

Dell EMC Ready Architectures for VDI

Designs for VMware Horizon on Dell EMC XC Family

April 2019

H17386.1

Validation Guide

Abstract

This validation guide describes the architecture and performance of the integration of VMware Horizon components for virtual desktop infrastructure (VDI) on Dell EMC XC Family devices.

Dell EMC Solutions

Copyright © 2018-2019 Dell Inc. or its subsidiaries. All rights reserved.

Published April 2019

Dell believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED “AS-IS.” DELL MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. USE, COPYING, AND DISTRIBUTION OF ANY DELL SOFTWARE DESCRIBED IN THIS PUBLICATION REQUIRES AN APPLICABLE SOFTWARE LICENSE.

Dell, EMC, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be the property of their respective owners. Published in the USA.

Dell EMC
Hopkinton, Massachusetts 01748-9103
1-508-435-1000 In North America 1-866-464-7381
www.DellEMC.com

CONTENTS

Chapter 1	Executive Summary	5
	Document purpose.....	6
	Audience.....	6
	We value your feedback.....	6
Chapter 2	Test Environment Configuration and Best Practices	7
	Validated hardware resources.....	8
	Enterprise hardware.....	8
	Graphics hardware.....	8
	Network hardware.....	8
	Validated software resources.....	9
	Validated system versions.....	9
	Virtual networking configuration.....	9
	Management server infrastructure.....	10
	NVIDIA GRID License Server.....	10
	SQL Server databases.....	11
	DNS.....	11
	High availability.....	11
	VMware Horizon architecture.....	12
Chapter 3	Solution Performance and Testing	13
	Testing process.....	14
	Resource monitoring.....	14
	Load generation.....	15
	Profiles and workloads.....	15
	Linked vs. Instant Clones.....	16
	Desktop VM test configurations.....	17
	Test results and analysis.....	18
	C7 configuration.....	19
Chapter 4	Conclusion	39
	Density recommendations.....	40
	Summary.....	40
Chapter 5	References	41
	Dell EMC documentation.....	42
	VMware documentation.....	42
	NVIDIA documentation.....	42

CONTENTS

CHAPTER 1

Executive Summary

This chapter presents the following topics:

- [Document purpose](#).....6
- [Audience](#).....6
- [We value your feedback](#).....6

Document purpose

This validation guide details the architecture, components, testing methods, and test results for Dell EMC XC Family devices with VMware Horizon 7. It includes the test environment configuration and best practices for systems that have undergone testing.

Audience

This guide is intended for architects, developers, and technical administrators of IT environments. It provides an in-depth explanation of the testing methodology and basis for VDI densities. It also validates the value of the Dell EMC Ready Architectures for VDI that deliver Microsoft Windows virtual desktops to users of VMware Horizon 7 VDI components on XC Family devices.

We value your feedback

Dell EMC and the authors of this document welcome your feedback on the solution and the solution documentation. Contact the Dell EMC Solutions team by [email](#) or provide your comments by completing our [documentation survey](#).

Authors: Dell EMC Ready Architectures for VDI Engineering Team, Donna Renfro

CHAPTER 2

Test Environment Configuration and Best Practices

This chapter presents the following topics:

- [Validated hardware resources](#)..... 8
- [Validated software resources](#)..... 9
- [Validated system versions](#)..... 9
- [Virtual networking configuration](#)..... 9
- [Management server infrastructure](#)..... 10
- [High availability](#)..... 11
- [VMware Horizon architecture](#)..... 12

Validated hardware resources

Dell EMC validated the solution with the specific hardware resources listed in this section.

Enterprise hardware

Review the test environment hardware.

We used the Dell EMC XC Family XC740xd-24 device with the components listed in the following table. We have designated the configuration as *C7*, which is referenced throughout the document.

Table 1 Validated hardware configurations

Configur ation	Enterprise platform	CPU	Memory	RAID controller	HD configuration	Network
C7	XC740xd-24	6138 Gold (20- core 2.0 GHz)	768 GB @ 2,666 MT/s	HBA 330	<ul style="list-style-type: none"> • 2 x 240 GB M.2 • 2 x 960 GB SSD • 4 x 1.8 TB HDD 	2 x Mellanox Connect X-4 LX 25 GbE SFP Rack NDC

Graphics hardware

We used the following NVIDIA GPU hardware in our tests for graphics-intensive workloads:

- NVIDIA Pascal P40—A dual-slot 10.5-inch PCI Express Gen3 graphics card featuring a single high-end NVIDIA Pascal GPU and a total of 24 GB GDDR5 memory per card

Network hardware

We used the following network hardware in our test environment:

- Dell Networking S4048 (10 GbE ToR switch)—A high-density, ultra-low-latency ToR switch that features 48 x 10 GbE SFP+ and 6 x 40 GbE ports and up to 720 Gbps switch fabric capacity
- Dell Networking S5248 (25 GbE ToR switch)—A high-density, high performance, open networking ToR switch that features 48 x 25 GbE SFP28, 4 x 100 GbE QFSP28 ports, 2 x 100 GbE QFSP28-DD ports and up to 2.0 Tbs switch fabric capacity

Validated software resources

Dell EMC validated this solution with the software components listed in the following table.

