Hitachi Infrastructure for Citrix Workspace with Hitachi Virtual Storage Platform E1090H

Reference Architecture Guide

MK-SL-252-00
September 2022

© 2022 Hitachi Vantara LLC. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including copying and recording, or stored in a database or retrieval system for commercial purposes without the express written permission of Hitachi, Ltd., or Hitachi Vantara LLC (collectively “Hitachi”). Licensee may make copies of the Materials provided that any such copy is: (i) created as an essential step in utilization of the Software as licensed and is used in no other manner; or (ii) used for archival purposes. Licensee may not make any other copies of the Materials. “Materials” mean text, data, photographs, graphics, audio, video and documents.

Hitachi reserves the right to make changes to this Material at any time without notice and assumes no responsibility for its use. The Materials contain the most current information available at the time of publication.

Some of the features described in the Materials might not be currently available. Refer to the most recent product announcement for information about feature and product availability, or contact Hitachi Vantara LLC at https://support.hitachivantara.com/en_us/contact-us.html.

Notice: Hitachi products and services can be ordered only under the terms and conditions of the applicable Hitachi agreements. The use of Hitachi products is governed by the terms of your agreements with Hitachi Vantara LLC.

By using this software, you agree that you are responsible for:

1. Acquiring the relevant consents as may be required under local privacy laws or otherwise from authorized employees and other individuals; and
2. Verifying that your data continues to be held, retrieved, deleted, or otherwise processed in accordance with relevant laws.

Notice on Export Controls. The technical data and technology inherent in this Document may be subject to U.S. export control laws, including the U.S. Export Administration Act and its associated regulations, and may be subject to export or import regulations in other countries. Reader agrees to comply strictly with all such regulations and acknowledges that Reader has the responsibility to obtain licenses to export, re-export, or import the Document and any Compliant Products.

Hitachi and Lumada are trademarks or registered trademarks of Hitachi, Ltd., in the United States and other countries.

AIX, AS/400e, DB2, Domino, DS6000, DS8000, Enterprise Storage Server, eServer, FICON, FlashCopy, GDPS, HyperSwap, IBM, Lotus, MVS, OS/390, PowerHA, PowerPC, RS/6000, S/390, System z9, System z10, Tivoli, z/OS, z9, z10, z13, z14, z/VM, and z/VSE are registered trademarks or trademarks of International Business Machines Corporation.

Active Directory, ActiveX, Bing, Excel, Hyper-V, Internet Explorer, the Internet Explorer logo, Microsoft, Microsoft Edge, the Microsoft corporate logo, the Microsoft Edge logo, MS-DOS, Outlook, PowerPoint, SharePoint, Silverlight, SmartScreen, SQL Server, Visual Basic, Visual C++, Visual Studio, Windows, the Windows logo, Windows Azure, Windows PowerShell, Windows Server, the Windows start button, and Windows Vista are registered trademarks or trademarks of Microsoft Corporation. Microsoft product screen shots are reprinted with permission from Microsoft Corporation.

All other trademarks, service marks, and company names in this document or website are properties of their respective owners.

Copyright and license information for third-party and open source software used in Hitachi Vantara products can be found in the product documentation, at https://www.hitachivantara.com/en-us/company/legal.html or https://knowledge.hitachivantara.com/Documents/Open_Source_Software.

Feedback

Hitachi Vantara welcomes your feedback. Please share your thoughts by sending an email message to SolutionLab@HitachiVantara.com. To assist the routing of this message, use the paper number in the subject and the title of this white paper in the text.

Revision history

Changes	Date
Initial release	September 20, 2022

Reference Architecture Guide

Why Virtual Desktop Infrastructure

With the increased demand for work-from-home employees in recent times, and the increased number of cyber attacks, such as ransomware, large enterprise companies have had to rethink their end-user desktop strategies. To address the seasonal burst of temporary contractors, or for short-term large demand spikes, corporations have embraced the use of a virtual desktop infrastructure (VDI) and application virtualization technologies.

For the shorter term, Desktop-as-a-Service (DaaS) offerings in the public cloud have been a quick and easy way to offer the remote workforce access to such digital workspaces. However, for longer term approaches, and to address the need of workspace mobility for all of their employees and not just seasonal workers, these companies also need a virtual app and a Virtual Desktop Infrastructure (VDI) on-premises strategy. Running VDI on-prem gives corporations better cost controls, better management over data privacy and security compliance, as well as better data proximity over in-house homebrewed applications with three-tier application servers hosted in the same datacenter as the VDI server hosts.

What we offer

The Hitachi Vantara solution portfolio of VDI appliances can address both small and large-scale deployments. In all cases, Hitachi Vantara has designed turnkey appliances that are aimed at accelerating the implementation of VDI in customer datacenters, having fully tested and validated the architecture for resiliency, scalability, availability, security, and performance, and shipping these appliances pre-cabled, pre-racked, and ready to be deployed by the customer or by our services organization.

This document describes the reference architecture for the implementation of Hitachi Infrastructure for Citrix Workspace. It shows the benefits of deploying Citrix Virtual Desktop Infrastructure (VDI) technology on a validated architecture built with Hitachi converged infrastructure leveraging Hitachi Advanced Server G2, Hitachi Virtual Storage Platform E1090H (VSP E1090H) with SSDs, and Cisco switches and management technologies.

Extensive testing evaluated the performance and capabilities of this integrated solution. Test results show that this Hitachi Infrastructure for Citrix VDI solution supports high virtual-desktop density (desktop sessions per host), and additional servers and storage can scale with near-linear performance.

Executive summary

The following are the key benefits of Citrix Workspace with HA810 G2 on E1090H storage:

- A flexible VDI workspace to improve performance, enhance simplicity, and reduce costs.
- A scalable remote user VDI solution for general-purpose, tier-1, and test/dev workloads.
- Maximized investments in storage performance and capabilities.
- Seamless mobility between server and storage device to provide the most optimal infrastructure for the workload that delivers on performance, availability, and costs.
- Optimized performance with no over-provisioning costs and a highly configurable and integrated infrastructure.

Contents of this reference architecture

This paper includes the results of a series of tests generated with Login VSI, the industry standard benchmarking tool for VDI workloads. For these tests, representative workload profiles such as Knowledge and Power Worker user workloads were selected and tested with hosted virtual desktops (HVD) and hosted shared desktops (HSD). See the following for detailed explanations of solution components and test results:

Note: These practices were developed in a lab environment. Many things affect production environments beyond prediction or duplication in a lab environment. Follow recommended practice by conducting proof-of-concept testing for acceptable results before implementing this solution in your production environment. Test the implementation in a non-production, isolated test environment that otherwise matches your production environment.

Solution overview

The Hitachi Infrastructure for Citrix Workspace solution combines VSP E1090H All-SSD storage, Hitachi Advanced Server HA810 G2 rack servers, Brocade Fibre Channel switches, Cisco Nexus switches, Citrix Hypervisor, and Citrix Virtual Apps and Desktops.

This is a solution designed and validated with compute, network, and storage best practices to provide a scalable, building-block style approach to deploying an infrastructure for VDI. The following figure shows the key elements of this solution.

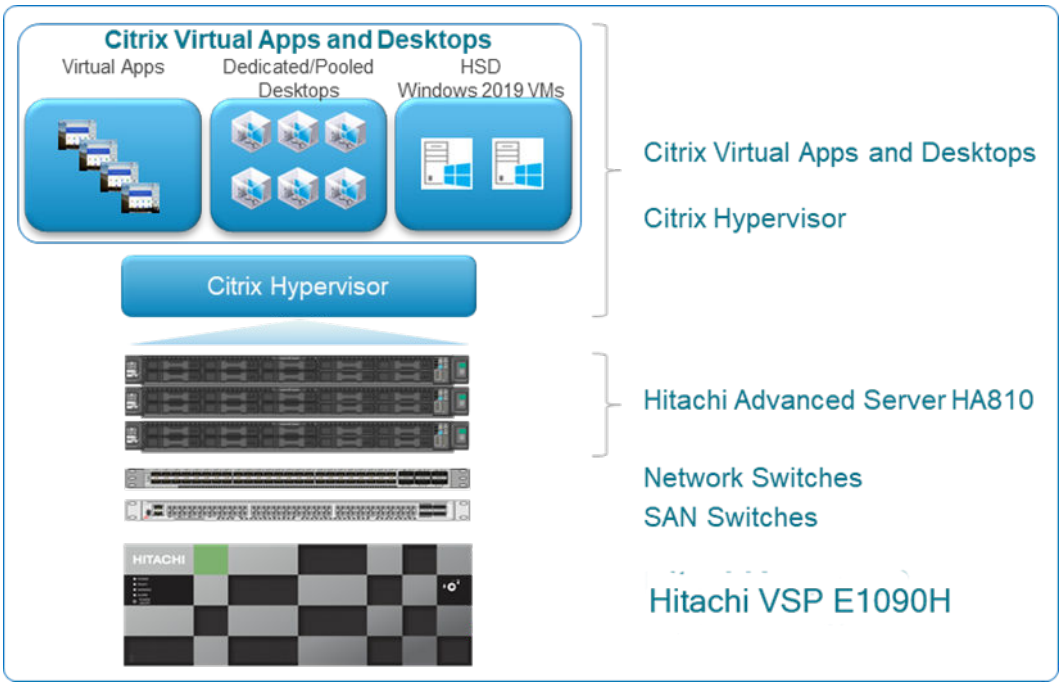

Solution components

This section describes the hardware and software components used to validate this reference architecture. For detailed component information, see Product descriptions (on page 36).

Hardware components

The following table lists the details of the hardware configuration used in a Hitachi converged infrastructure to run the different test cases for this reference architecture.

Vendor	Hardware	Detail Description	Version	Quantity
Hitachi Vantara	HA810G2	▪ VDI Management: HA810 G2 Server Specifications: • 2 × Intel ® Xeon ® Gold 5318Y CPUs (20c, 2.1GHz, 125W) • Memory : 384 GB DDR4 • 1 × 10/25 GbE dual-port SFP28 • 1 × 1GbE RJ45 port for OOB (iLO) • 1 × Smart Array P816i-a SR Controller ▪ VDI User Workloads: HA810 G2 Server Specifications: • 6 × Intel ® Xeon ® Gold 6348 CPUs (24c, 2.6GHz, 205W) • Memory: 1024 GB DDR4 • 1 × 10/25GbE dual-port SFP28 • 1 × 1GbE RJ45 port for OOB (iLO) • 1× Smart Array P408i-a SR controller • 1 × 32 Gb HBA	iLO:5	2 for Management 4 for Workloads

Vendor	Hardware	Detail Description	Version	Quantity
Hitachi Vantara	VSP E1090H	▪ VSP E1090H • SSDs: 14 × 4 TB • RAID 6: 12D+2P+1S • 32 Gbps 4-port CHB: 2 pairs • CTL: 1 pair • DBN: 1 • Cache: 1024 GB		
Cisco	Cisco Nexus	▪ Nexus 93180YC-FX • 48 × 10/25 GbE fiber ports • 6 × 40/100 Gbps QSFP28 ports	NXOS 9.3.3	2
		▪ Nexus 92348GC-X • 1 GE 48-Port Gb Ethernet switch	NXOS 9.3.3	1
Brocade	Brocade Fibre Channel switches	▪ G620 • 48 SFP+ and 4 QSFP • Fibre Channel ports • 16/32 Gbps SFPs	V8.2.1e	2

Software components

The following table lists the Unified Compute Platform Converged Infrastructure (UCP CI) solution software components used to run the different test cases for this reference architecture.

Software	Version
Citrix Hypervisor	8.2.0 Long Term Service Release CU1
Citrix Virtual Apps and Desktops	71912 LTSR Cumulative Release CU4
Citrix XenCenter	8.2.1
Microsoft® Windows Server® 2019	Datacenter
Microsoft SQL Server®	2019
Microsoft Windows® 10	2004 Enterprise Edition
Login VSI	4.1.40.1
Additional software	▪ Microsoft Office Professional Plus 2016 64-Bit ▪ Citrix VM Tools for Windows 9.0.42 (64-bit)

Solution design

The infrastructure servers for Citrix that were used for this solution were placed in separate infrastructure clusters with dedicated resources, also called management pools clusters. The compute pool was dedicated to Citrix desktops.

The following figure shows the key hardware and software components used for the reference architecture tests.

