4x HPE SimpliVity 380 Gen10 Small Enterprise All-Flash servers
  - Intel Xeon Gold 6150 (2.70 GHz, 18-core, 2 sockets per server)
  - 768 GB memory
- 4x HPE ProLiant DL380 Gen10 compute nodes
  - Intel Xeon Gold 6150 (2.70 GHz, 18-core, 2 sockets per server)
  - 768 GB Memory
- 699 GB usable memory per HPE SimpliVity 380 system
- 768 GB memory per HPE DL380 compute node
- 4 x 2 TB datastores for each cluster
- 10GbE networking

SOLUTION COMPONENTS

Table 1 provides an overview of the configuration for Citrix Virtual Apps and Desktops 7 1808, Login VSI 4.1.32.1, and the tested HPE SimpliVity 380 plus HPE ProLiant DL380 Gen10 compute node building block.

TABLE 1. Solution components

Parameter	Setting
Users per host/HPE SimpliVity 380 node	
Hosted desktops w/o compute nodes	188 per host
Hosted desktops with compute nodes	175 per host
Machine catalogs deployed	
Hosted desktops w/o compute nodes	1x MCS-deployed catalog w/7500 Windows 10 VMs
Hosted desktops with compute nodes	1x MCS-deployed catalog w/1400 Windows 10 VMs
Machine catalog configuration	
Hosted desktops	Random, user data not saved, 256 MB RAM cache, 10 GB disk cache per VM
Login VSI configuration	
Benchmark mode	Enabled
Hosted desktop workload	Knowledge Worker
Hosted desktop connection	Direct Desktop Connector (DDC)
Number of launchers	6 launchers

BEST PRACTICES AND CONFIGURATION GUIDANCE FOR THE SOLUTION

In this section, configuration guidance and best practices will be shown for deploying Citrix Virtual Apps and Desktops 7 1808 on HPE SimpliVity 380 infrastructure. The guidance for the infrastructure is required for both management and desktop workload.

Management infrastructure

Table 2 and Table 3 show the infrastructure deployed to support the Citrix Virtual Apps and Desktops 7 1808 workload for the testing.

TABLE 2. Management infrastructure workloads

Workload	Version	vCPUs	RAM	Disk	OS
VMware vCenter Server® Appliance™—Desktop (large)	6.5	16	32 GB	640 GB	VMware® Virtual Appliance
vCenter Server Appliance—Mgmt (small)	6.5	4	16 GB	290 GB	VMware Virtual Appliance
Arbiter—Desktop	3.7.6	2	4 GB	40 GB	Windows Server 2016
Arbiter—Mgmt (lives outside of cluster)	3.7.6	2	4 GB	40 GB	Windows Server 2016
Microsoft SQL Server x 2 (Always On Availability Group)	2014 U1	4	8 GB	100 GB	Windows Server 2016
File Share Witness (Always On Availability Group)	N/A	2	4 GB	40 GB	Windows Server 2016
AD DC/DHCP/DNS x 2	N/A	2	4 GB	40 GB	Windows Server 2016
Citrix Virtual Apps and Desktops Controller Server x 2	7 1808	4	8 GB	40 GB	Windows Server 2016
Citrix Virtual Apps and Desktops StoreFront Server x 2	7 1808	4	4 GB	40 GB	Windows Server 2016
Citrix Virtual Apps and Desktops Licensing Server	7 1808	4	4 GB	40 GB	Windows Server 2016

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 an HPE SimpliVity 380 Small Enterprise All-Flash.
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 data centers	1	A single vSphere cluster is present, and no data center-level separation is necessary.
VMware vSphere® High Availability configuration	HA enabled Admission control enabled % of cluster resources reserved: 50% Isolation response: Leave powered on	Enabled to restart VMs in the event of an VMware 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: 8507 MB das.vmcpuminmhz: 1000 MHz	Both are set to averages of the workload in the cluster. This serves to set the percentage of cluster resources in HA calculation to that of an average VM.
vSphere version	6.5.0	Latest supported release

HPE SimpliVity servers

To support the management workload outlined in this document, a 2-host vSphere cluster, comprising a pair of HPE SimpliVity 380 nodes, which is recommended. Unlike other hyperconverged infrastructure solutions, HPE SimpliVity fully supports a 2-host cluster in its minimum configuration. Using HPE SimpliVity allows you to start small, with only the infrastructure you need, and then scale out as your VDI environment grows.

vCenter Servers

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

Virtual Apps and Desktops Delivery Controllers

A single Delivery Controller can support up to 5000 users. To ensure high availability, two Delivery Controllers were deployed in an N+1 configuration.

Virtual Apps and Desktops StoreFront Servers

A single StoreFront Server can support up to 10,000 users. To ensure high availability, two StoreFront Servers were deployed in an N+1 configuration.

Virtual Apps and Desktops License Server

Only a single License Server is required.

Infrastructure Services (domain controllers/DNS/DHCP)

These services are co-located on the same virtual machines and no CPU or memory pressure was observed during testing. Active Directory design and recommendations are outside the scope of this document. 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 was 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 4.

TABLE 4. Required Microsoft SQL databases

Database	Authentication	Size	Recovery mode
Virtual Apps and Desktops Delivery Controller DBs	Windows authentication	Default	Full

Sizing: Compute, storage, and network resources for each infrastructure VM are selected using Citrix best practices as a baseline and modified based on their observed performance on the HPE SimpliVity 380 systems.

HPE SimpliVity Arbiter placement

The HPE SimpliVity Arbiter should always be deployed outside the HPE SimpliVity infrastructure that it manages. The arbiter instance that supports the desktop VCSA is deployed as a virtual machine in the management cluster. The arbiter instance supporting the management VCSA should be deployed outside the management cluster. In this instance, it is deployed as a small virtual machine outside this configuration.