Component	Description/Version
Hypervisor	ESXi 6.7
Broker technology	VMware Horizon 7 version 7.7
Broker database	Microsoft SQL Server 2016
Management VM OS	Microsoft Windows Server 2016 (Connection Server & DB)
Virtual desktop OS	Microsoft Windows 10 Enterprise
Office application suite	Microsoft Office Professional 2016
Login VSI test suite	Version 4.1.32.1
Platform	Nutanix AOS version 5.10.0.2
NVIDIA GRID software (for graphics testing)	7.1

Validated system versions

Dell EMC validated this solution using the system versions that are listed in the following table.

Table 2 Version matrix for tested system

Server configuration	Nvidia vGPU version	Hyper-visor	Hyper-visor version	Hyper-visor build	Bios	AOS version	Windows 10 version	Windows 10 patches
C7	N/A	ESXi	6.7	10764712	1.6.11	5.100.2	1803 OS Build 17134.523	KB44100347
C7 + 3 x P40	7.1							KB4480966

Virtual networking configuration

The network configuration for the Dell EMC XC Family devices uses a 25 Gb converged infrastructure model.

All required VLANs traverse two 25 Gb NICs configured in an Active/Active team. For larger scaling, we recommend that you separate the infrastructure management virtual machines (VMs) from the compute VMs to aid in predictable compute host scaling.

We used the following VLAN configurations for the compute hosts, management hosts, and iDRAC in this solution model:

- Compute hosts
 - Management VLAN: Configured for hypervisor infrastructure traffic—L3 routed by using the spine layer
 - Live Migration VLAN: Configured for Live Migration traffic—L2 switched by using the leaf layer
 - VDI VLAN: Configured for VDI session traffic—L3 routed by using the spine layer
- Management hosts
 - Management VLAN: Configured for hypervisor management traffic—L3 routed by using the spine layer
 - Live Migration VLAN: Configured for Live Migration traffic—L2 switched by using the leaf layer
 - VDI Management VLAN: Configured for VDI infrastructure traffic—L3 routed by using the spine layer
- VLAN iDRAC: Configured for all hardware management traffic—L3 routed by using the spine layer

Management server infrastructure

The following table lists the sizing requirements for the management server components.

Table 3 Sizing for XC Family devices, RDSH, and NVIDIA GRID license server (optional)

Component	vCPU	RAM (GB)	NIC	OS + data vDisk (GB)	Tier 2 volume (GB)
VMware vCenter Appliance	2	16	1	290	
Platform Services Controller	2	2	1	30	
Horizon Connection Server	4	16	1	40	
SQL Server	4	8	1	40	210 (VMDK)
File server	1	4	1	40	2,048 (VMDK)
Nutanix CVM	8	32	3	0	
RDSH VM	8	32	1	80	
NVIDIA GRID License Server	2	4	1	40 + 5	

NVIDIA GRID License Server

When using NVIDIA vGPU cards, graphics-enabled VMs must obtain a license from a GRID License Server on your network to be entitled for vGPU.

We installed the GRID License Server software on a system running a Windows 2016 operating system to test vGPU configurations.

We made the following changes to the GRID License Server to address licensing requirements:

- Used a reserved fixed IP address
- Configured a single MAC address
- Applied time synchronization to all hosts on the same network

SQL Server databases

During validation, a single dedicated SQL Server 2016 VM hosted the VMware databases in the management layer. We separated SQL data, logs, and tempdb into their respective volumes, and created a single database for Horizon Connection Server.

We adhered to VMware best practices for this testing, including alignment of disks to be used by SQL Server with a 1,024 KB offset and formatted with a 64 KB file allocation unit size (data, logs, and tempdb).

DNS

DNS is the basis for Microsoft Active Directory and also controls access to various software components for VMware services. All hosts, VMs, and consumable software components must have a presence in DNS. We used a dynamic namespace integrated with Active Directory and adhered to Microsoft best practices.