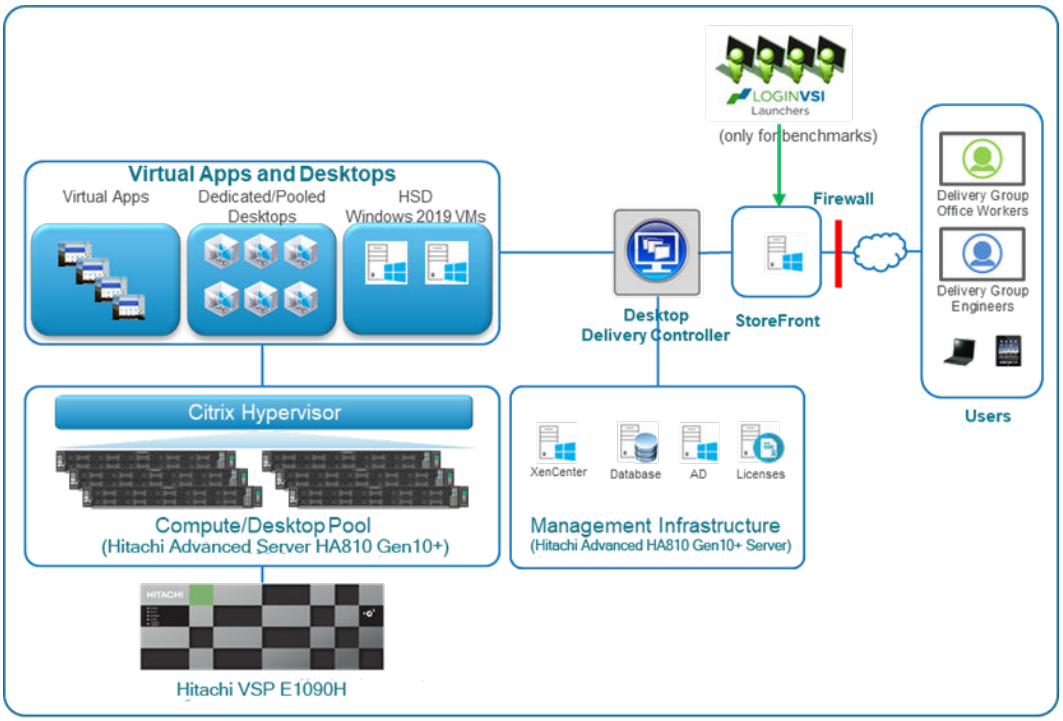

Test environment configuration

The entire infrastructure was configured following best practices and guidelines using the Hitachi Unified Compute Platform with redundancy contained at every level:

- Hitachi Virtual Storage Platform with dual controllers
- Redundancy at the network switch level: 2 × Cisco Nexus switches
- Redundancy at the SAN switch level: 2 × Brocade switches
- Redundancy at the server level

Solution infrastructure

The following figure provides a high-level overview of the solution infrastructure. A 4-node Hitachi Advanced Server HA810 G2 compute pool was used to host the Citrix VDI desktops, a 2-node HA810 G2 pool was used to host Citrix infrastructure management VMs, and Hitachi VSP E1090H SSDs were used for storage. For details, see Hardware components (on page 5).

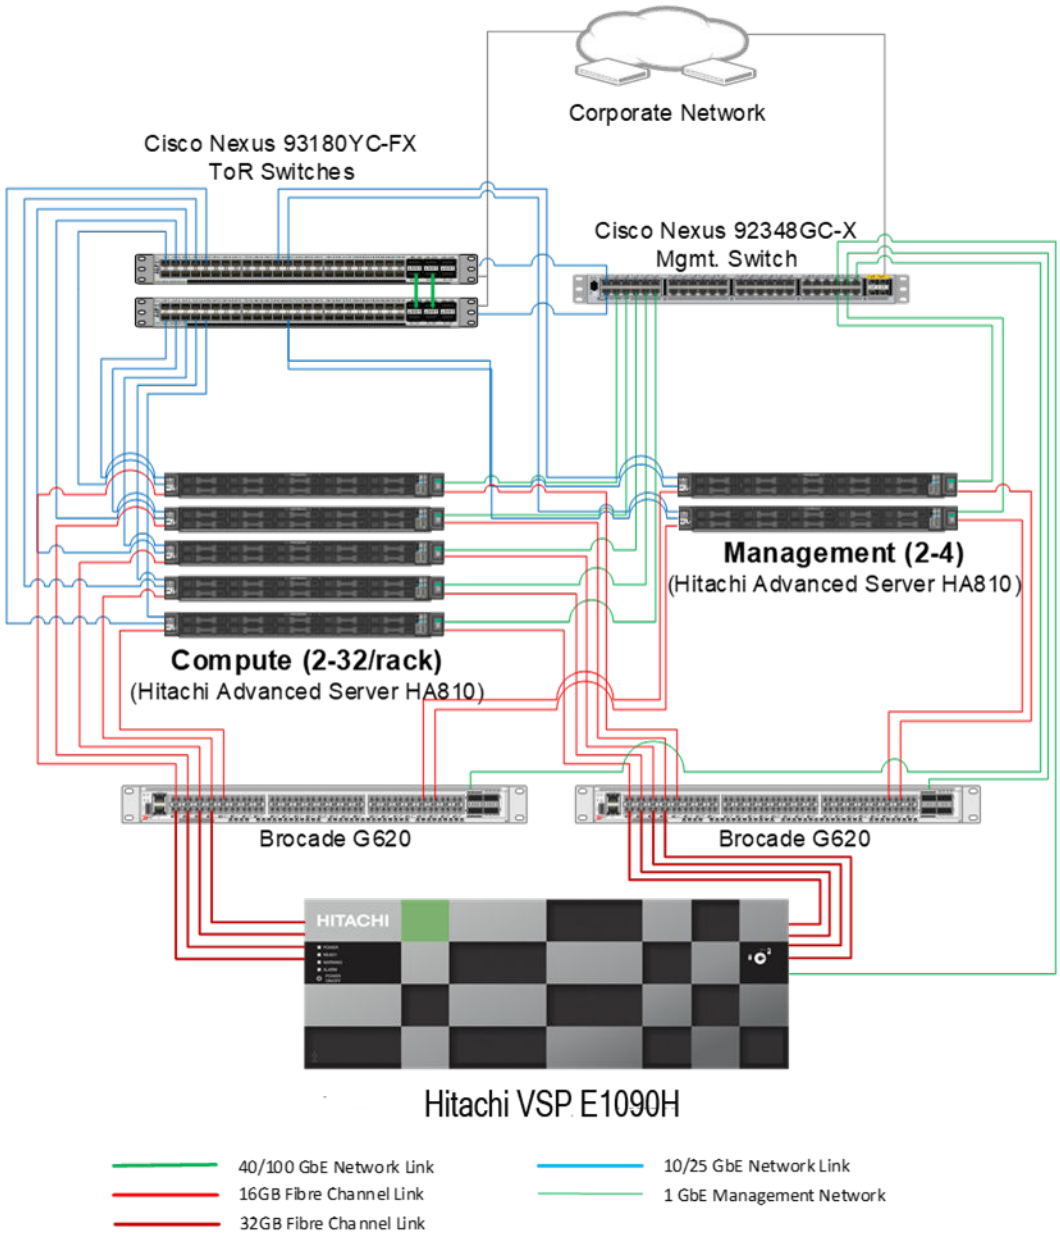

Storage configuration

The storage for this test was configured with 14 LUNs. A single dynamic provisioning pool was configured with one RAID 6 (12D+2P) parity group.

The following figure and table show the layout of the storage provisioning and the virtual volumes (DP-VOL) created and presented to the Citrix hypervisors in the management resource and compute resource pools.

Adaptive Data Reduction (ADR) was enabled for all the virtual volumes to improve storage utilization with inline compression and deduplication. Zoning was configured following Hitachi Unified Compute Platform CI for Citrix VDI best practices.

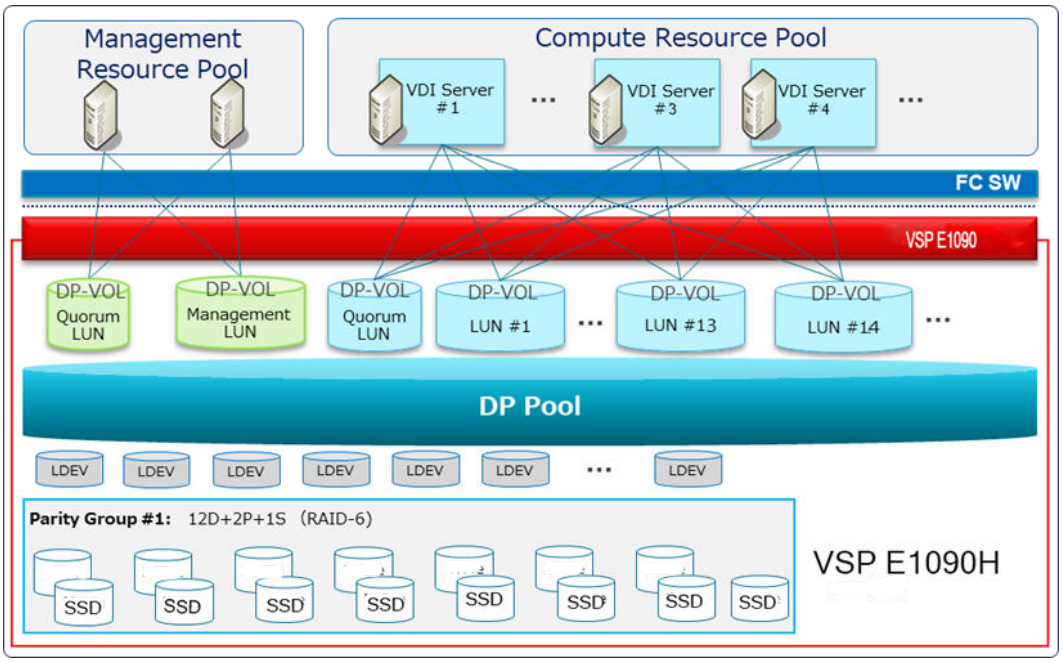

Note: The capacity of the volumes and naming of LDEVs in this table is just an example of the configuration used in the lab.

Virtual volume LDEV ID	Capacity	LDEV Name	Citrix Compute Pool/Cluster
00:0D:00	400 MB	Quorum_Mgmt_lun	Management Pool
00:0D:01	600 GB	UCP_Management_lun	
00:01:00	400 MB	Quorum_Compute_lun	Compute Pool
00:01:01 - 00:01:0C	3.3 TB	UCP_citrixvdi_lun1- 12	

The following figure shows information about the Hitachi VSP Storage system.

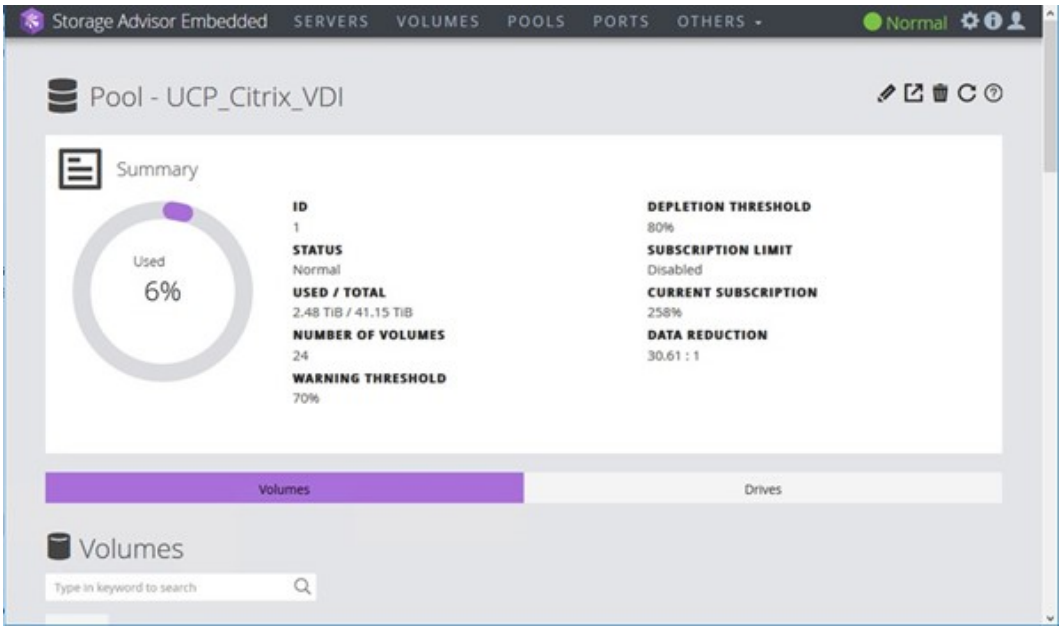

For additional details and features of the Hitachi Virtual Storage Platform E1090H see Hitachi Virtual Storage Platform E Series.