VMware vSphere Storage APIs - 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 plug-in for HPE SimpliVity is installed during deployment, so no manual intervention is required.

Datastores

A single datastore per HPE SimpliVity 380 server is required to ensure even storage distribution across cluster members. This is less important in a 2-node HPE SimpliVity server configuration. However, if customers had to grow the environment, they should follow the best practice to grow from 2 node to 3 or more. Additionally, when provisioning an MCS pool, an equal number of desktops should be deployed to each datastore. This is accomplished by enabling only one datastore at a time and deploying an even number of desktops to that datastore. Proceed in the same manner for the remaining datastores, again enabling only one datastore as a storage resource at a time. This best practice has been shown 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.

Networking

The following best practices are utilized in the vSphere networking design:

- HPE OVC networking is segregated from ESXi host and virtual machine network traffic
- 10GbE is used wherever possible for HPE 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. For the remaining traffic (virtual machines, HPE SimpliVity Federation, HPE SimpliVity storage, VMware vSphere vMotion®), a single VMware vSphere Distributed Switch™ is deployed. The switch configurations are shown in Table 5 and Table 6.

TABLE 5. 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

TABLE 6. 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 (more importantly, when configuring an MTU size ensure consistency of the MTU size between the desktops and infrastructure network segments to prevent costly packet fragmentation)
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

Desktop infrastructure

The following configurations was utilized for the testing. Table 7 shows the VM configuration used for the Knowledge Worker VM when running with a stand-alone HPE SimpliVity deployment. Table 8 shows the corresponding configuration when running in a deployment with HPE SimpliVity combined with compute nodes.

Figure 10 shows Machine Creation Services desktop configuration.

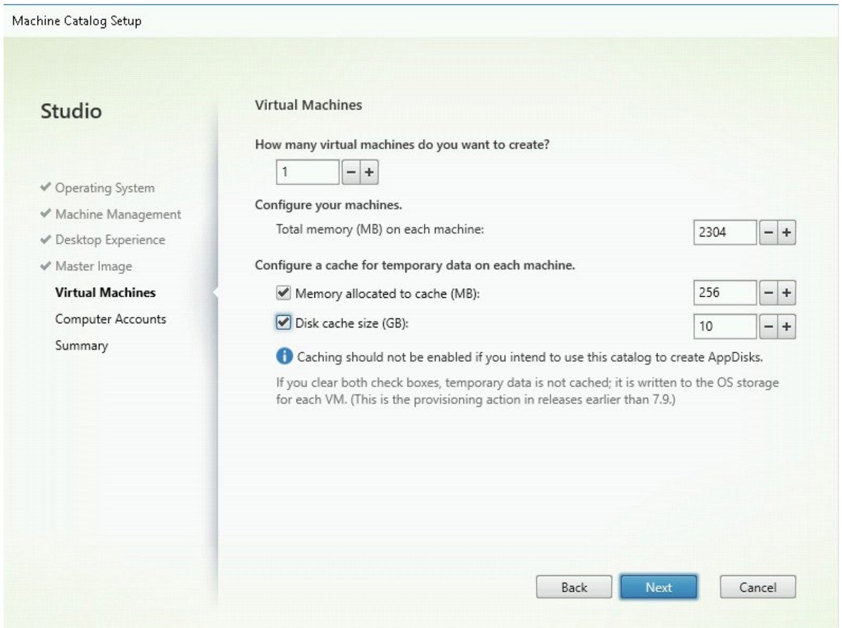

FIGURE 10. MCS desktop configuration

TABLE 7. Knowledge Worker VM configuration—hosted desktops without compute nodes

Attribute	Specification
Operating system	Windows 10 LTSB 64-bit
Virtual hardware	VM virtual hardware version 13
VMware Tools™	10.1.5
Number of vCPUs	2
Memory, including MCS RAM cache	2304 MB
MCS RAM cache size	256 MB
MCS Disk cache size	10 GB
Virtual disk: VMDK	40 GB
NTFS cluster alignment	8 KB
SCSI controller	VMware Paravirtual
Virtual floppy drive	Removed
Virtual CD/DVD drive	Removed
NIC vendor and model Number of ports/NIC x speed	VMXNET3 1x 10 Gigabit Ethernet
OS page file	1.5 GB starting and max.
Number deployed	750

TABLE 8. Knowledge Worker VM configuration—hosted desktops with compute nodes

Attribute	Specification
Operating system	Windows 10 LTSB 64-bit
Virtual hardware	VM virtual hardware version 13
VMware Tools	10.1.5
Number of vCPUs	2
Memory, including MCS RAM cache	2304 MB
MCS RAM cache size	256 MB
Virtual disk: VMDK	40 GB
NTFS cluster alignment	4 KB
SCSI controller	VMware Paravirtual
Virtual floppy drive	Removed
Virtual CD/DVD drive	Removed
NIC vendor and model Number of ports/NIC x speed	VMXNET3 1x 10 Gigabit Ethernet
OS page file	1.5 GB starting and max.
Number deployed	1400

VMware vSphere Storage APIs - 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 plug-in for HPE SimpliVity is installed during deployment, so no manual intervention is required.

Datastores

You must deploy a datastore for each HPE SimpliVity 380 system in each vSphere cluster. In this configuration, four datastores were created for each vSphere cluster. This helps to evenly distribute storage load across the HPE SimpliVity systems in the vSphere cluster, as well as increase the likelihood of any given desktop that has locality with its VMDK disk.

Each datastore contains a virtual machine template and write cache files for every virtual machine. The write cache file contains all disk writes of a target device when using a write-protected vDisk (Standard Image).