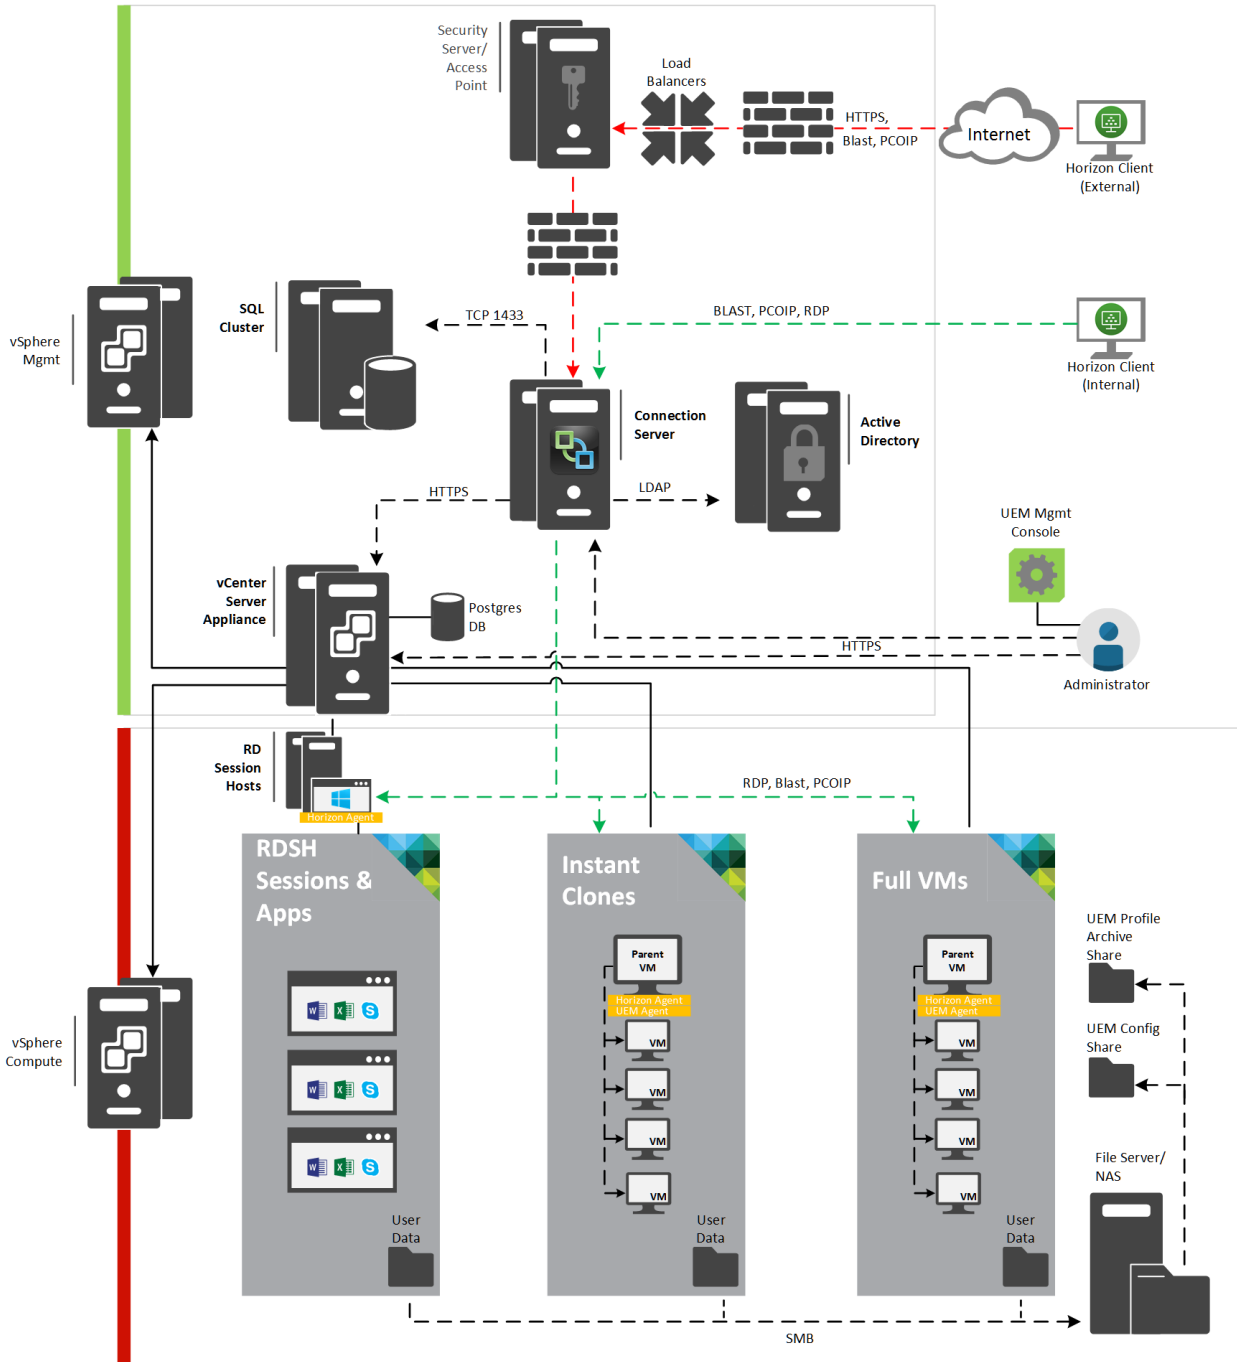
High availability

Although we did not enable high availability (HA) during the validation that is documented in this guide, we strongly recommend that HA be factored into any VDI design and deployment. This process involves following the N+1 model with redundancy at both the hardware and software layers. The design guide for this architecture provides additional recommendations for HA.

