

REFERENCE ARCHITECTURE

Dell EMC Ready System for VDI on XC Series

Integration of VMware Horizon with Dell EMC XC Series Hyper-Converged Appliances

[Abstract](#)

A Reference Architecture for integrating Dell EMC XC Series Hyper-Converged Appliances and VMware Horizon brokering software on VMware ESXi hypervisor to create virtual application and virtual desktop environments on 14th generation Dell EMC PowerEdge Servers.

January 2018

Revisions

Date	Description
January 2018	Initial release

Acknowledgements

This paper was produced by the following members of the Dell EMC storage engineering team:

Authors: Peter Fine – Chief Architect

Geoff Dillon – Sr. Solutions Engineer

Andrew Breedy – Sr. Solutions Engineer

Jonathan Chamberlain – Solution Engineer

Support: David Hulama – Sr. Technical Marketing Advisor

The information in this publication is provided “as is.” Dell Inc. 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 software described in this publication requires an applicable software license.

© 2018 Dell Inc. or its subsidiaries. All Rights Reserved. Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners.

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

Table of contents

Revisions.....	2
Acknowledgements.....	2
Executive summary.....	5
1 Introduction.....	6
1.1 Objective.....	6
1.2 What's new	6
2 Solution architecture overview	7
2.1 Introduction	7
2.2 Dell EMC XC Series Hyper-Converged appliances	7
2.3 Distributed Storage Fabric.....	8
2.4 App Mobility Fabric	8
2.4.1 Nutanix architecture.....	8
2.5 Nutanix Hyper-Converged Infrastructure.....	10
2.6 Nutanix all-flash	12
2.7 Dell EMC XC Series - VDI solution architecture.....	12
2.7.1 Networking.....	12
2.7.2 XC Series – Enterprise solution pods.....	13
3 Hardware components	16
3.1 Network.....	16
3.1.1 Dell Networking S3048 (1Gb ToR switch).....	16
3.1.2 Dell Networking S4048 (10Gb ToR switch).....	17
3.2 Dell EMC XC Series Hyper-Converged appliances	18
3.2.1 Dell EMC XC640.....	19
3.2.2 Dell EMC XC740xd.....	22
3.3 NVIDIA Tesla GPUs	24
3.3.1 NVIDIA Tesla M10	24
3.3.2 NVIDIA Tesla M60	25
3.4 Dell Wyse Endpoints	26
3.4.1 Wyse 3040 Thin Client (ThinOS, ThinLinux)	26
3.4.2 Wyse 5040 AIO Thin Client (ThinOS).....	26
3.4.3 Wyse 5060 Thin Client (ThinOS, ThinLinux, WES7P, WIE10)	26
3.4.4 Wyse 7020 Thin Client (WES 7/7P/8, WIE10, ThinLinux).....	27
3.4.5 Wyse 7040 Thin Client (WES7P, WIE10)	27
4 Software components.....	28
4.1 VMware.....	28

4.1.1	VMware Horizon 7	28
4.1.2	VMware vSphere 6	29
4.2	Microsoft RDSH	29
4.2.1	NUMA architecture considerations	29
4.3	NVIDIA GRID vGPU	33
4.3.1	vGPU profiles.....	33
5	Solution architecture for Horizon	40
5.1	Management role configuration	40
5.1.1	VMware Horizon management role requirements	40
5.1.2	RDSH on vSphere	40
5.1.3	NVIDIA GRID license server requirements	41
5.1.4	SQL databases	41
5.1.5	DNS	42
5.2	Storage architecture overview	42
5.2.1	Nutanix containers	43
5.3	Virtual networking	44
5.3.1	vSphere	44
5.4	Scaling guidance	45
5.5	Solution high availability	47
5.6	Communication flow for Horizon.....	48
6	Solution performance and testing.....	49
6.1	Summary	49
6.2	Test and performance analysis methodology.....	49
6.2.1	Testing process	49
6.2.2	Resource monitoring	51
6.2.3	Resource utilization	52
6.3	Test configuration details.....	52
6.3.1	Compute VM configurations	53
6.4	Standard VDI test results and analysis	54
6.4.1	XC740xd-C7	56
6.5	vGPU test results and analysis.....	65
6.5.1	XC740xd-C7 with Tesla M60.....	67
A	Related resources	81

Executive summary

This document provides the reference architecture for integrating Dell EMC XC Series Hyper-Converged Appliances and VMware Horizon software to create virtual application and virtual desktop environments.