HPE OVC configuration

For the desktop-supporting HPE SimpliVity 380 nodes, HPE OVC should be tuned to use six virtual CPUs instead of the standard configuration of four virtual CPUs for the HPE SimpliVity 380 Small Enterprise All-Flash model. This increases the numbers of VMs supported per HPE SimpliVity 380 node and is required to support the desired number of compute nodes.

Compute nodes

This design calls for the use of compute nodes, which are vSphere servers with no local HPE SimpliVity 380 components. The compute nodes in this design are attached 1:1 with HPE SimpliVity nodes within the vSphere cluster. Therefore, each HPE SimpliVity node in the vSphere cluster will have a single HPE ProLiant DL380 compute node attached via NFS mount. This requires the mapping of a VMkernel port per compute node to an individual HPE OVC in order to mount its storage. No single HPE OVC should have more than one compute node mapped. In the tested configuration, the compute node to HPE OVC mapping was done as shown in Table 9.

TABLE 9. vSphere VMkernel port to HPE OVC storage IP mapping

Compute node (VMkernel IP)	HPE SimpliVity 380 (VMkernel IP)	HPE OVC mapped (storage IP)
HPE DL380 Gen10 #1 (10.111.33.121/22)	HPE SimpliVity 380 #1 (10.111.33.111/22)	HPE OVC #1 (10.111.32.111/22)
HPE DL380 Gen10 #2 (10.111.33.122/22)	HPE SimpliVity 380 #2 (10.111.33.112/22)	HPE OVC #2 (10.111.32.112/22)
HPE DL380 Gen10 #3 (10.111.33.123/22)	HPE SimpliVity 380 #3 (10.111.33.113/22)	HPE OVC #3 (10.111.32.113/22)
HPE DL380 Gen10 #4 (10.111.33.124/22)	HPE SimpliVity 380 #4 (10.111.33.114/22)	HPE OVC #4 (10.111.32.114/22)

Networking

The following best practices are utilized in the vSphere networking design:

- HPE OVC networking is segregated from ESXi host and virtual machine network traffic
- 10GbE is used wherever possible, for HPE 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. For the remaining traffic (virtual machines, HPE SimpliVity Federation, HPE SimpliVity storage, vMotion), a single vSphere Distributed Switch is deployed. The switch configurations are shown in Table 10 and Table 11.

TABLE 10. Desktop 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

TABLE 11. Desktop 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
Number of ports	4096
Number of uplinks	2
Network adapters	10GbE NICs on each host
Network I/O control	Disabled

Parameter	Setting
VMkernel ports/VM networks	vmk1 – vMotion vmk2 – Storage vMotion – Active/Standby – MTU 9000 Federation – Standby/Active – MTU 9000
VMkernel ports/VM networks	STORAGE – STANDBY/Active – MTU 9000 Desktop 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 the users working environment from their client device is huge, if restoration of that decoupled environment is not planned to the latter. Without planning, there exists a risk of removing a user 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. They are:

- An active/active configuration across multiple sites leveraging non-persistent desktops
- An active/active configuration across multiple sites with a tertiary site for regional DR
- An active/active configuration across multiple sites leveraging MCS 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. This is not an all-inclusive list. These scenarios and the designs detailed therein are flexible and can be adjusted to meet most of the failure and recovery scenario.

Scenario #1 – Active/Active Site Recovery and Failback with non-persistent Desktops

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

- Desktop vCenter Server
- Citrix Virtual Apps and Desktops Controller Server
- Citrix Virtual Apps and Desktops StoreFront 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 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.

Citrix non-persistent desktop 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.

Citrix user profile manager is configured on each site 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). The following steps explains the failure scenario for Site Recovery and Failback.

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 can 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 the following 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 cases where the DNS A and PTR records are 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 the recovery of the file server and re-IP are complete. This is 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 Citrix pods themselves (vCenter Server + Controller + Storefront) should be there within a site. Given the pod and block architecture, there is no reason to fail over of 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 Citrix 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 11 shows an active/active configuration across multiple sites leveraging non-persistent desktops.