VMware Horizon architecture

The following figure shows the Horizon communication flow.

Figure 1 VMware Horizon architecture



CHAPTER 3

Solution Performance and Testing

This chapter presents the following topics:

- [Testing process](#)..... 14
- [Test results and analysis](#)..... 18

Testing process

To ensure the optimal combination of end-user experience (EUE) and cost-per-user, we conducted performance analysis and characterization testing on this solution using the Login VSI load-generation tool. Login VSI is a carefully designed, holistic methodology that monitors both hardware resource utilization parameters and EUE during load-testing.

We tested each user load against four runs: a pilot run to validate that the infrastructure was functioning and valid data could be captured, and three subsequent runs to enable data correlation.

During testing, while the environment was under load, we logged in to a session and completed tasks that correspond to the user workload. While this test is subjective, it helps to provide a better understanding of the EUE in the desktop sessions, particularly under high load. It also helps to ensure reliable data gathering.

Resource monitoring

To ensure that the user experience was not compromised, we monitored the following important resources:

- **Compute host servers**—VMware vCenter (for VMware vSphere-based solutions) or Microsoft Performance Monitor (for Hyper-V-based solutions) gathers key data (CPU, memory, disk, and network usage) from each of the compute hosts during each test run. This data is exported to .csv files for single hosts, and then consolidated to show data from all hosts. While the report does not include specific performance metrics for the management host servers, these servers are monitored during testing to ensure that they are performing at an expected level with no bottlenecks.
- **Hardware resources**—Resource contention, which occurs when hardware resources have been exhausted, can cause poor EUE. We monitored the relevant resource utilization parameters and applied relatively conservative thresholds, as shown in the following table. Thresholds are carefully selected to deliver an optimal combination of good EUE and cost-per-user while also providing burst capacity for seasonal or intermittent spikes in usage.

Table 4 Parameter pass/fail thresholds

Parameter	Pass/fail threshold
Physical host CPU utilization	85% ^a
Physical host memory utilization	85%
Network throughput	85%
Storage I/O latency	20 ms
Login VSI Failed Session	2%

- a. The Ready Solutions for VDI team recommends that average CPU utilization not exceed 85% in a production environment. A 5% margin of error was allocated for this validation effort. Therefore, CPU utilization sometimes exceeds our recommended percentage. Because of the nature of LoginVSI testing, these exceptions are reasonable for determining our sizing guidance.

- **GPU resources**—vSphere Client monitoring collects data about the GPU resource use from a script that is run on ESXi 6.7 and later hosts. The script runs for the duration of the test and contains NVIDIA System Management Interface commands. The commands query each GPU and log the GPU processor, temperature, and memory use to a .csv file.

Load generation

Login VSI from Login VSI, Inc. is the industry-standard tool for testing VDI environments and Remote Desktop Session Host (RDSH) environments.

Login VSI installs a standard collection of desktop application software (for example, Microsoft Office, Adobe Acrobat Reader) on each VDI desktop. It then uses launcher systems to connect a specified number of users to available desktops within the environment. When the user is connected, a logon script starts the workload, configures the user environment, and starts the test script. Each launcher system can launch connections to a number of VDI desktops (target machines). A centralized management console configures and manages the launchers and the Login VSI environment.

In addition, we used the following login and boot paradigm:

- Users were logged in within a login timeframe of 1 hour, except when testing low-density solutions such as GPU/graphic-based configurations, in which users were logged in every 10 to 15 seconds.
- All desktops were started before users logged in.
- All desktops ran an industry-standard anti-virus solution. Windows 10 machines used Windows Defender.

Profiles and workloads

Machine profiles and user workloads determine the density numbers that the solution can support. Each profile and workload is bound by specific metrics and capabilities, with two targeted at graphics-intensive use cases.

Profiles and workloads are defined as follows:

- **Profile**—The configuration of the virtual desktop; the number of vCPUs and the amount of RAM that is configured on the desktop and available to the user
- **Workload**—The set of applications that is used

We load-tested two profiles by using a workload that is representative of the profile. The following table describes each use case.

Table 5 Virtual desktop profiles and workloads

Profile name/ workload	Workload description
Task worker	The least intensive of the standard workloads. This workload primarily runs Microsoft Excel and Microsoft Internet Explorer, with some minimal Microsoft Word activity, as well as Microsoft Outlook, Adobe, and copy and zip actions. The applications are started and stopped infrequently, which results in lower CPU, memory, and disk I/O usage.
Knowledge worker	Designed for virtual machines with 2 vCPUs. This workload includes the following activities: <ul style="list-style-type: none"> • Outlook—Browse messages.