Network configuration

The network switches were configured following the best practices documented as part of Hitachi UCP solutions.

All the hosts in the management and compute pools (or clusters) were configured with NIC bonding in active-active mode, with each host contributing two 25 GbE network interfaces. Separate virtual networks were configured for management/XenMotion traffic, VDI desktops, and infrastructure virtual machines. The following figure shows a logical diagram of the network configuration at the hypervisor level.

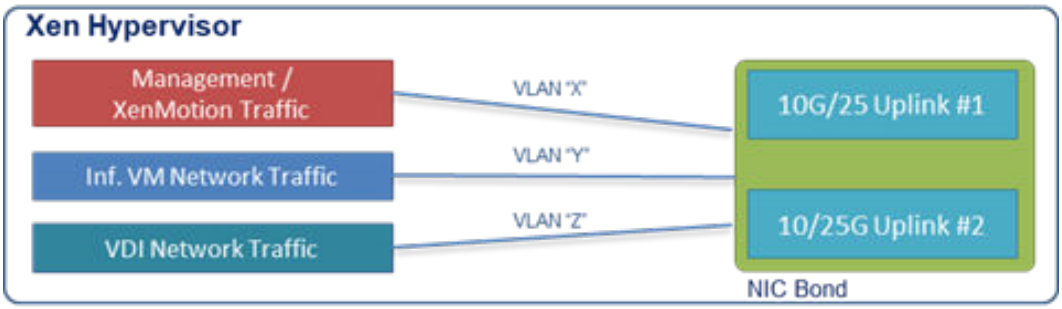

The following figure shows additional details about the network configuration for the compute pool as seen from XenCenter.

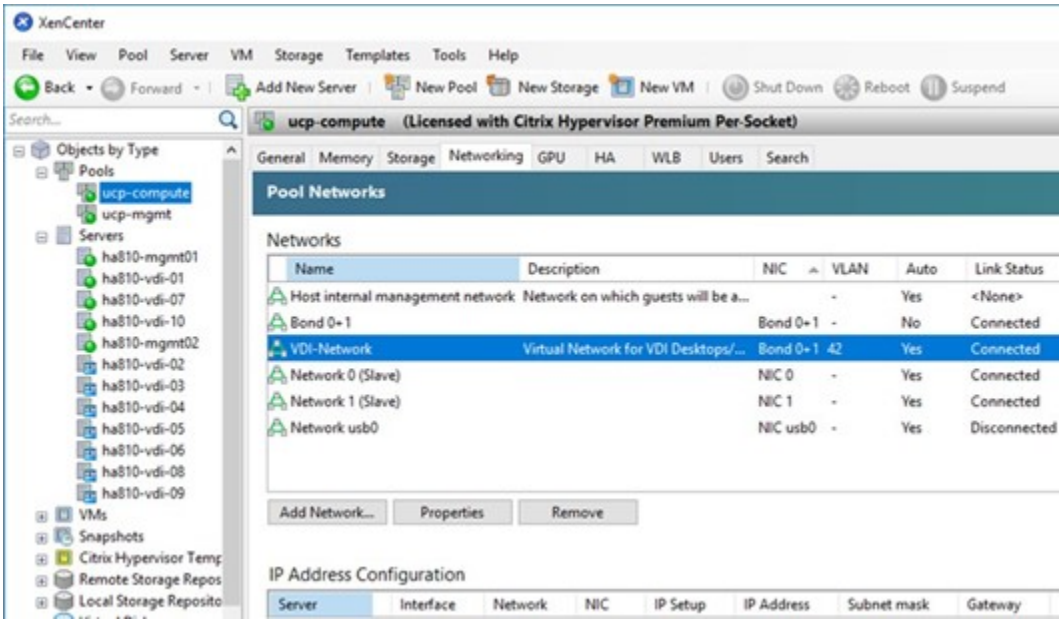

Infrastructure virtual machines

Citrix infrastructure virtual machines

The following table shows the configuration of different Citrix infrastructure virtual machines that support this solution.

Server Name	vCPU	Memory	Disk Size	Operating System
Database	4	16 GB	60 GB (OS) 200 GB(data)	Microsoft Windows Server 2019 Microsoft SQL Server 2019
Citrix Delivery Controller	4	12 GB	100 GB	Microsoft Windows Server 2019
Citrix Storefront	4	8 GB	100 GB	Microsoft Windows Server 2019
Citrix License Server	2	8 GB	40 GB	Microsoft Windows Server 2019
Citrix XenCenter	2	8 GB	40 GB	Microsoft Windows Server 2019
Domain Controller	4	8 GB	100 GB	Microsoft Windows Server 2019

The domain controller was deployed to support user authentication and domain services for the Citrix infrastructure and Login VSI.

Hosted virtual desktops (HVD) virtual machine configuration

The following table shows the configuration profiles of the Windows 10 master VM templates.

VM Profile	VM OS	OS vDisk	vCPU	VM Memory (configured)	VM Memory (reserved)
Profile 1	Windows 10 v2004 Enterprise	36 GB	4	8 GB	8 GB

Hosted shared desktops (HSD) virtual machine configuration

The following table shows the configuration of the HSD master VM. The Windows Server 2019 VM for HSD was configured to not spam physical NUMA nodes and thus ensured optimum performance.

VM Profile	VM OS	OS vDisk	vCPU	VM Memory (configured)	VM Memory (reserved)
HSD VM Profile	Windows Server 2019	60 GB	8	32 GB	32 GB

The number of deployed Citrix virtual desktops is based on the CPU/memory resources on the servers of the compute pool cluster.

Solution deployment

To deploy this solution, complete the following procedures.

Hitachi VSP Storage deployment

The storage infrastructure for the compute and management resource pool cluster was configured following the best practices for Hitachi VSP storage. For more information on how to configure VSP E1090H for this solution, see the *VSP E1090 Hardware Reference* at https://knowledge.hitachivantara.com/Documents/Storage/VSP_E_Series/93-06-4x/VSP_E1090_Hardware_Reference.

Hitachi infrastructure servers deployment

The Citrix Hypervisor was installed on the management and compute nodes following the Citrix Installation Guide and best practices available in the Citrix Product Documentation Install inforrmation at https://docs.citrix.com/en-us/citrix-hypervisor/install.html.

The installation media for Citrix Hypervisor and XenCenter is available on the Citrix Hypervisor downloads website at https://www.citrix.com/downloads/citrix-hypervisor/:

- Citrix Hypervisor 8.2.0 Base Installation ISO
- XenCenter 8.2.1 Windows Management Console Installer

The following resource pools were configured and managed by XenCenter to support the solution and test environment:

- Citrix XenCenter
- VDI Management Pool
  - Two Hitachi Advanced Server HA810 G2 nodes were configured on the management pool.
  - This pool hosted the following Citrix infrastructure VMs: XenCenter, Domain Controller, Database, Delivery Controller, StoreFront, and License Server.
- VDI Compute Pool
  - Four Hitachi Advanced Server HA810 G2 nodes were configured on the compute pool.
  - This pool hosted the Citrix VDI persistent desktops with Windows 10 and HSD with Windows Server 2019. All of these desktops were provisioned with Citrix Machine Creation Services (MCS).

Both compute and management resource pools were configured with High Availability (HA), multipathing for Fibre Channel storage, and NIC bonding for the network traffic.

Citrix infrastructure component deployment

These procedures assume that virtual machine templates have been prepared with the Windows 10 and Windows Server 2019 operating systems, including the latest patches.

This deployment process consists of the following:

1. Configure a Citrix License Server
2. Configure Citrix Delivery Controller for MCS provisioning
   - Create a Site
   - Configure a XenDesktop Site Hosting Connection
3. Configure Citrix StoreFront
   - Create a store to provide users access to their virtual desktops
4. Prepare the Citrix XenDesktop master images
   - Master image for Windows 10, for Hosted Virtual Desktops (HVD)
   - Master image for Windows Server 2019, for Hosted Shared Desktops (HSD)
5. Deploy the MCS catalog
   - Deploy the persistent desktop catalog with a Windows 10 image
   - Deploy the HSD catalog with Windows Server 2019
6. Create a delivery group, and assign a machine catalog to the delivery group

Login VSI deployment for load generation

The Login VSI infrastructure was hosted in a separate compute/storage environment and connected to the same set of switches as the target systems running the VDI desktops. The launchers and Login VSI environment were configured and managed by a centralized management console, and followed Login VSI best practices:

- Knowledge Worker Workload — The Knowledge Worker workload is an intensive workload that balances the system stresses smoothly.
- Power Worker Workload — The standard Power Worker workload is a very intensive workload that puts maximum stress on the system. Many applications in this workload use larger files and higher resolution media.

Solution testing and results

This section describes the test methodology, tools, and configuration used to validate the Hitachi Infrastructure for Citrix Workspace solution using Hitachi Virtual Storage Platform E1090H.

Test methodology

Load generation tool

Login VSI was used to generate Knowledge Worker and Power Worker workloads on the desktops. Login VSI launchers were configured to initiate no more than 24 sessions to the Citrix StoreFront/Delivery Controller Servers to simulate an end-to-end execution of the entire Citrix VDI infrastructure stack.

Login VSI Knowledge and Power Worker workload profiles were used. The Login VSI Pro Library package was used for tests with persistent desktops/full clones and Power Worker workloads.

The following table lists the applications that Login VSI exercised during workload testing of both workload profiles.

Workload Type/ Applications	Knowledge Workload	Power Workload (with Pro Library)
Applications Exercised	▪ Adobe Acrobat ▪ Freemind/Java ▪ Microsoft Internet Explorer ▪ Microsoft Excel ▪ Microsoft Outlook ▪ Microsoft PowerPoint ▪ Microsoft Word	▪ Adobe Acrobat ▪ Freemind/Java ▪ Microsoft Internet Explorer ▪ Microsoft Excel ▪ Microsoft Outlook ▪ Microsoft PowerPoint ▪ Microsoft Word ▪ Notepad

Workload Type/ Applications	Knowledge Workload	Power Workload (with Pro Library)
	▪ Notepad ▪ Photo Viewer	▪ Photo Viewer ▪ Simulated application installs ▪ 7-Zip Compression ▪ Microsoft Windows Media Player ▪ Login VSI Pro Library

VM configuration and workloads

Each image was optimized with the Login VSI tuning templates and the Citrix Optimizer tool at https://support.citrix.com/article/CTX224676/citrix-optimizer-tool, in conformance with Login VSI test standards. In addition, all Windows 10 desktops were running Windows Defender antivirus software.

The following table shows the configuration profiles for the Windows 10 VM templates used for the Power Worker workloads.

VM Profile	Login VSI Workload	VM OS	vCPU	VM Memory	Resolution
Profile 1	Power Worker	Windows 10 v2004 Enterprise	4	8 GB	1920 × 1080

The following table shows the configuration of the HSD VM. The Windows Server 2019 VM for HSD was configured to not spam physical NUMA nodes to ensure optimum performance.

VM Profile	Login VSI Workload	VM OS	vCPU	VM Memory	Resolution
HSD VM Profile	Knowledge Worker	Windows Server 2019	8	32 GB	1920 × 1080

Test cases

The tests for this reference architecture were designed to capture the performance capabilities and user experience (based on Login VSI workloads) when using VSP E1090H.

The following table shows a summary of the test cases, and the total number of hosts and desktops tested in each case.

Test Case	Login VSI Workload	Test	Total VMs /Sessions
Persistent VDI Desktops/Full Clones	Power Worker	2 hosts	240 VMs (120/hosts)
		Scalability test up to 4 hosts	480 VMs
HSD Test	Knowledge Worker	1 host	12 VMs/300 sessions
		Full-scale test with 4 hosts	48 VMs/1200 sessions

The following tests cases were designed to determine the maximum number of desktops/ sessions per host.

ServerN+1 fault tolerance was not factored in for the compute pool when running these tests. Always consider high availability for a production environment by adding one extra host per compute pool, reducing the number of virtual machines, or reducing the virtual CPU/memory allocated to the virtual machines.

Persistent VDI Desktops: Power Worker Workload

For this test, the desktop catalog was created using Citrix Machine Creation Services (MCS) with single-session Windows 10 OS, a persistent desktop experience, and full copy (or full clones).

The virtual machine template used for full clones was configured for a Power Worker workload. All tests were run with 120 VMs per host; this is the maximum number of VMs that can be run on a single host based on the host and VM configuration defined for this profile.

The desktop pool was configured to use 14 storage repositories (SRs), which were presented to all the hosts in the compute cluster/pool from 14 volumes in the Hitachi VSP Storage Pool.

Virtual desktops were evenly distributed across all SRs and Citrix hypervisors. For each test, the same number of VMs were started on each host in the compute pool to obtain accurate user experience metrics during testing.