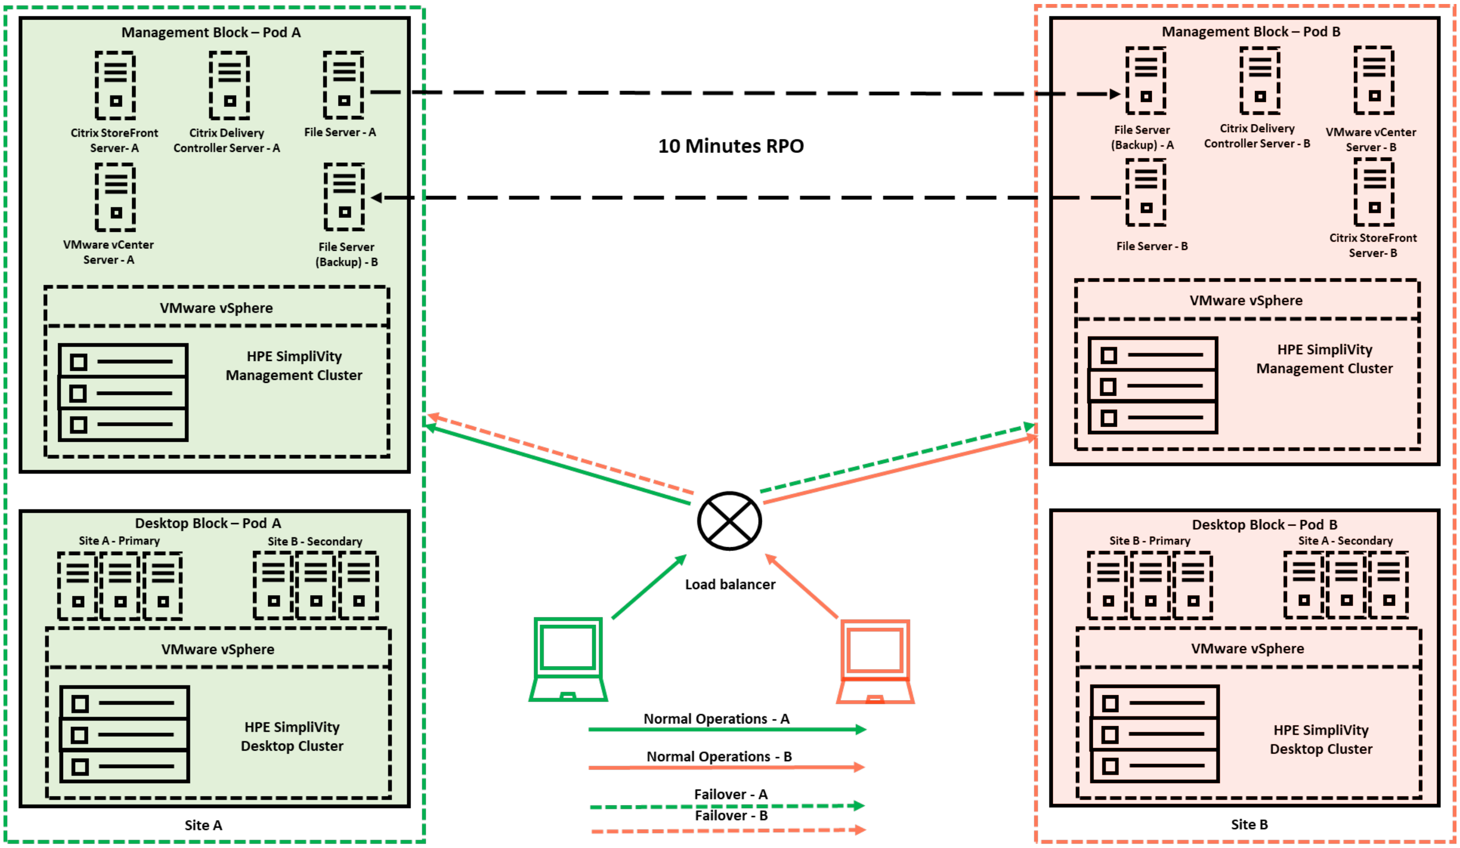

FIGURE 11. Active/Active Site Recovery with Linked Clones

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 Citrix pods. These include, per site:

- Desktop vCenter Server
- Citrix Virtual Apps and Desktops Controller Server
- Citrix Virtual Apps and Desktops StoreFront 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, are 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).

Citrix non-persistent Desktop 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). The following steps explains the failure scenario for Site Failure – Disaster recovery with a tertiary site.

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

Figure 12 shows an active/active configuration across multiple sites with a tertiary site for regional DR.