Table 5 Virtual desktop profiles and workloads (continued)

Profile name/ workload	Workload description
	<ul style="list-style-type: none"> • Internet Explorer—Browse websites and open a YouTube style video (480p movie trailer) three times in every loop. • Word—Start one instance to measure response time and another to review and edit a document. • Doro PDF Printer and Acrobat Reader—Print a Word document and export it to PDF. • Excel—Open a large randomized sheet. • PowerPoint—Review and edit a presentation. • FreeMind—Run a Java-based Mind Mapping application. • Other—Perform various copy and zip actions.
Power worker	<p>The most intensive of the standard workloads. The following activities are performed with this workload:</p> <ul style="list-style-type: none"> • Begin by opening four instances of Internet Explorer and two instances of Adobe Reader, which remain open throughout the workload. • Perform more PDF printer actions than in the other workloads. • Watch a 720p and a 1080p video. • Reduce the idle time to two minutes. • Perform various copy and zip actions.
Graphics performance configuration/ multimedia	<p>A workload that is designed to heavily stress the CPU when using software graphics acceleration. GPU-accelerated computing offloads the most compute-intensive sections of an application to the GPU while the CPU processes the remaining code. This modified workload uses the following applications for its GPU/CPU-intensive operations:</p> <ul style="list-style-type: none"> • Adobe Acrobat • Google Chrome • Google Earth • Microsoft Excel • HTML5 3D spinning balls • Internet Explorer • MP3 • Microsoft Outlook • Microsoft PowerPoint • Microsoft Word • Streaming video



Linked vs. Instant Clones

Horizon supports two provisioning methods that deliver space-optimized virtual desktop pools: Linked Clones and Instant Clones. The user density per host is not

impacted by using one over the other. The differences in the test graphs for these two methods are a result of the following processes:

- For Linked Clones, all the VMs are rebooted before the test starts to simulate a boot storm.
- For Instant Clones, the VMs are rebooted after the session logs off, because when a user logs out of the instant clone, the clone is destroyed and recreated for the next user.

The following figure shows the differences.



Desktop VM test configurations

The following table summarizes the compute VM configurations for the profiles and workloads that we tested.

Table 6 Desktop VM specifications

User profile	vCPUs	ESXi configured memory	ESXi reserved memory	Screen resolution	Operating system
Task worker	2 ^a	3 GB	1.5 GB	1280 x 720	Windows 10 Enterprise 64-bit
Knowledge worker	2	4 GB	2 GB	1920 x 1080	Windows 10 Enterprise 64-bit
Power worker	4	8 GB	4 GB	1920 x 1080	Windows 10 Enterprise 64-bit
Multimedia	4	8 GB	8 GB	1920 x 1080	Windows 10 Enterprise 64-bit

Table 6 Desktop VM specifications (continued)

- a. Dell EMC has validated the LoginVSI Task worker workload with two vCPUs assigned per VM, although LoginVSI lists the typical VM vCPU profile for this workload as being a single vCPU. Dell EMC diverges from this definition to deliver virtual desktops with great user experience. Increasing the vCPU count to 2 in the vCPU profile associated with the Task worker workload does have a minor impact on densities but generates improved user experience in return.

Test results and analysis

We used the Login VSI test suite to simulate the user experience for several profile types under the typical workload for that type. The following table summarizes the test results for the compute hosts using the various workloads and configurations.

Note

Dell EMC is aware of the vulnerabilities known as Meltdown, Spectre, and Foreshadow/L1TF, which affect many modern microprocessors. Ensure that you read the information in the following links:

- <http://www.dell.com/support/article/SLN308588>
- <http://www.dell.com/support/article/SLN308587>
- <http://dell.com/support/L1-terminal-fault>

Consider this information in combination with the version information in [Table 2](#) on page 9 to understand the vulnerability mitigation status of the environment we used to derive the test results shown in the following table.

Table 7 Test results summary

Server Configuration	Workload	User Density	Avg CPU	Avg Mem Active	Avg IOPS / User
C7	Task Worker	120	86.5%	122 GB	2.7
C7	Knowledge Worker	90	84%	125 GB	3.5
C7	Power Worker	65	84%	146 GB	4.5
C7 + 3 x P40	Multimedia (Virtual PC: P40-1B)	40	87% ^a	360 GB	5.5

- a. The CPU utilization threshold of 85% is relaxed when testing with graphics cards. This test represents maximum utilization of the graphical resources available to the system as well as full user concurrency. Ideally, in a production environment, you would decrease the user density slightly or use higher bin processors to bring the CPU utilization closer to the 85% threshold. All LoginVSI tests completed successfully without reaching VSI maximum, indicating that user experience was good.

The table headings are defined as follows:

- **User density**—The number of users per compute host that successfully completed the workload test within the acceptable resource limits for the host. For clusters, this number reflects the average of the density achieved for all compute hosts in the cluster.
- **Avg CPU**—The average CPU usage over the steady-state period. For clusters, this number represents the combined average CPU usage of all compute hosts. On the latest Intel processors, the ESXi host CPU metrics exceed the rated 100 percent for the host if Turbo Boost is enabled, which is the default setting. An additional 35 percent of CPU is available from the Turbo Boost feature, but this additional CPU headroom is not reflected in the VMware vSphere metrics where

the performance data is gathered. Therefore, CPU usage for ESXi hosts is adjusted and each CPU graph includes a line indicating the potential performance headroom that is provided by Turbo boost.