HSD Test: Knowledge Worker Workload

For this test, the desktop catalog was created using Citrix Machine Creation Services (MCS) with multi-session Windows Server 2019. This catalog provides Hosted Shared Desktops (HSD) for large-scale deployment of standardized Windows multi-session machines.

The virtual machine template was configured. All tests were run with 290 sessions per host; this is the maximum number of sessions that can be run on a single host based on the host and HSD/VM configuration defined for this profile.

The desktop pool was configured to use 14 storage repositories (SRs), which were presented to all the hosts in the compute cluster/pool from 14 volumes from the Hitachi VSP storage pool.

The Windows Server 2019 VMs were evenly distributed across all the SRs and Citrix hypervisors. For each test, the same number of VMs were started on each host in the compute pool to obtain accurate user experience metrics during testing.

Persistent VDI Desktops: Power Workers

Workload testing

Login VSI was configured to launch a Power Worker workload profile. The test environment had the following configuration:

- All desktops were restarted before each test. After all desktops were active and registered on Citrix Studio, there was at least a 15-20 minute wait time before starting each test.
- For all tests, Login VSI was configured to stagger all user logins during a period of 48 minutes (the 2880-second standard benchmark launch rate), followed by 20 minutes of steady workload, and then 10-20 minutes for logoff of all sessions.
- All Login VSI Launchers were restarted before each test.
- The number of launchers used for each test was sufficient to not exceed 24 sessions per launcher.
- The connection resolution used for the sessions was 1920 × 1080.

Two Server Test: Persistent Desktops and Power Worker

Login VSI - test results

The Login VSI Max user experience score for the test with 240 Power Worker workloads was not reached. This means that the system was not saturated, and it was able to complete the test without exceeding the average response time latency threshold.

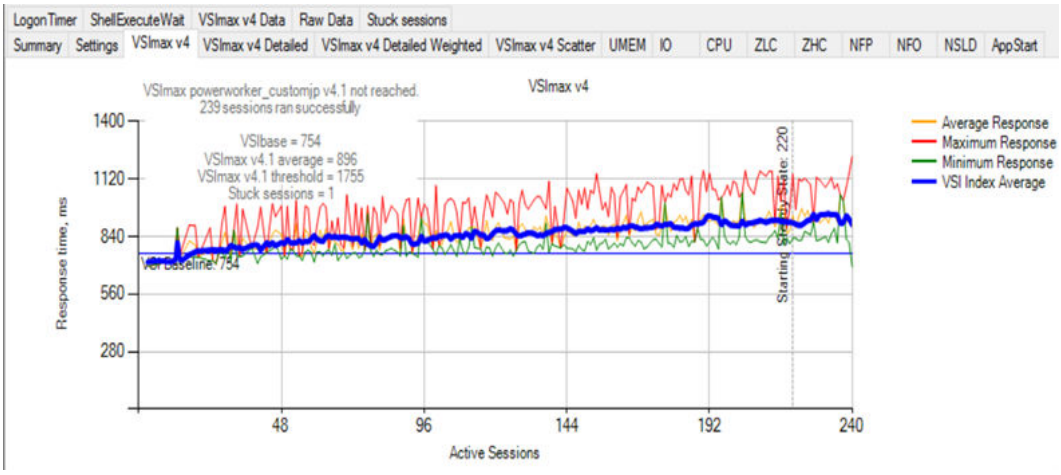

- The test completed with 239 of the 240 successful power worker sessions.
- With this number of sessions, the Login VSI baseline performance score was 754, which was Very Good based on the Login VSI baseline performance rating.
- This indicates that the number of power workloads tested did not put any strain on the system resources.

Hypervisor performance

The following figure illustrates the performance of one of the hosts in the compute pool/ cluster running a Power Worker workload during login, steady state, and logoff operations.

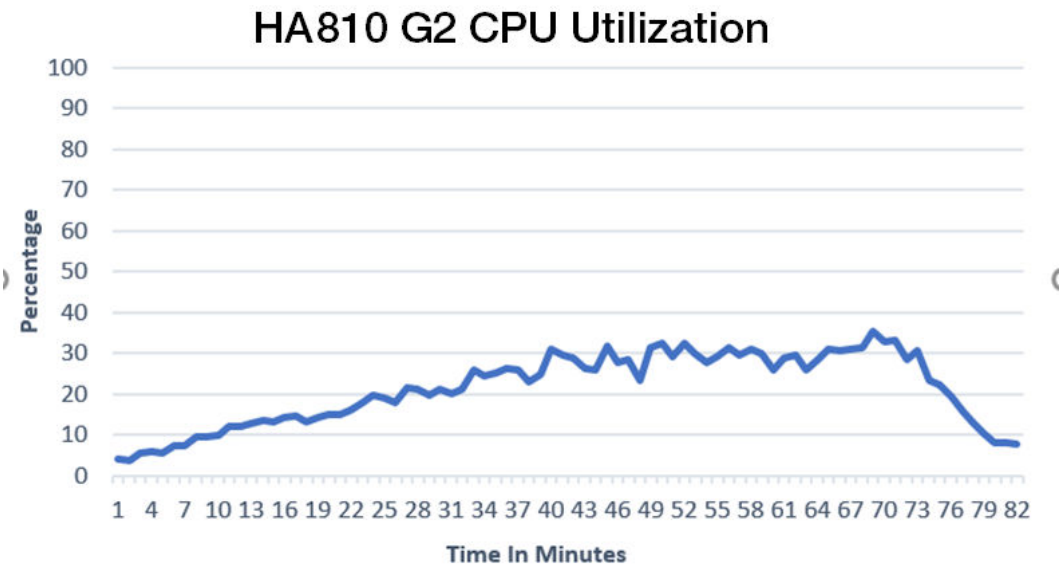

- The performance metrics of the hypervisor show the following:
  - CPU utilization does not rise above 40% during all the tests when 120 power VDI profile VMs are running on one host.
  - Memory utilization was more than 90% because the VMs were configured with full memory reservation.
  - Considering that each hypervisor has 2 × 25 GbE ports, these results show that network bandwidth is not an issue and that there is plenty of bandwidth available for additional workloads.
  - Also, this shows that there is still plenty of headroom from the CPU perspective on the compute pool to support bursts in workloads while maintaining performance.

VSP E1090 performance

The following figures show IOPS, latency, and MPU usage at the VSP storage pool level during the logon, steady state, and logoff phases for all scalability tests.

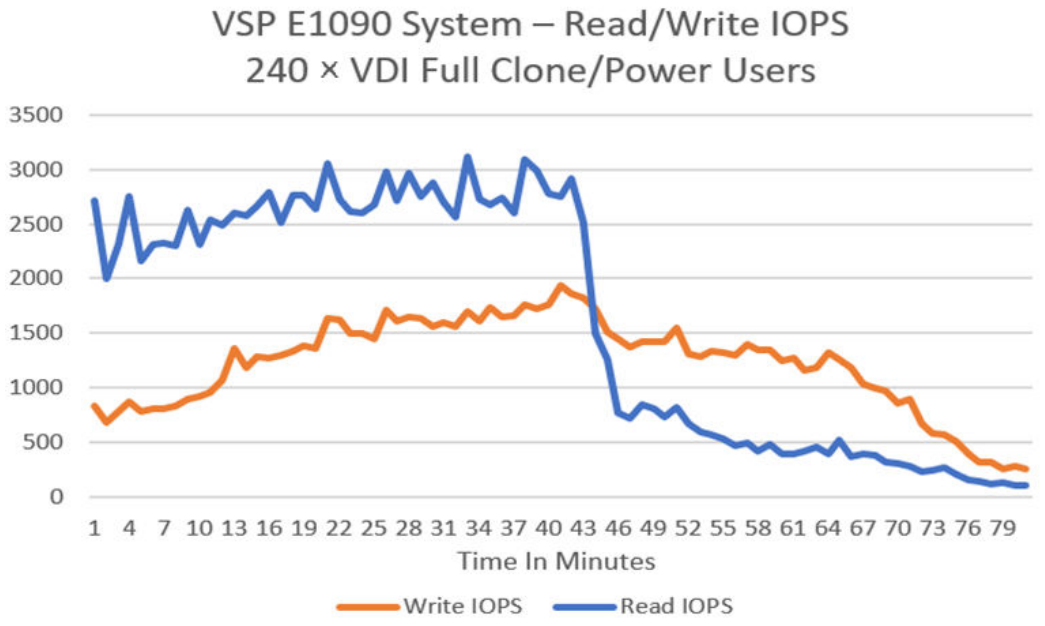

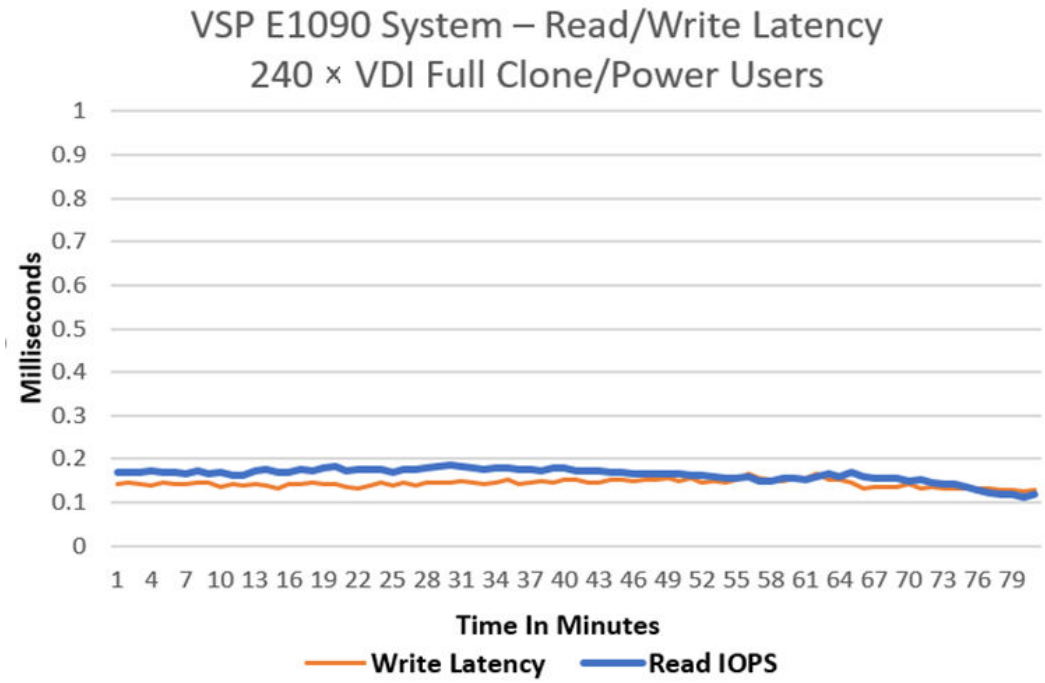

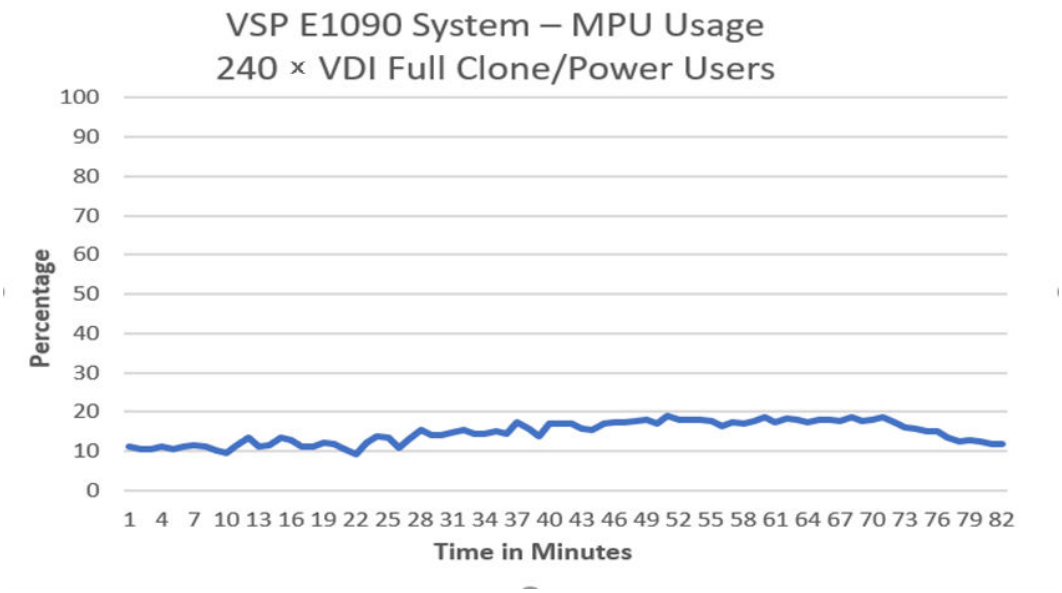

Here are some of the key metrics for the test with 2 hosts and 240 Power Worker workloads:

- The cluster reached a maximum peak of 2000 Write IOPS during the logon phase.
- The cluster reached a maximum peak of 3100 Read IOPS during the logon phase.
- The read response time did not rise above 0.2 milliseconds during the duration of the test.
- The write response time was ~0.2 milliseconds during the duration of the test.
- The MPU usage reached a maximum peak of 18%.