The Dell EMC XC Series is a hyper-converged solution that combines storage, compute, networking, and virtualization using industry-proven Dell EMC PowerEdge™ server technology and Nutanix software. By combining the hardware resources from each appliance into a shared-everything model for simplified operations, improved agility, and greater flexibility, Dell EMC and Nutanix together deliver simple, cost-effective solutions for enterprise workloads.

VMware Horizon provides a complete end-to-end virtualization solution delivering Microsoft Windows virtual desktops or server-based hosted shared sessions to users on a wide variety of endpoint devices.

1 Introduction

This document addresses the architecture design, configuration and implementation considerations for the key components required to deliver virtual desktops or shared sessions via VMware Horizon® on VMware vSphere® 6 running on the Dell EMC XC Series Hyper-Converged infrastructure platform.

For manuals, support info, tools, and videos, please visit: www.Dell.com/xcseriesmanuals.

1.1 Objective

Relative to delivering the virtual desktop environment, the objectives of this document are to:

- Define the detailed technical design for the solution.
- Define the hardware requirements to support the design.
- Define the constraints which are relevant to the design.
- Define relevant risks, issues, assumptions and concessions – referencing existing ones where possible.
- Provide a breakdown of the design into key elements such that the reader receives an incremental or modular explanation of the design.
- Provide solution scaling and component selection guidance.

1.2 What's new

- XC Series Appliances launched on 14th generation Dell EMC PowerEdge platforms
- NVIDIA Tesla M60 and vGPU testing

2 Solution architecture overview

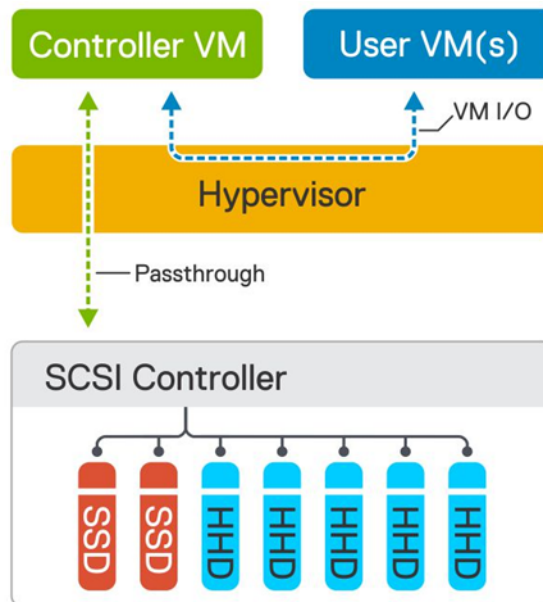
2.1 Introduction

Dell EMC customers benefit in leveraging this integrated solution for their primary workload data protection needs. This integrated solution offers Virtual Machine (VM) deployment and lifecycle management for the combined solution offering. Protection for newly deployed and existing VMs. Usage of policies and best practices and the consequent streamlining of the data protection workflow are the primary goals for this solution. This section will provide an overview of the products used to validate the solution.

2.2 Dell EMC XC Series Hyper-Converged appliances

Dell EMC XC Series hyper-converged appliances start with the proven Dell EMC PowerEdge 14th generation server platform and incorporate many of the advanced software technologies that power leading web-scale and cloud infrastructures. Backed by Dell EMC global service and support, these 1- and 2U appliances are preconfigured for specific virtualized workloads, and are designed to maintain data availability in case of node and disk failure.

The XC Series infrastructure is a scalable cluster of high-performance appliances, or servers, each running a standard hypervisor and containing processors, memory, and local storage (consisting of solid state disk (SSD) flash for high performance and high-capacity disk drives), hybrid or all-flash. Each appliance runs virtual machines just like a standard hypervisor host as displayed below.