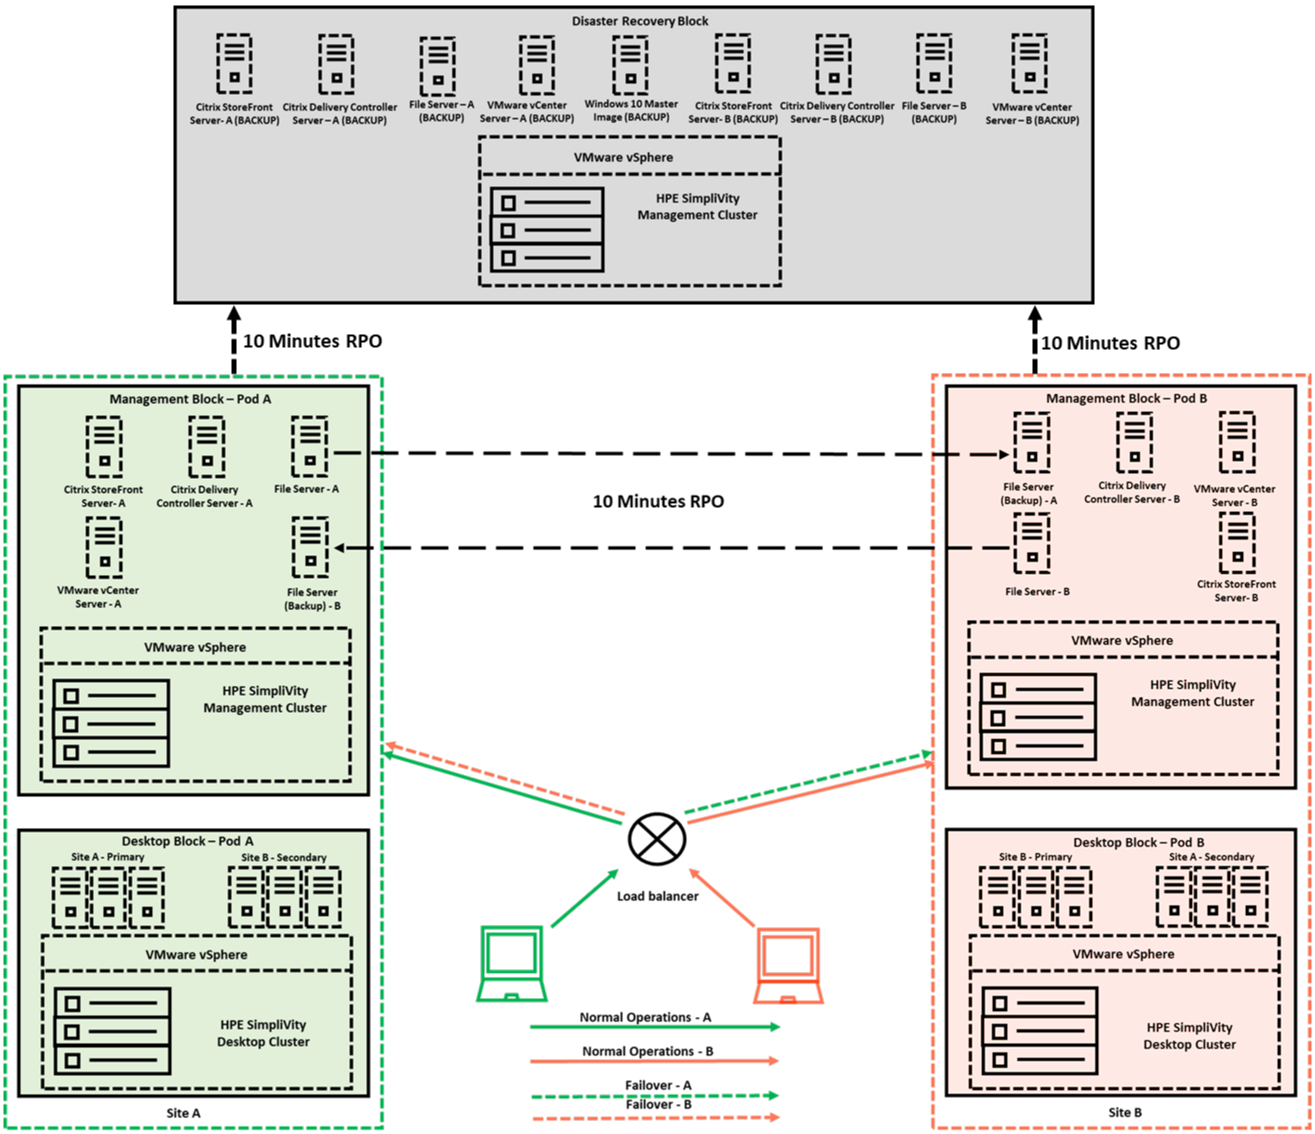

FIGURE 12. Active/Active Site Failure – DR with a Tertiary Site

Best practices

The requirements to fulfill this scenario for DR to and from a tertiary site are typically minimal if given the value. In general, 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 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 (MCS full clone desktops)

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

- Desktop vCenter Server
- Citrix Virtual Apps and Desktops Controller Server
- Citrix Virtual Apps and Desktops StoreFront Server
- 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 MCS full clone desktops workload within its management domain but in the opposite data center. This allows for failover of MCS full clone desktops to the opposite data center while maintaining the relationship between MCS full clone desktops and its vCenter Server and Citrix 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. Citrix MCS full clone desktops 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 MCS full clone desktops. User data and applications are unique per desktop. The following steps explains the failure scenario for Site Recovery and Failback.

1. Desktops and management servers backed up across sites.
2. Site A fails.
3. Recovery of management workloads at site B.
4. Recovery of MCS full clone desktops at site B.
5. Site A users can 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 MCS 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 13 shows an active/active configuration across multiple sites leveraging MCS full clone desktops.