Scalability Test: Persistent Desktops and Power Worker

As part of the scalability tests, two additional tests were run with 480 desktops running on 4 hosts, and 960 desktops running on 8 hosts, while maintaining the same configuration on the Virtual Storage Platform storage system.

Login VSI - test results

The Login VSI Max user experience score for the Power Worker test was not reached. This means that the system was not saturated, and it completed the test without exceeding the average response time latency threshold.

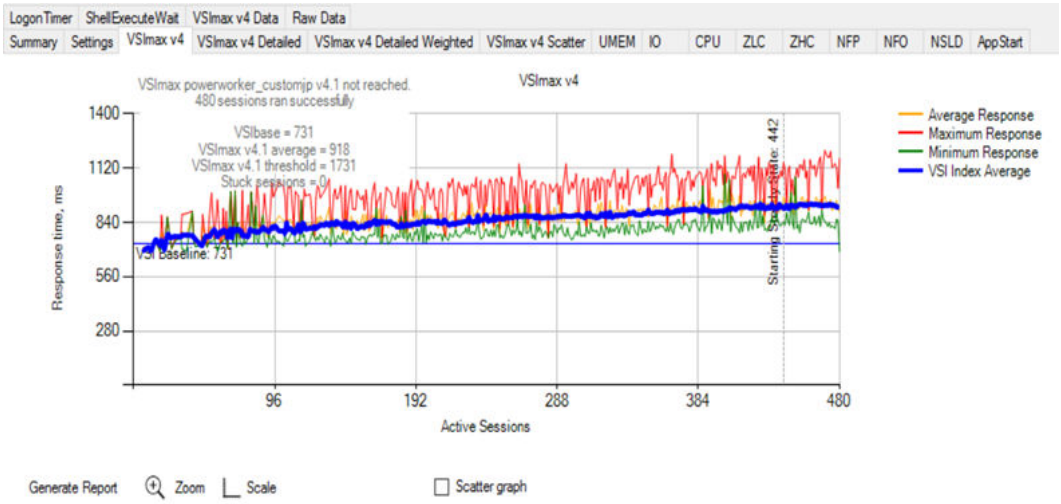

- The test with 4 hosts and 480 desktops completed with 480 successful power worker sessions.
- With this number of sessions, the maximum capacity VSI max (v4.1) was not reached and the Login VSI baseline performance score was 731. This score is rated Very Good based on the Login VSI baseline performance rating.

This indicates that the number of Power Worker workloads tested did not strain the system resources.

The following figure shows a summary of the results of the scalability tests with 1, 2, and 4 hosts. It clearly shows linear scalability. Moreover, it shows that the Login VSI baseline performance score was not impacted at all, even when doubling the number of hosts and desktops. Overall, the average response time remained nearly identical across all of the tests.

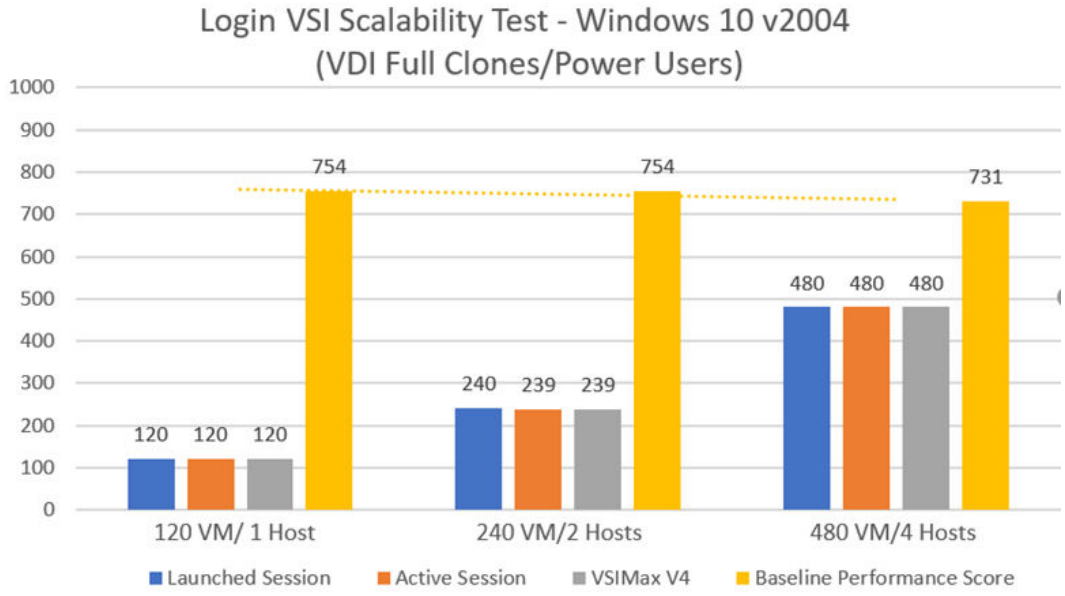

Hypervisor performance

The following figure illustrates the performance of one of the hosts in the compute pool/ cluster running Power Worker workloads during login, steady state, and logoff operations.

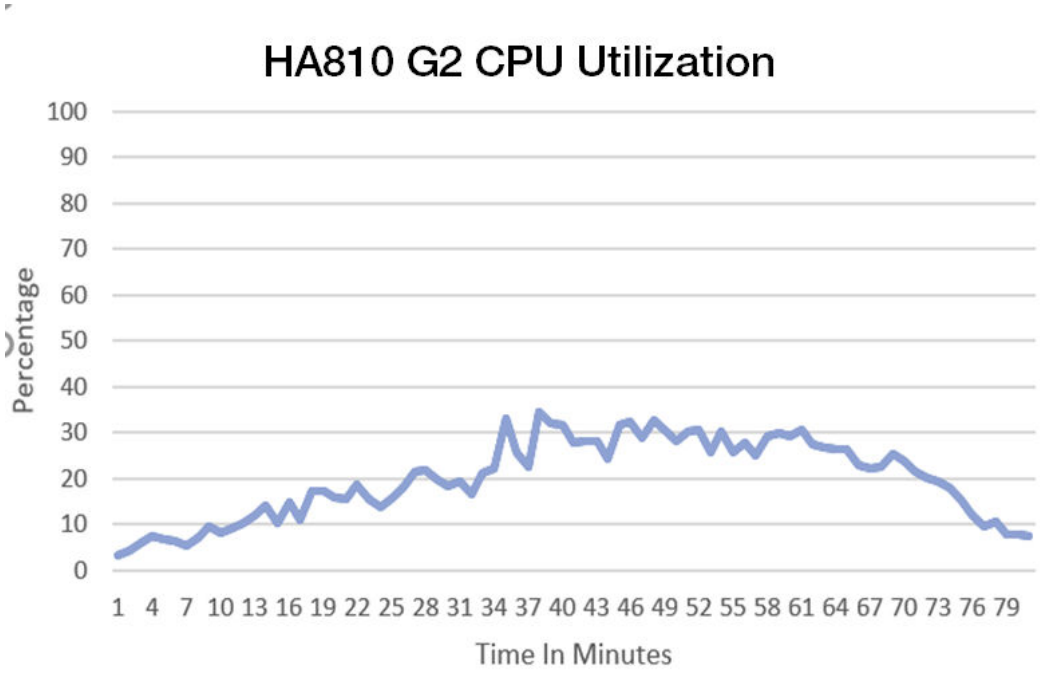

- The performance metrics indicate that CPU utilization does not rise above 40% during all the tests. In addition, the following was observed:
  - Memory utilization was more than 90% because all VMs were configured with full memory reservation.
  - Considering that each hypervisor has 2 × 25 GbE ports, these results show that network bandwidth is not an issue and that there is plenty of bandwidth available for additional workloads.
- This shows that there is still plenty of headroom from the CPU perspective on the compute pool to support bursts in workloads while maintaining user performance.

VSP E1090H performance

The following figures show IOPS, latency, and MPU usage at the VSP storage pool level during the logon, steady state, and logoff phases for all scalability tests.

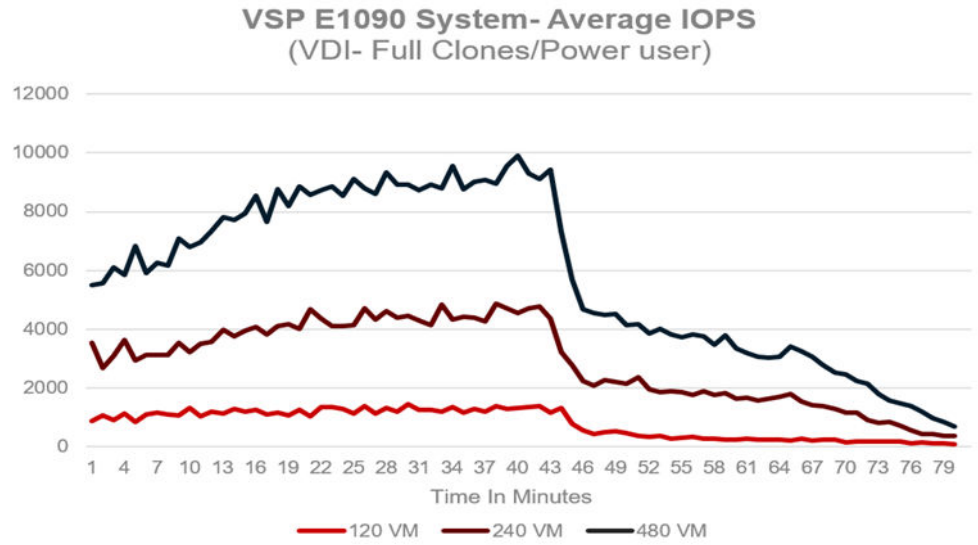

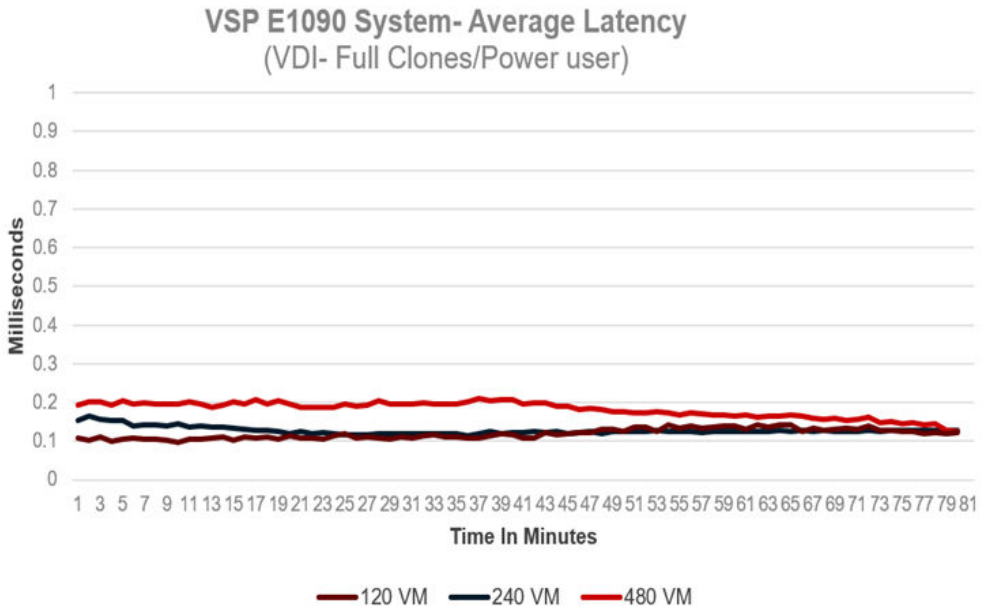

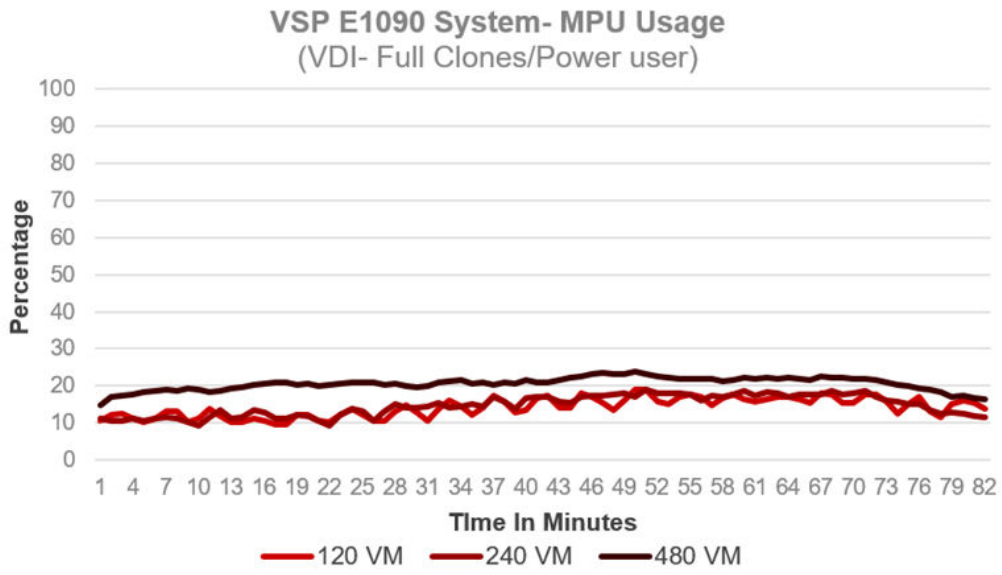

The following are some of the key metrics for the test with 1, 2, and 4 hosts with 120, 240, and 480 VMs respectively for the Power Users profile:

- The cluster reached a maximum peak of 9907 total IOPS during the logon phase.
- The average response time did not rise above 0.25 milliseconds.
- The MPU usage did not rise above 22% during all the tests.
- Considering that the storage system for these tests was configured with a single parity group with 14 LUNs, these IOPS/CPU and latency results indicate that there is plenty of headroom for additional workloads.

HSD with Windows Server 2019: Knowledge Workers

Workload testing

Login VSI was configured to launch a Knowledge Worker workload profile. The test was executed with the following configuration:

- All Windows Server 2019 VMs were restarted before each test. After all VMs were active and registered on Citrix Studio, there was at least a 15 minute wait time before starting each test.
- For all tests, Login VSI was configured to stagger the logins of all users during a period of 48 minutes (the 2880-second standard benchmark launch rate), followed by 20 minutes of steady workload, and then 10-20 minutes for logoff of all the sessions.
- All Login VSI launchers were restarted before each test.
- The number of launchers was sufficient to not exceed 24 sessions per launcher.
- The connection resolution used for the sessions was 1920 × 1080.

Single Server Test: HSD Sessions and Knowledge Workers

Login VSI - Test Results

The Login VSI Max user experience score for the test with 290 Knowledge Worker workloads was not reached. This means that the system was not saturated, and that it was able to complete the test without exceeding the average response time latency threshold.

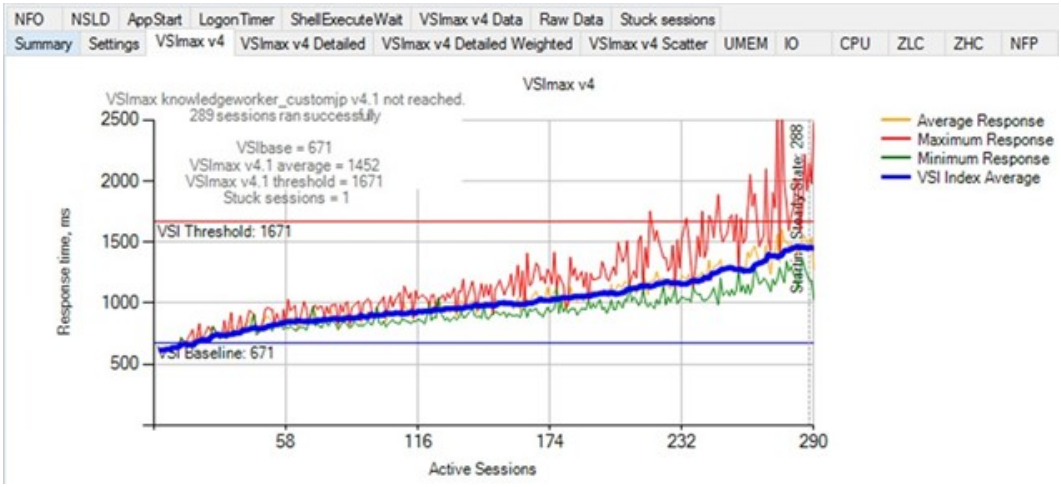

- The test completed with 289 of the 290 successful knowledge worker sessions.
- With this number of sessions, the Login VSI baseline performance score was 671, which is rated Very Good based on the Login VSI baseline performance rating.

Hypervisor performance

The following figure illustrates the performance of one of the hosts in the compute pool/ cluster running HSD/Windows Server 2019 with Knowledge Worker workloads during login, steady state, and logoff operations.

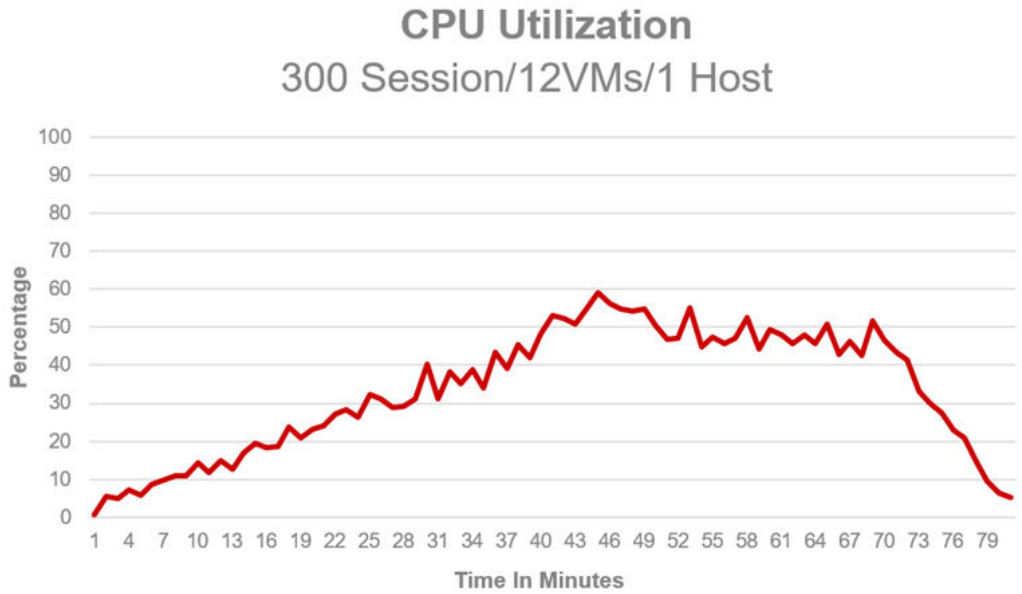

The performance metrics indicate that CPU utilization does not rise above 60% during all the tests. In addition, the following was observed:

- Memory utilization was less than 50%, considering that only 12 × Windows Server 2019 VMs were configured per host.
- Write latency at the host/HBA level did not rise above 2.3 ms.

Windows 2019 VM performance

The following figure illustrates the performance for one of the Windows Server VMs running Knowledge Worker workloads.

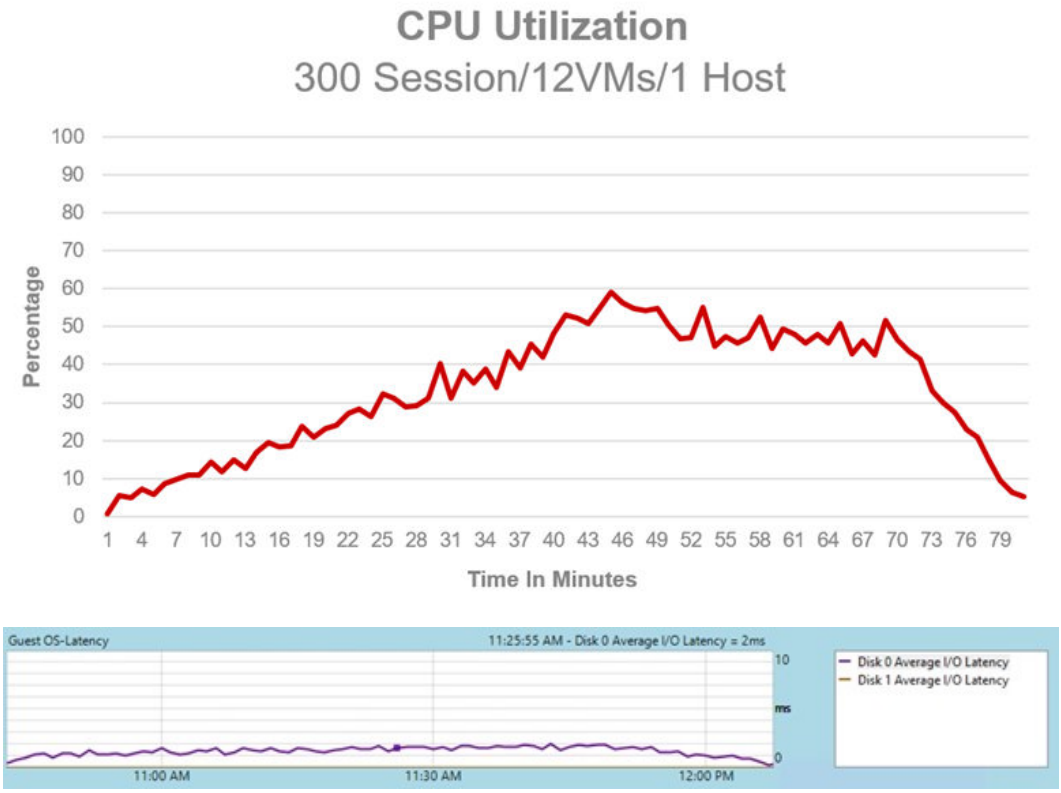

The performance metrics from XenCenter show the following:

- CPU utilization at the Guest OS level does not rise above 60% during all the tests.
- Memory utilization at the Guest OS level was not more than 20 GB (62%) of utilization.
- The latency at the Guest OS level did not rise above 2 ms.

VSP E1090H performance

The following figures show IOPS, latency, and MPU usage at the VSP storage pool level during logon, steady state, and logoff phases.