2.3 Distributed Storage Fabric

The Distributed Storage Fabric (DSF) delivers enterprise data storage as an on-demand service by employing a highly distributed software architecture. Nutanix eliminates the need for traditional SAN and NAS solutions while delivering a rich set of VM-centric software-defined services. Specifically, the DSF handles the data path of such features as snapshots, clones, high availability, disaster recovery, deduplication, compression, and erasure coding.

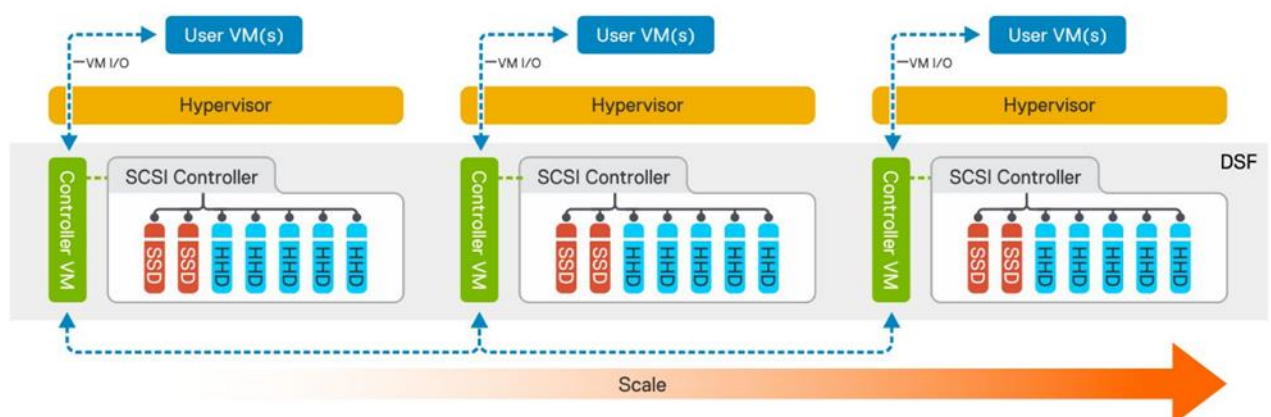
The DSF operates via an interconnected network of Controller VMs (CVMs) that form a Nutanix cluster, and every node in the cluster has access to data from shared SSD, HDD, and cloud resources. The hypervisors and the DSF communicate using the industry-standard NFS, iSCSI, and SMB3 protocols, depending on the hypervisor in use.

2.4 App Mobility Fabric

The App Mobility Fabric (AMF) collects powerful technologies that give IT professionals the freedom to choose the best environment for their enterprise applications. The AMF encompasses a broad range of capabilities for allowing applications and data to move freely between runtime environments, including between Nutanix systems supporting different hypervisors, and from Nutanix to public clouds. When VMs can migrate between hypervisors, administrators can host production and development or test environments concurrently on different hypervisors and shift workloads between them as needed. AMF is implemented via a distributed, scale-out service that runs inside the CVM on every node within a Nutanix cluster.

2.4.1 Nutanix architecture

Nutanix software provides a hyper-converged platform that uses DSF to share and present local storage to server nodes within a cluster while creating a clustered volume namespace accessible to all nodes. The figure below shows an overview of the Nutanix architecture including, user VMs, the Nutanix storage CVM, and its local disk devices. Each CVM connects directly to the local storage controller and its associated disks. Using local storage controllers on each host localizes access to data through the DSF, thereby reducing storage I/O latency. The DSF replicates writes synchronously to at least one other XC Series node in the system, distributing data throughout the cluster for resiliency and availability. Replication factor 2 (RF2) creates two identical data copies in the cluster, and replication factor 3 (RF3) creates three identical data copies.



DSF virtualizes local storage from all appliances into a unified pool. DSF uses local SSDs and capacity disks from all appliances to store virtual machine data. Virtual machines running on the cluster write data to DSF as if they were writing to local storage. Nutanix data locality ensures that the XC Series node providing CPU and