- **Avg active memory**—For ESXi hosts, the amount of memory that is actively used, as estimated by the VMkernel based on recently touched memory pages. For clusters, this is the average amount of guest physical memory that is actively used across all compute hosts over the steady-state period.
- **Avg IOPS per user**—IOPS calculated from the average disk IOPS over the steady state period divided by the number of users.

C7 configuration

We performed the following multimedia performance testing on the XC740xd system in the C7 configuration described in [Validated hardware resources](#) on page 8.

Notes:


- We deployed Horizon and vSphere management roles within the cluster on a single host that also hosted desktops. This optional configuration is beneficial for POCs or small deployments looking to maximize user density.
- We allocated 8 vCPUs and 32 GB of memory to the Nutanix Controller VM (CVM) when we configured the Nutanix cluster.
- We set the NVIDIA Pascal P40 GPU accelerator card with vGPU scheduling policy to **Fixed Share Scheduler**.

Task Worker, 360 Users, ESXi 6.7, Horizon 7.7

We ran the following tests on this workload.

CPU



The graph shows the performance data for 360 user sessions on three compute hosts. The CPU reached a steady state average of 86% during the test cycle when all users were logged on to the compute host. The management host reached a CPU max of 22 percent during logoff and averaged 16 percent during steady state.



Memory


No memory consumption constraints occurred on the compute hosts. Of a total of 768 GB available memory per node, the compute host reached a maximum memory

consumption of 399 GB, with active memory usage reaching a max of 134 GB during the Steady State phase. Each user session consumed 1.1 GB of memory and .34 GB of active memory. No memory ballooning or swapping occurred on any host.




Network usage

Network bandwidth was not an issue on this test run with a steady state peak of approximately 524 Mbps. The busiest period for network traffic was during the recreate phase. The steady state average was 472 Mbps. Each desktop produced network throughput of 1.3 Mbps in steady state.




IOPS

The Cluster IOPS reached a maximum of 6,050 Disk IOPS during the recreate clones and averaged 998 Disk IOPS during steady state. Based on these numbers, each user session generated 2.7 IOPS in steady state.




The Cluster Controller IOPS reached a maximum of 41,820 disk IOPS during the recreate clones and averaged 2,909 disk IOPS during Steady State. Based on these numbers each user session generated eight IOPS in Steady State.




I/O latency

The peak cluster I/O latency was 3.49 ms during recreate clones and the average cluster I/O latency during steady state was 0.37 ms. The chart clearly shows a very steady and very low level of I/O Latency throughout the test run.



User experience

The Login VSI Max user experience score shown below for this test was not reached. When manually interacting with the sessions during steady state the mouse and window movement was responsive and video playback was good. The baseline performance of 929 indicated that the user experience for this test run was very good. The index average reached 1,122, which was well below the VSI maximum threshold of 1,929.




Knowledge Worker, 270 Users, ESXi 6.7, Horizon 7.7

We ran the following tests on this workload.



CPU

The following graph shows the performance data for 270 user sessions on three compute hosts. The CPU reaches a steady state average of 84 percent during the test cycle when all users were logged on to compute host. The management host reached a CPU maximum of 16 percent during logoff and averaged nine percent during steady state.




Memory

No memory consumption constraints occurred on the compute hosts. Of a total of 768 GB available memory per node, the compute host reached a maximum memory consumption of 364 GB, with active memory usage reaching a maximum of 134 GB during the Steady State phase. Each user session consumed 1.3 GB of memory and .46 GB of active memory. No memory ballooning or swapping occurred on any host.




Network usage

Network usage peaked at 647 Mbps during boot storm on one host, and the average network usage for all hosts was 597 Mbps during steady state. Each desktop produced network throughput of 2.1 Mbps in steady state.




IOPS

The Cluster IOPS reached a maximum of 7,185 disk IOPS during the Recreate Clones and averaged 936 Disk IOPS during steady state. Based on these numbers each user session generated 3.5 IOPS in steady state.




The Cluster Controller IOPS reached a maximum of 19,926 disk IOPS during the recreate clones and averaged 3,708 disk IOPS during Steady State. Based on these numbers each user session generated 14 IOPS in Steady State.




I/O latency

The peak cluster I/O latency was 0.6 ms during logon and the average cluster I/O latency during steady state was 0.41 ms. The chart clearly shows a very steady and very low level of I/O Latency throughout the test run.



User experience

The Login VSI Max user experience score shown below for this test was not reached. When manually interacting with the sessions during steady state the mouse and window movement was responsive and video playback was good. The baseline performance of 942 indicated that the user experience for this test run was very good. The index average reached 1,187, which was well below the VSI maximum threshold of 1,943.