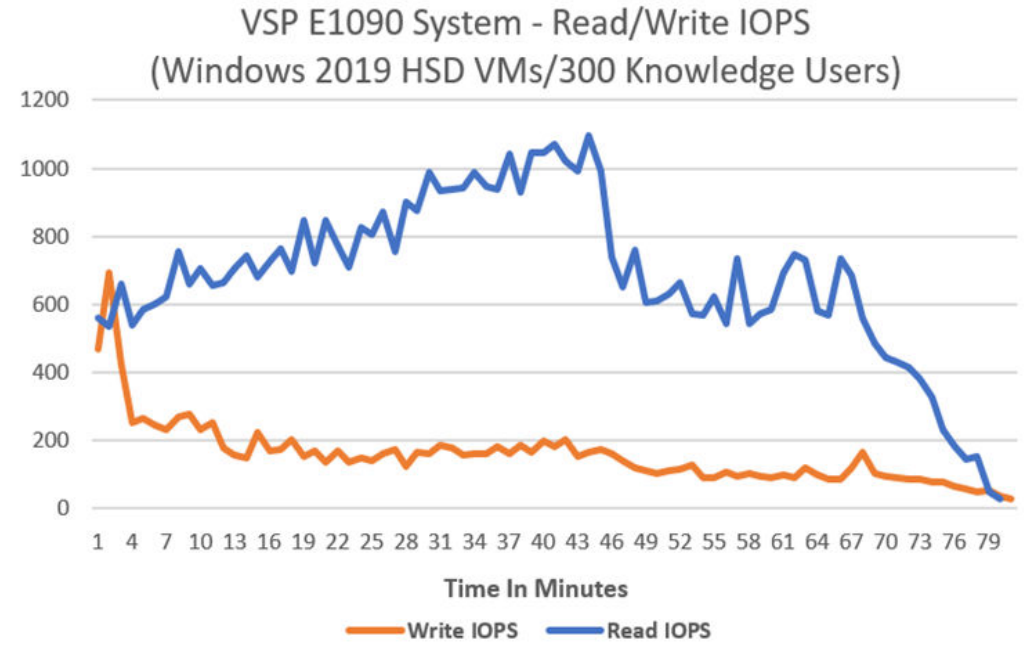

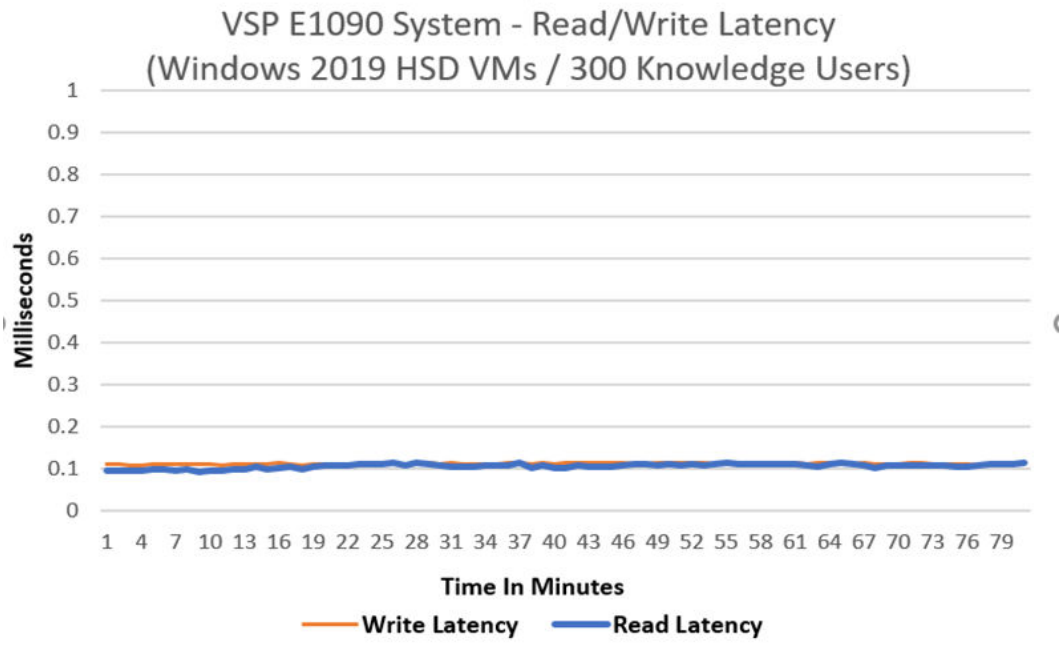

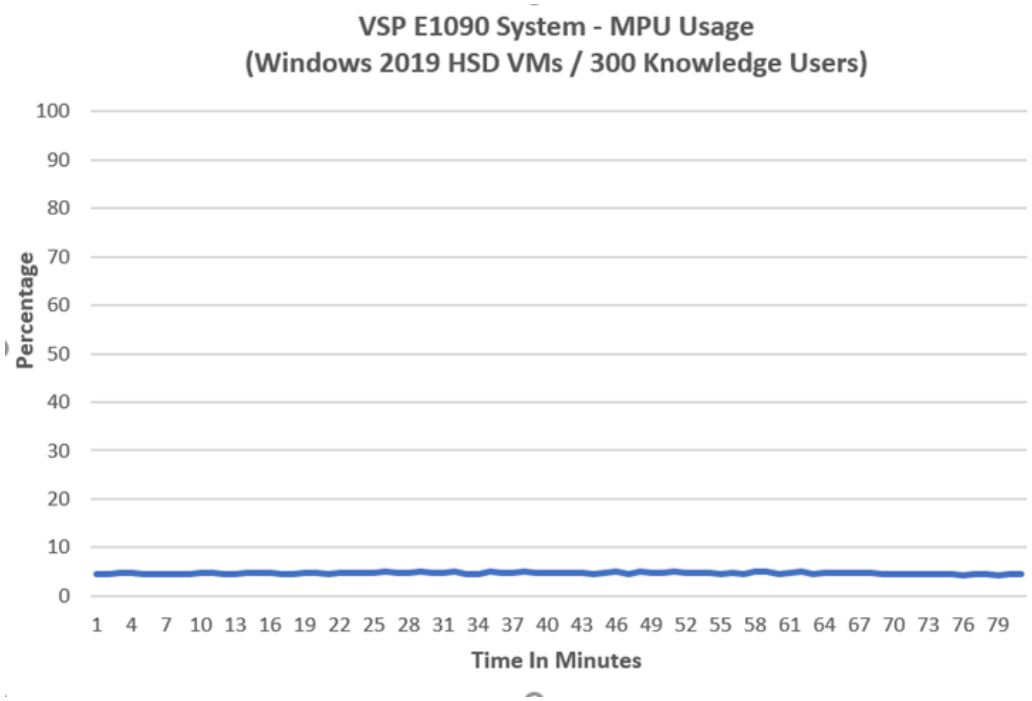

The following are some of the key metrics for the large test with 1 host and 290 Knowledge Worker workloads:

- The system reached a maximum peak of 1720 read IOPS during the logon phase.
- The system reached a maximum peak of 890 write IOPS during the logon phase.
- The average response time did not rise above 0.3 milliseconds.
- The MPU usage did not rise above 10% during all of the tests.

Considering that the storage system for these tests was configured with a single parity group (14 NVMe devices), these IOPS/CPU and latency results indicate that there is plenty of headroom for additional workloads.

Full-scale Test: HSD Sessions and Knowledge Workers

This test was set up with 1160 sessions using only 4 hosts, while maintaining the same configuration on the VSP storage system.

Login VSI - test results

The Login VSI Max user experience score for the test with 1200 Knowledge Worker workloads was not reached. This means that the system is not saturated, and that it was able to complete the test without exceeding the average response time latency threshold.

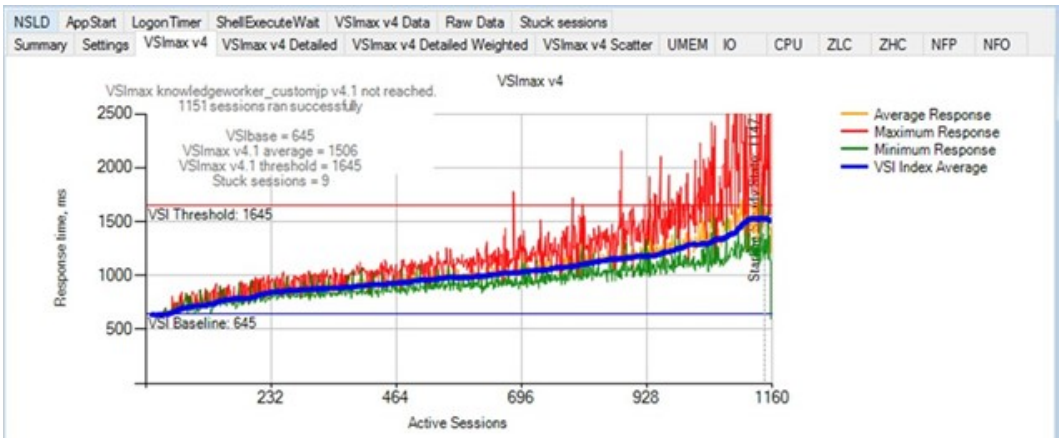

- The test completed with 1151 of the 1200 successful Knowledge Worker workload sessions.
- With this number of sessions, the Login VSI baseline performance score was 645, which is Very Good based on the Login VSI baseline performance rating.

This indicates that the number of Knowledge Worker workloads tested did not put any strain on the system resources.

VSP E1090 performance

The following figures show IOPS, latency, and MPU usage at the VSP storage pool level during the full-scale test.

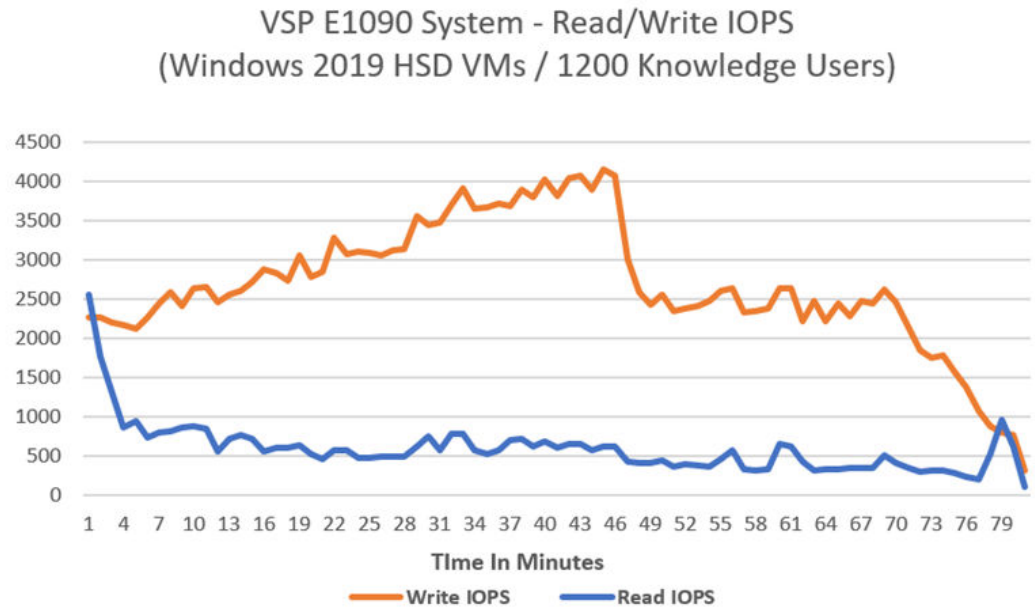

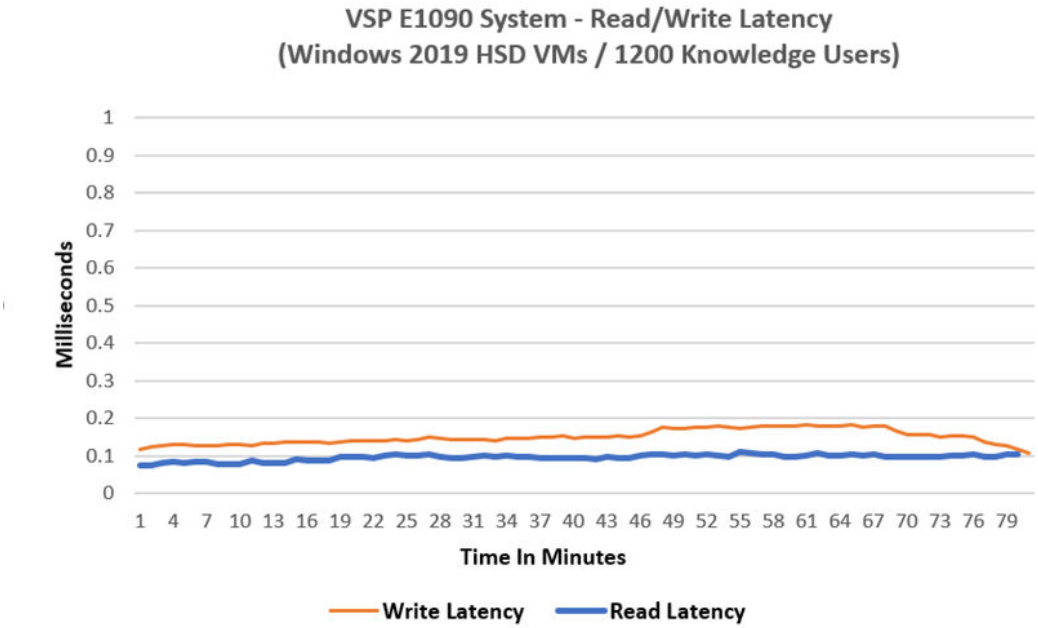

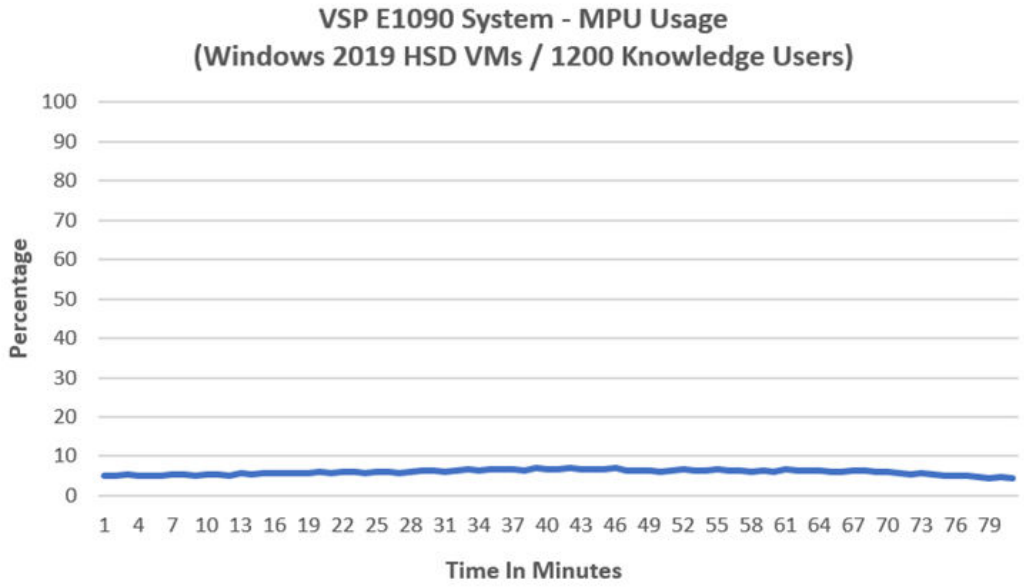

The following are some of the key metrics for the test with 4 hosts and 12000 knowledge users:

- The system reached a maximum peak of 2633 Read IOPS during the logon phase.
- The system reached a maximum peak of 4500 Write IOPS during the logon phase.
- The average response time did not exceed 0.2 milliseconds.
- The MPU usage did not rise above 10% during all the tests.