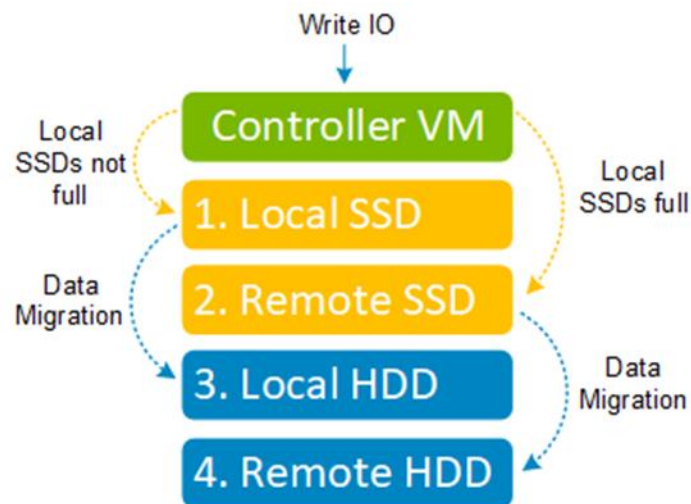
memory to a VM also provides its disk as well, thus minimizing IO that must cross the network. XC Series supports multiple hypervisors and provides choice and flexibility to customer.

XC Series offers customer choice of hypervisors without being locked-in. The hypervisors covered in this reference architecture are:

- VMware® ESXi®

In addition, the solution includes the Nutanix Controller VM (CVM), which runs the Nutanix software and serves I/O operations for the hypervisor and all VMs running on that host. Each CVM connects directly to the local storage controller and its associated disks thereby reducing the storage I/O latency. The data locality feature ensures virtual machine I/Os are served by the local CVM on the same hypervisor appliance, improving the VM I/O performance regardless of where it runs.

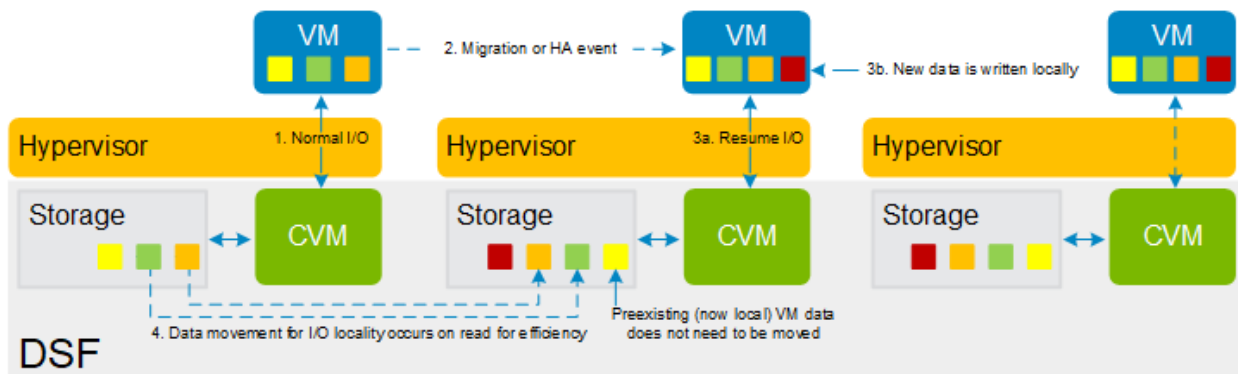
The Nutanix solution has no LUNs to manage, no RAID groups to configure, and no complicated storage multipathing to set up since there is no reliance on traditional SAN or NAS. All storage management is VM-centric, and the DSF optimizes I/O at the VM virtual disk level. There is one shared pool of storage that includes flash-based SSDs for high performance and low-latency HDDs for affordable capacity. The file system automatically tiers data across different types of storage devices using intelligent data placement algorithms. These algorithms make sure that the most frequently used data is available in memory or in flash for optimal performance. Organizations can also choose flash-only storage for the fastest possible storage performance. The following figure illustrates the data I/O path for a write in a hybrid model with a mix of SSD and HDD disks.