Power worker, 195 users, ESXi 6.7, Horizon 7.7

We ran the following tests on this workload.


CPU usage

The following graph shows the performance data for 195 user sessions on three compute hosts. The CPU reached a steady state average of 84 percent during the test cycle when all users were logged on to the compute host. The management host reached a CPU max of 16 percent during logoff and averaged 9 percent during steady state.




Memory

No memory consumption constraints occurred on the compute hosts. Of a total of 768 GB available memory per node, the compute host reached a maximum memory consumption of 504 GB with active memory usage reaching a max of 173 GB during the Steady State phase. Each user session consumed 2.1 GB of memory and .75 GB of active memory. No memory ballooning or swapping occurred on any host.



Network usage


Network usage peaked at 762 Mbps on one host during the steady state phase, and the average network usage for all hosts during steady state was 629 Mbps. Each desktop accounted for 3.2 Mbps in steady state.




IOPS

The IOPS graphs and numbers came from the Nutanix Prism web console. They clearly display the initial logon of the desktops, the steady state, and finally the logoff phase. The graphs show IOPS data for the cluster on the GPU host.

The cluster reached a maximum of 3,299 disk IOPS during the boot storm and averaged 885 IOPS during steady state. Based on these numbers, each user session generated 4.5 disk IOPS during steady state.




The cluster controller reached a maximum of 27,376 disk IOPS during the recreate clones and averaged 3,602 disk IOPS during steady state. Based on these numbers, each user session generated 18.4 IOPS in steady state.



I/O latency


The latency graphs and latency numbers came from the Nutanix Prism web console. They display the initial logon of the desktops, the steady state, and then the logoff

phase. The graphs show latency data for the cluster and the GPU host. The cluster latency reached a maximum of .77 ms during the logoff and averaged 0.46 ms during steady state.



User experience

The Login VSI Max user experience score shown below for this test was not reached. When manually interacting with the sessions during steady state the mouse and window movement was responsive and video playback was good. The baseline performance of 884 indicated that the user experience for this test run was very good. The index average reached 1,127, which was well below the threshold of 1,884.




Multimedia workload, 40 vGPU users, ESXi 6.7, Horizon 7.7

We ran the following tests on this workload.

CPU usage


The GPU-enabled compute host was populated with 40 vGPU-enabled virtual machines and used the NVIDIA P40-1B profile. With all user virtual machines powered on and before starting the test, the CPU usage was approximately 8 percent on the GPU-enabled compute host.

The following figure shows the performance data for 40 user sessions on the management and GPU-enabled compute hosts. The CPU reached a steady state average of 87 percent during the test cycle when all users were logged on to the GPU-enabled compute host. The management host reached a CPU maximum of 11 percent during steady state.





GPU usage

We gathered the GPU metrics from the vSphere Web Client. The GPU usage during the steady state period averaged approximately 23 percent and reached a peak usage of 34.3 percent with the multimedia workload.




Memory

With regard to memory consumption for this test run, there were no constraints on the management or GPU-enabled compute hosts. Of a total of 768 GB available memory per node, the GPU compute host reached a maximum memory consumption of 374 GB with active memory usage reaching a maximum of 360 GB during the steady state phase. Each user session consumed 9.3 GB of memory and 89.0 GB of active memory. No variations in memory usage occurred throughout the test because all vGPU-enabled VM memory was reserved. No memory ballooning or swapping occurred on either host.



Network usage

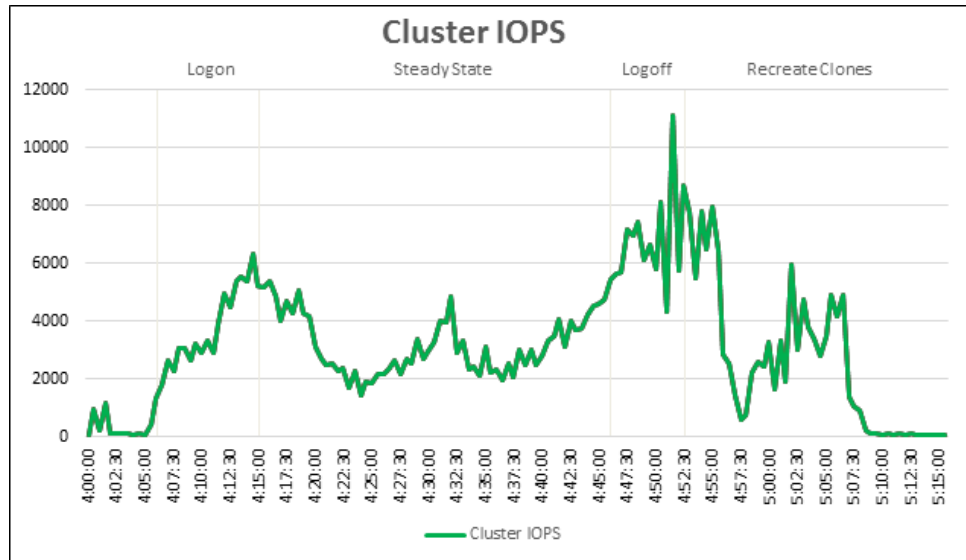
Network bandwidth was not an issue on this test run, with a steady state peak of approximately 854 Mbps. The busiest period for network traffic was during the logoff phase. The compute / GPU host reached a peak of 866 Mbps at the start of logoff.