Considering that the storage system for these tests was configured with a single parity group, these IOPS/CPU and latency results indicate that there is plenty of headroom for additional workloads.

Adaptive data reduction - optimizing storage efficiency

VSP E1090H provides adaptive data reduction (ADR) technology derived from Hitachi Storage Virtualization Operating System RF. A first for the VSP E series, hardware-assisted data reduction technology offers improved data reduction, reducing the number of drives, guaranteeing more usable capacity, and providing up to 82% adaptive data reduction throughput improvement over existing VSP E series.

This technology supports selectable inline compression and deduplication that can be set at the volume level, enabling you to optimize the system for specific service level agreements (SLAs) and capacity needs. In short, your organization will store more with less storage, which significantly impacts your ability to improve storage utilization, reduce your storage footprint, and control costs.

The following figures show comparisons between ADR-enabled vs ADR-disabled IOPS, and latency at the VSP storage pool level during the logon, steady state, and logoff phases for all scalability tests using 4 hosts and 480 full clone VDI.

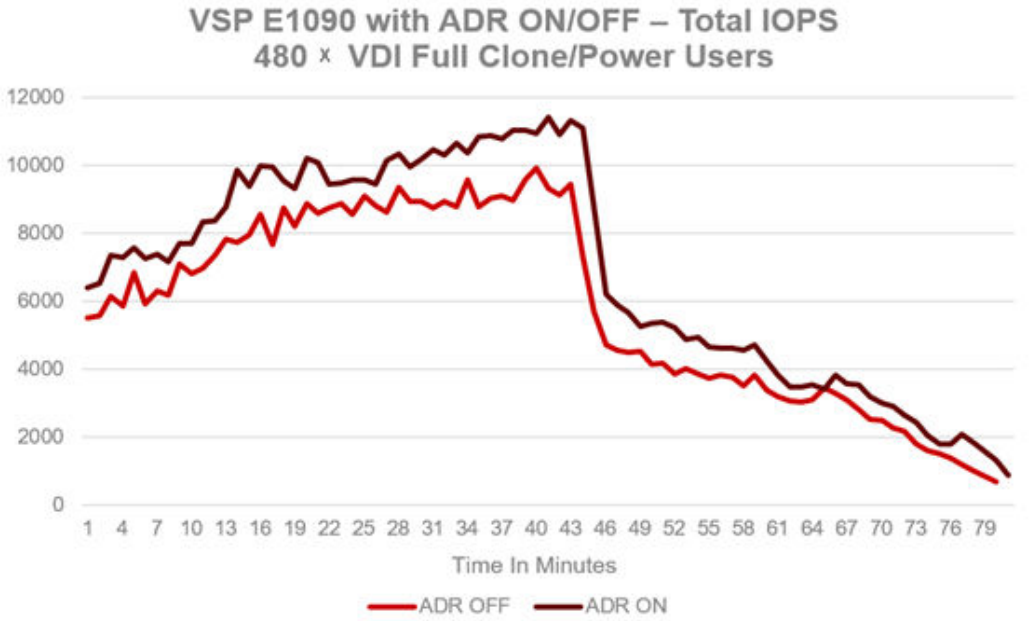

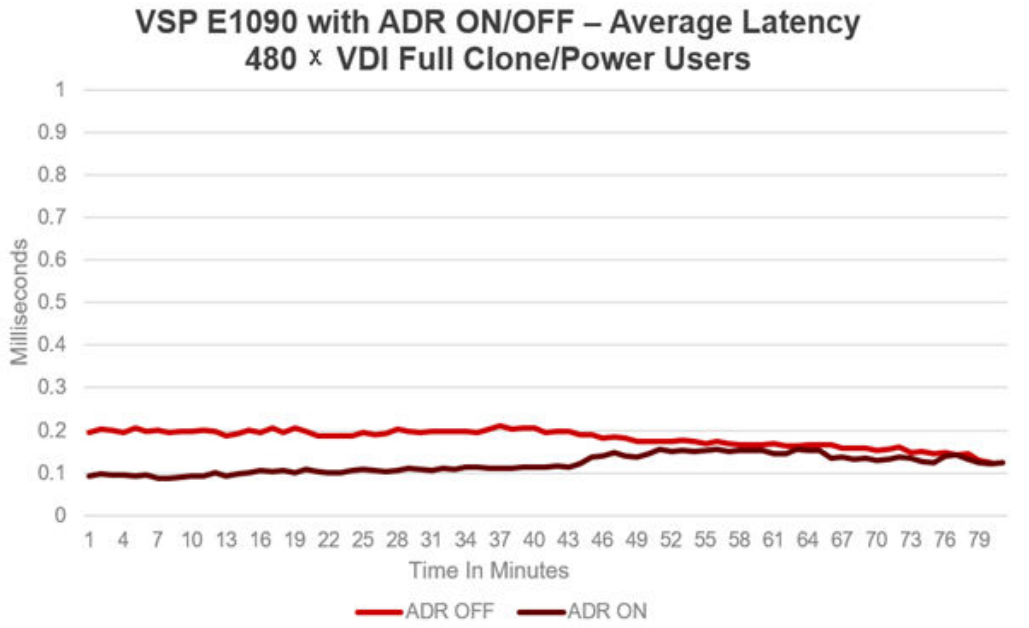

The storage performance metrics show the following:

- Login VSI baseline performance scores are very similar across the different tests run with VSP E1090 with ADR Enabled or Disabled.
- Average latency ON VSP E1090 ADR On is ~50% lower than ADR OFF.
- Total IOPS rises by 20% when tested with ADR On compared with ADR Off.
- While the user experience (Login VSI scores) is not impacted when ADR is ON or OFF, the big difference is in data reduction.

For example, 480 VDI desktops full clones require ~18 TB of space with ADR Off compared to only ~0.83 TB when using ADR On (a data reduction ratio of 30:1).

Conclusion

This reference architecture provides guidance on how to design and implement the Hitachi Infrastructure for Citrix Workspace solution. It also describes the tests performed by Hitachi Vantara to validate and measure the performance and capabilities of the recommended solution, including third-party validated performance testing from Login VSI, the industry standard benchmarking tool for virtualized workloads.

The Hitachi Infrastructure for Citrix Workspace solution delivers a platform for enterprise end-user computing deployments using the capabilities of Hitachi Virtual Storage Platform E1090H SSD storage, Hitachi Advanced Server HA810 G2 rack servers, Brocade Fibre Channel Switches, Cisco Nexus Switches, Citrix Hypervisor, and Citrix Virtual Apps and Desktops. This is a solution designed and validated with compute, network, and storage best practices to provide performance, high availability, scalability, and flexibility for your environment.

Organizations looking to deploy enterprise-class virtual desktop infrastructure (VDI) and hosted shared desktop (HSD) solutions can confidently deploy these solutions on top of a Hitachi Infrastructure for Citrix Workspace system to ensure high performance and a positive user experience.

All performance metrics analyzed show that running Citrix VDI desktops on Hitachi Unified Compute Platform support the following:

- The ability to start small and scale out in affordable increments — from pilot to production.
- Very low latency ensuring a positive user experience for VDI or HSD.
- High virtual desktop density (desktop/sessions per host), and additional servers and storage that scale with near- linear performance.
- Independently validated, with a score rated Good or Very Good based on the Login VSI baseline performance rating for all user workload profiles tested.
- Enterprise-class data protection and resiliency.

Product descriptions

The following information describes components used in this reference architecture.

Hitachi Virtual Storage Platform E1090H

Hitachi Virtual Storage Platform E1090H supercharges business application performance with SSDs storage. It uses Hitachi Ops Center, so you can improve IT operations with the latest AI and ML capabilities. Advanced data reduction in Virtual Storage Platform E1090H enables you to run data reduction with even the most performance hungry applications.

The VSP E1090H hybrid models can be configured with both SSDs and hard disk drives (HDDs). The hybrid architecture allows for greater scalability and provides data-in-place migration support

With Virtual Storage Platform E1090H and the rest of Hitachi midrange storage family, you have agile and automated data center technology. These systems allow you to cost-effectively meet your current digital expectations and give you the ability to address future challenges, as your application data needs and service levels evolve. With time-tested, proven availability and scalability, Hitachi Vantara delivers infrastructure solutions that help you maximize your data center advantage.

Citrix

Citrix delivers people-centric solutions that power a better way to work by offering secure apps and data on any device, network or digital workspace.

- Citrix Virtual Apps and Desktops is the leading solution for applications and desktop delivery. It enables secure and remote access to Windows applications and desktops from any device, anywhere.
- Citrix Hypervisor is optimized for Citrix Virtual Apps and Desktops and simplifies your operational management, ensuring a high definition user experience for intensive workloads.

Login VSI

Login VSI (www.loginvsi.com) is the industry standard in VDI performance testing.

Login VSI offers a complete suite of proven software solutions to design, build, and safeguard the optimal performance, scalability, availability and compatibility of desktops and applications running in any type of (centralized) Windows environment, including SBC, VDI, DaaS, and fat clients.

Typical customers are enterprises with centralized desktop environments and/or business critical applications running in VDI, and all the major IT vendors that offer well-performing solutions for VDI.

The Login VSI Enterprise Edition offers a unique combination of synthetic load-testing and pro-active monitoring capabilities, allowing enterprises to design, build, and maintain VDI environments (both infrastructure and applications) that can provide, and safeguard, the optimal end-user experience.

The Login VSI load-testing solution generates a large number of synthetic users to test and protect the performance and scalability of new and existing VDI, SBC, and DaaS deployments. The Login PI active monitoring solution uses a single synthetic user running 24/7, to safeguard performance and availability of virtual desktop infrastructures and applications. The Login AT application compatibility testing solution checks the availability and health of large numbers of applications, fast and efficient.

For more information, or a free trial, please visit http://www.loginvsi.com.

Login VSI accepts no responsibility regarding this publication in any way and cannot be held accountable for any damages following from, or related to, any information contained within this publication, or any conclusions that may be drawn from it.

Cisco Nexus switches

The Cisco Nexus switch product line provides a series of solutions that make it easier to connect and manage disparate data center resources with software-defined networking (SDN). Leveraging the Cisco Unified Fabric, which unifies storage, data and networking (Ethernet/IP) services, the Nexus switches create an open, programmable network foundation built to support a virtualized data center environment.

Brocade switches from Broadcom

Brocade and Hitachi Vantara have partnered to deliver storage networking and data center solutions. These solutions reduce complexity and cost, as well as enable virtualization and cloud computing to increase business agility.

Brocade Fibre Channel switches deliver industry-leading performance, simplifying scale-out network architectures. Get the high-performance, availability, and ease of management you need for a solid foundation to grow the storage network you want.

For more information

Hitachi Vantara Global Services offers experienced storage consultants, proven methodologies and a comprehensive services portfolio to assist you in implementing Hitachi products and solutions in your environment. For more information, see the Services website.

Demonstrations and other resources are available for many Hitachi products. To schedule a live demonstration, contact a sales representative or partner. To view on-line informational resources, see the Resources website

Hitachi Academy is your education destination to acquire valuable knowledge and skills on Hitachi products and solutions. Our Hitachi Certified Professional program establishes your credibility and increases your value in the IT marketplace. For more information, see the Hitachi Vantara Training and Certification website.

For more information about Hitachi products and services, contact your sales representative, partner, or visit the Hitachi Vantara website.

Hitachi Vantara

Corporate Headquarters
2535 Augustine Drive
Santa Clara, CA 95054 USA

HitachiVantara.com/contact