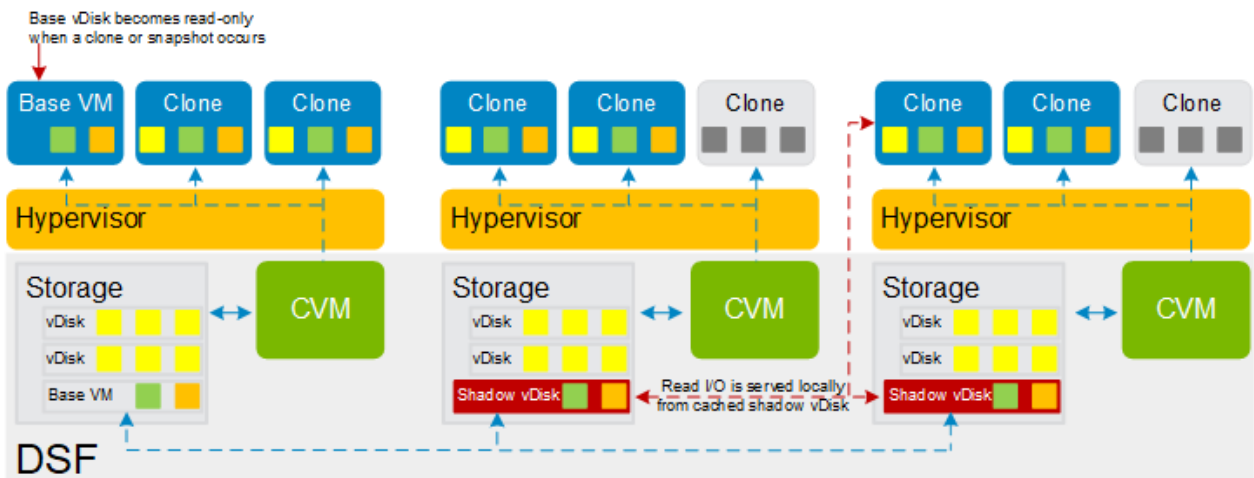
Local storage for each XC Series node in the architecture appears to the hypervisor as one large pool of shared storage. This allows the DSF to support all key virtualization features. Data localization maintains performance and quality of service (QoS) on each host, minimizing the effect noisy VMs have on their neighbors' performance. This functionality allows for large, mixed-workload clusters that are more efficient and more resilient to failure when compared to traditional architectures with standalone, shared, and dual-controller storage arrays.

When VMs move from one hypervisor to another, such as during live migration or a high availability (HA) event, the now local CVM serves a newly migrated VM's data. While all write I/O occurs locally, when the local CVM reads old data stored on the now remote CVM, the local CVM forwards the I/O request to the remote CVM. The DSF detects that I/O is occurring from a different node and migrates the data to the local

node in the background, ensuring that all read I/O is served locally as well. The next figure shows how data follows the VM as it moves between hypervisor nodes.



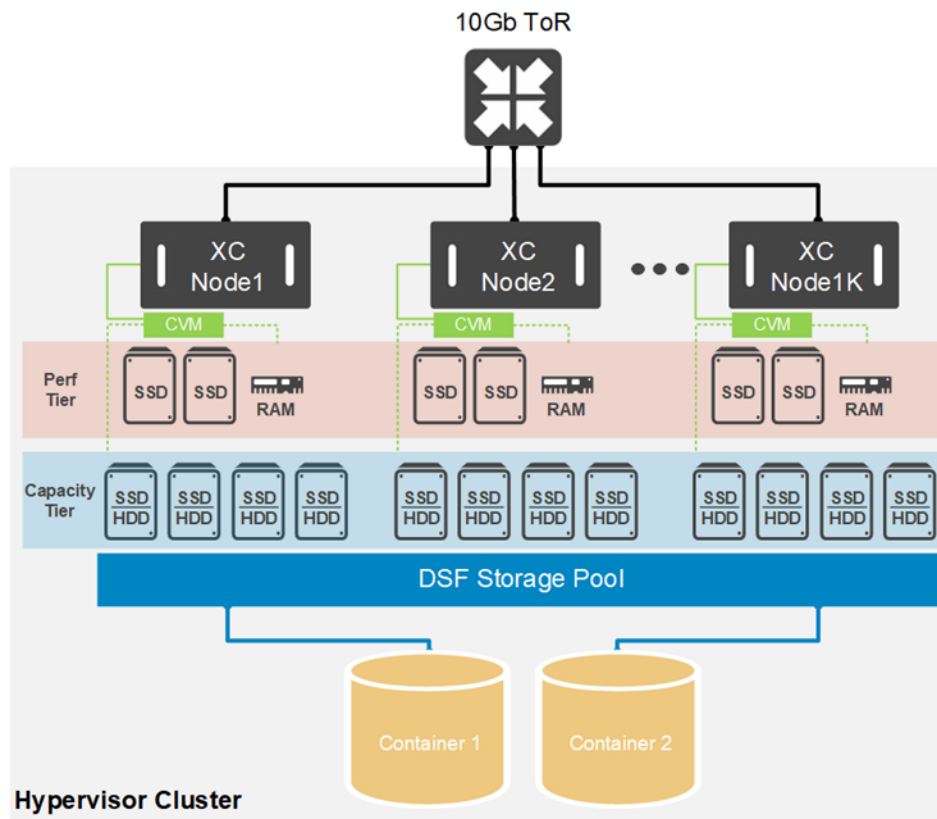
Nutanix Shadow Clones delivers distributed localized caching of virtual disks performance in multi-reader scenarios, such as desktop virtualization using VMware Horizon or Microsoft Remote Desktop Session Host (RDSH). With Shadow Clones, the CVM actively monitors virtual disk access trends. If there are requests originating from more than two remote CVMs, as well as the local CVM, and all of the requests are read I/O the virtual disk will be marked as immutable. When the disk is immutable, each CVM then caches it locally, so local storage can now satisfy read operations.