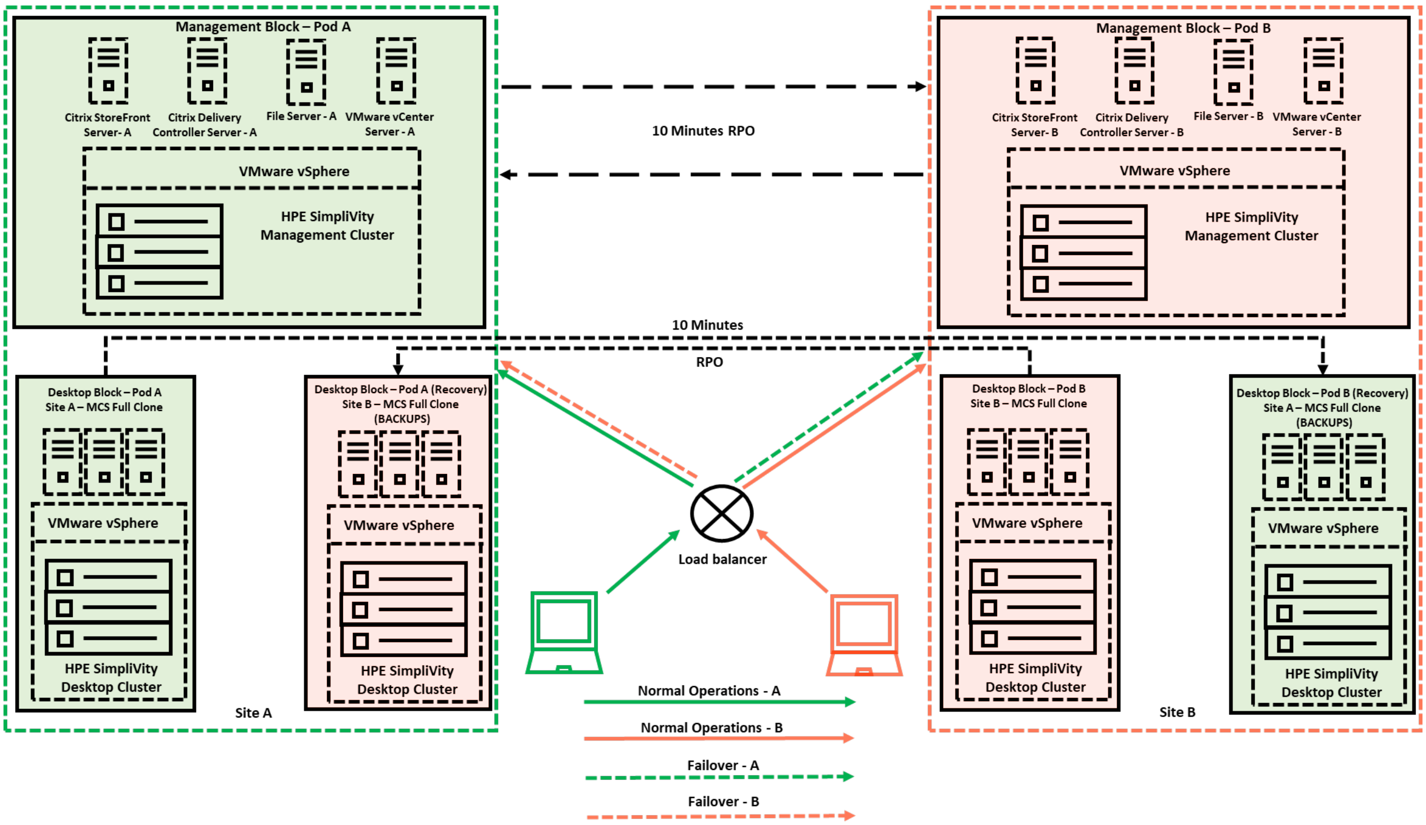

FIGURE 13. Active/Active Site Recovery and Failback (Full Clones)

Best Practices

To fail over the management workloads from one physical site to another, one of the below 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.

MCS full clone desktops are 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 MCS full clone desktops stay within their respective management domains (vCenter Server + Citrix Controller pod), no re-registration of desktops to another pod or failing back to the original pod is necessary, significantly it reduces 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. 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

This section discusses the sizing requirements for running the standard Login VSI workload sizes. Details of the test workloads are provided, while test results are analyzed to provide recommendations for optimizing performance.

HPE SimpliVity Federation and vSphere cluster sizing

The solution is architected such that workloads are split into multiple vSphere clusters, with the 750 hosted desktop (performance-critical) workloads in one vSphere cluster and the 1400 hosted desktop (balanced economy with compute nodes) workloads in the other.

To support multiple vSphere clusters in a single vCenter Server, both clusters must belong to a single HPE SimpliVity Federation. A vCenter Server supports a single HPE SimpliVity Federation. For Citrix desktop workload supporting vSphere infrastructure, vCenter Server Linked Mode is not supported by Citrix. An overview of the vSphere setup is shown in Table 12.

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 accurately.

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 an HPE SimpliVity 380 Small Enterprise All-Flash.
Number of vSphere clusters	2	The hosted desktop and hosted desktop with compute nodes workloads are split into separate vSphere clusters.
Number of vSphere data centers	1	With HPE OmniStack 3.7.6, the fault domain for an HPE SimpliVity Cluster has been moved to the vSphere cluster level and multiple vSphere clusters per data center are supported.
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: 12%	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: 2304 MB das.vmcpuminmhz: 187 MHz	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.
vSphere version	6.5.0	Latest supported release (n-1)

HPE SimpliVity servers

Two vSphere clusters consist of a 4+0 (four HPE SimpliVity 380 systems) to support the Knowledge Worker hosted desktop and a 4+4 (four HPE SimpliVity nodes plus four compute nodes) to support the other Knowledge Worker hosted desktop workload. The following design patterns were observed:

- **Limit physical CPU to virtual CPU oversubscription:** In this configuration, each HPE SimpliVity 380 system has 36 physical cores. HPE OVC takes six physical cores (when attached with compute nodes), leaving 30 per system for desktop workloads to use. The HPE ProLiant DL380 compute nodes have 36 usable physical cores each. Each hosted desktop VM has two vCPUs.
- **Do not overcommit memory:** In this configuration, each HPE SimpliVity 380 system has 768 GB of available physical memory whereas 69 GB of memory is reserved for HPE OVC on each system. It leaves 699 GB available per HPE SimpliVity 380 system for desktop workloads. The HPE ProLiant DL380 compute nodes have 768 GB of usable memory each. Each hosted desktop VM has 2.25 GB of memory.
- Ensure the server hardware is running the latest HPE SimpliVity approved BIOS settings (for example, High Performance mode) and firmware.

Table 13 shows the desktop infrastructure requirements.

TABLE 13. Desktop infrastructure resource requirements

Node attribute	HPE SimpliVity 380	HPE DL380 compute node	Total usable	Required	Ratio/overage
Hosted desktop CPU	30 x 4 = 120 pCPUs	N/A	120 pCPUs	1500 vCPUs	12.5:1 (16.7:1 failover)
Hosted desktop w/compute node CPU	30 x 4 = 120 pCPUs	36 x 4 = 144 pCPUs	264 pCPUs	2800 vCPUs	10.6:1 (12.3:1 failover)
Hosted desktop RAM	699 GB x 4 = 2796 GB	N/A	2796 GB	1688 GB	1108 GB spare (409 GB failover)
Hosted desktop w/compute node RAM	699 GB x 4 = 2796 GB	768 GB x 4 = 3072 GB	5868 GB	3150 GB	2718 GB spare (1950 GB failover)

Workload description

Login VSI has standard workloads to use in benchmark testing, including the Knowledge Worker workload. The percentage values for CPU usage, disk reads, and disk writes are relative to the Knowledge Worker workload. These standard per-user workloads, as defined by Login VSI, are shown in Table 14 and Table 15.