IOPS

The IOPS graphs and numbers came from the Nutanix Prism web console. They clearly display the initial logon of the desktops, the steady state, and finally the logoff phase. The graphs show IOPS data for the cluster on the GPU host.

The cluster reached a maximum of 2,557 disk IOPS during the boot storm and averaged 220 IOPS during steady state. Based on these numbers, each user session generated 5.5 disk IOPS during steady state.




The cluster controller reached a maximum of 11,132 disk IOPS during the boot storm and averaged 3,239 disk IOPS during steady state. Based on these numbers, each user session generated 80 IOPS in steady state.



GPU host disk IOPS


The GPU host reached a maximum of 2,557 disk IOPS during the boot storm and averaged 220 disk IOPS during steady state. Based on these numbers, each user session generated 5.5 IOPS in steady state.




Latency

The latency graphs and latency numbers came from the Nutanix Prism web console. They display the initial logon of the desktops, the steady state, and then the logoff phase. The graphs show latency data for the cluster and the GPU host.

The cluster latency reached a maximum of 93 ms during the boot storm and averaged 0.47 ms during steady state.




The GPU host latency reached a maximum of 1.8 ms during the boot storm and averaged 0.647 ms during steady state.



User experience

The following figure shows that the user experience score did not reach the Login VSI maximum for this test. When manually interacting with the sessions during steady state the mouse and window movement was responsive and video playback was good. The baseline performance of 911 indicated that the user experience for this test run was good. The Index average reached 1,569, which was well below the threshold of 1,912.



Notes

No disk latency issues occurred during testing.

CHAPTER 4

Conclusion

- [Density recommendations](#).....40
- [Summary](#).....40

Density recommendations

We tested all configurations with Microsoft Windows 10 and Microsoft Office 2016. Test results provide recommended user densities, as shown in the following table.

Table 8 User density recommendations

Server configuration	Workload	User density
C7	Task worker	120
C7	Knowledge worker	90
C7	Power worker	65
C7 + 3 x P40	Multimedia (virtual PC: P40-1B)	40

Summary

The configurations for the XC Family devices have been optimized for VDI. We selected the memory and CPU configurations that provide optimal performance. You can change these configurations to meet your own requirements, but keep in mind that changing the memory and CPU configurations from those that have been validated in this document affects the user density per host. We applied mitigations to the Spectre and LT1F vulnerabilities during the validation.

With the introduction of the six-channels-per-CPU requirement for Skylake, the Density Optimized memory configuration recommendation has increased from the previous guidance of 512 GB to 768 GB. This change was necessary to ensure a balanced memory configuration and optimized performance for your VDI solution. The additional memory is advantageous, considering the resulting increase in operating system resource utilization and the enhanced experience for users when they have access to additional memory allocations.

CHAPTER 5

References

This chapter presents the following topics:

- [Dell EMC documentation](#).....42
- [VMware documentation](#)..... 42
- [NVIDIA documentation](#).....42

Dell EMC documentation

The following Dell EMC documentation provides additional and relevant information. Access to these documents depends on your login credentials. If you do not have access to a document, contact your Dell EMC representative. Also see the [Dell EMC VDI Information Hub](#) for a complete list of VDI resources.

- [Dell EMC Virtual Desktop Infrastructure](#)
- [Dell EMC XC Series and XC Core Technical Resource Center](#)

This document is part of the documentation set for this architecture, which includes the following:

- [Dell EMC Ready Architectures for VDI: Designs for VMware Horizon on XC Family Design Guide](#)
- [Dell EMC Ready Architectures for VDI: Designs for VMware Horizon on XC Family Deployment Guide](#)
- [Dell EMC Ready Architectures for VDI: Designs for VMware Horizon on XC Family Validation Guide](#)

VMware documentation

The following VMware documentation provides additional and relevant information:

- [VMware vSphere documentation](#)
- [VMware Horizon 7 documentation](#)
- [Best Practices for Published Application and Desktops in VMware Horizon Apps and VMware Horizon 7](#)
- [VMware Compatibility Guide](#)
- [Horizon 7 Enterprise Edition Reference Architecture](#)
- [Horizon 7 Enterprise Edition Multi-Site Reference Architecture](#)

NVIDIA documentation

The following NVIDIA documentation provides additional and relevant information:

- [NVIDIA Virtual GPU Software Quick Start Guide](#)