TABLE 14. Office and Knowledge Worker

Parameter	Office Worker setting (for reference)	Knowledge Worker setting
Apps open	5–8	5–9
CPU usage	82%	100%
Disk reads	90%	100%
Disk writes	101%	100%
IOPS	8.1	8.5
Memory	1.5 GB	1.5 GB
CPU	1 vCPU	2 vCPU

TABLE 15. Applications

Applications installed
Microsoft Word 2016
Microsoft Excel 2016
Microsoft Outlook 2016
Microsoft PowerPoint 2016
Internet Explorer 11
MindMap
Adobe® Flash Player
Doro PDF Printer
Photo Viewer

All performance testing that is documented here utilizes the Login VSI (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, which are being performed within a desktop workload in a scripted loop. The baseline is the measurement of the response time for specific operations performed in the desktop workload, measured in milliseconds (ms). The desktops in all the following tests are rebooted and ready to service users. Desktop provisioning, customization, and boot time are not measured in these results.

Login VSI has independently validated all performance results published in this document. These results are reviewed for accuracy and to ensure that they meet the strict performance requirements. Login VSI has a set of requirements to be validated by their program. Figure 14 shows the validation by Login VSI.

FIGURE 14. Validated by Login VSI

There are two values that are important to note:

- **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.

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

Analysis and recommendations

Table 16 summarizes the test results.

TABLE 16. Summary of Login VSI test results

Test case	VSIbase	VSImax average	VSImax reached?
750 Knowledge Workers: Hosted	596 ms	1134 ms	No
1400 Knowledge Workers: Hosted	617 ms	1114 ms	No

Hosted desktops without compute nodes: 750 Knowledge Worker users

VSIbase for the environment was 596 ms, and VSImax has not reached in any run, indicating that the system consistently delivered a high-performance user experience with plenty of performance headroom left on the system. VSImax average was 1134 ms, and VSImax threshold was 1597 ms. Latency is consistently low, confirming that HPE SimpliVity infrastructure could easily handle the 750 Knowledge Worker users as shown in Figure 15.

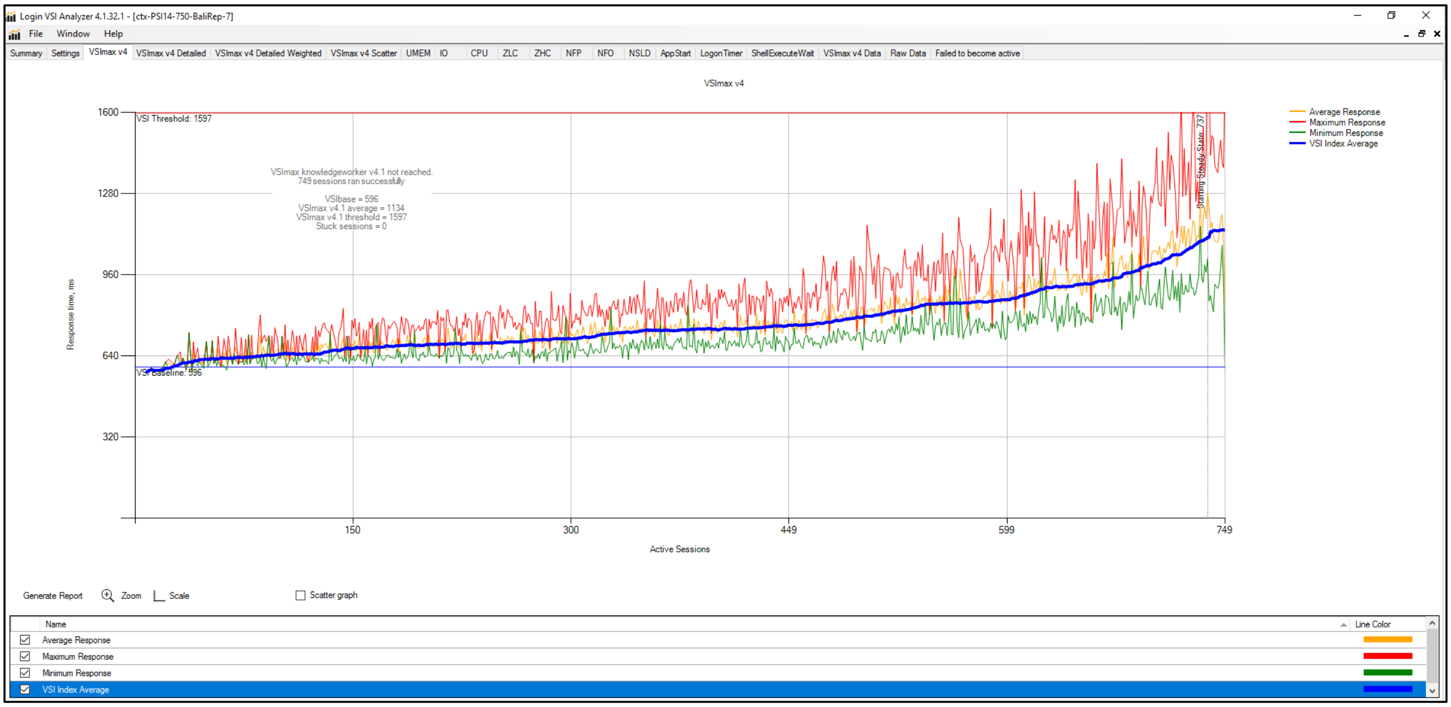

FIGURE 15. 750 Knowledge Worker results

Hosted desktops with compute nodes: 1400 Knowledge Worker users

VSIbase for the environment was 617 ms, and VSImax has not reached in any run, indicating that the system consistently delivered a high-performance user experience with plenty of performance headroom left on the system. VSImax average was 1114 ms, and VSImax threshold was 1618 ms. Latency is consistently low, confirming that the HPE SimpliVity 380 infrastructure could easily handle the 1400 Knowledge Worker users in a hosted desktop configuration with Windows 10, as shown in Figure 16. Most importantly, this test was performed using only four nodes of storage resources.

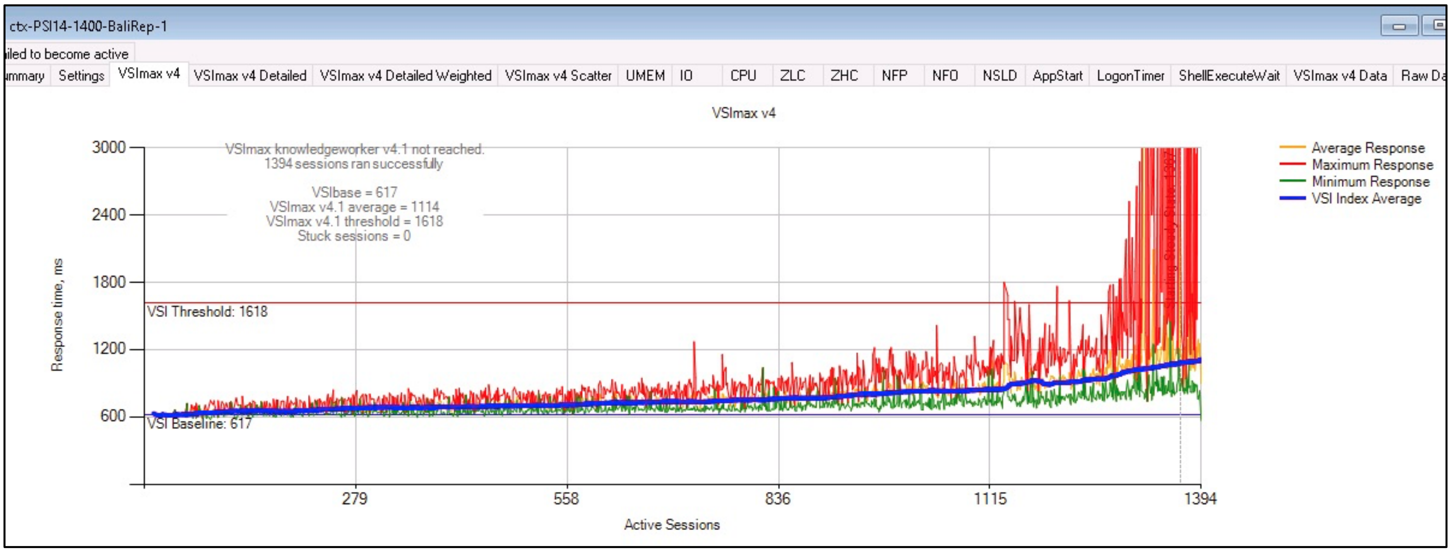

FIGURE 16. 1400 Knowledge Worker results

SUMMARY

This Reference Architecture provides guidance to organizations implementing Citrix Virtual Apps and Desktops 7 1808 on HPE SimpliVity 380 Gen10 hyperconverged infrastructure. It describes the tests performed by Hewlett Packard Enterprise to validate and measure the operation and performance of the recommended solution. These tests included third-party validated performance testing from Login VSI, the industry-standard benchmarking tool for virtualized workload.

Organizations are looking to client virtualization solutions like Citrix Virtual Apps and Desktops to reduce software licensing, distribution, and administration expenses, and to improve security and compliance. The HPE SimpliVity 380 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 for client virtualization provides:

- Simplified deployment with hyperconverged infrastructure providing incremental building blocks
- Ability to start small and scale out in affordable increments from pilot to production
- Independently validated client virtualization performance
- Deployment of full-clone desktops with the same data efficiency as linked clones
- Enterprise-class data protection and resiliency

This solution testing was performed in December 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. Hewlett Packard Enterprise 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 a Hewlett Packard Enterprise 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 Hewlett Packard Enterprise Reseller or Hewlett Packard Enterprise Sales Representative for more details. 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 0D1	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 0D1	Factory Integrated
1	P01366-B21	HPE 96W Smart Storage Battery (up to 20 Devices) with 145mm Cable Kit
1	P01366-B21 0D1	Factory Integrated
1	804331-B21	HPE Smart Array P408i-a SR Gen10 (8 Internal Lanes/2GB Cache) 12G SAS Modular Controller
1	804331-B21 0D1	Factory Integrated
1	700751-B21	HPE FlexFabric 10Gb 2-port 534FLR-SFP+ Adapter
1	700751-B21 0D1	Factory Integrated
2	830272-B21	HPE 1600W Flex Slot Platinum Hot Plug Low Halogen Power Supply Kit
2	830272-B21 0D1	Factory Integrated
1	BD505A	HPE iLO Advanced 1-server License with 3yr Support on iLO Licensed Features
1	BD505A 0D1	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 0D1	Factory Integrated
1	867809-B21	HPE Gen10 2U Bezel Kit
1	867809-B21 0D1	Factory Integrated
1	826703-B21	HPE DL380 Gen10 SFF Systems Insight Display Kit
1	826703-B21 0D1	Factory Integrated
1	733660-B21	HPE 2U Small Form Factor Easy Install Rail Kit
1	733660-B21 0D1	Factory Integrated
1	H1K92A3	HPE 3Y Proactive Care 24x7 SVC
1	H1K92A3 R2M	HPE iLO Advanced Non-Blade - 3yr Support

Qty	Product #	Product Description
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

HPE Servers, hpe.com/servers

HPE Storage, hpe.com/storage

HPE Networking, hpe.com/networking

HPE Technology Consulting Services, hpe.com/us/en/services/consulting.html

To help us improve our documents, please provide feedback at hpe.com/contact/feedback.

LEARN MORE AT

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

Hewlett Packard Enterprise

a50001114enw, version 2.0, March 2025