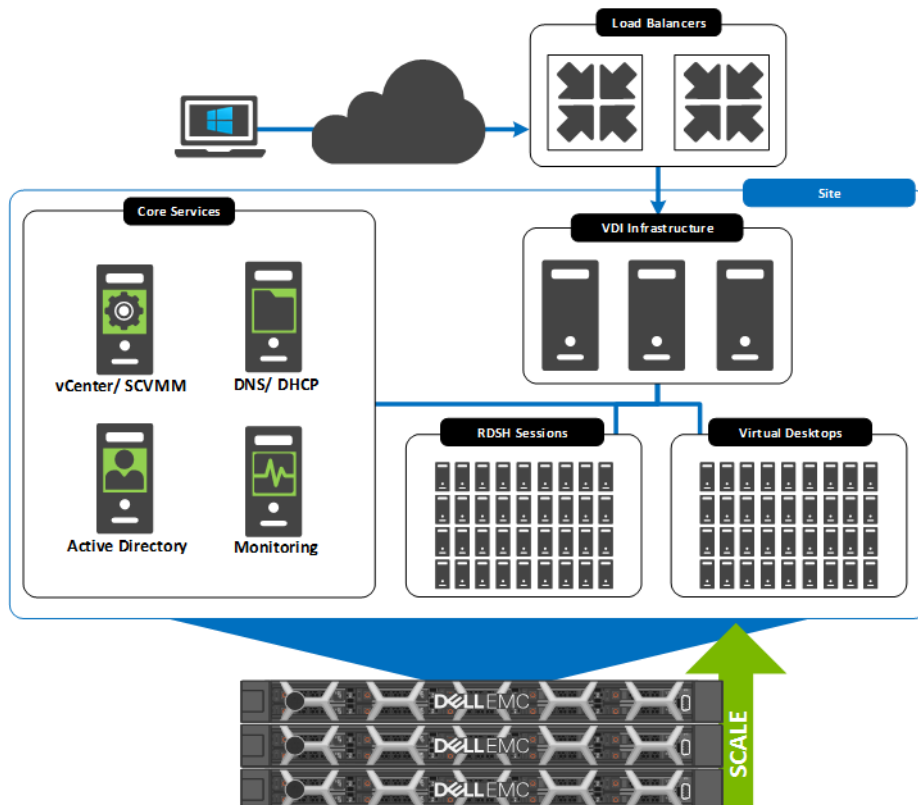
2.5 Nutanix Hyper-Converged Infrastructure

The Nutanix hyper-converged infrastructure provides an ideal combination of both high-performance compute with localized storage to meet any demand. True to this capability, this reference architecture has been validated as optimized for the VDI use case.

The next figure shows a high-level example of the relationship between an XC Series node, storage pool, container, pod and relative scale out:



This solution allows organizations to deliver virtualized or remote desktops and applications through a single platform and support end users with access to all of their desktops and applications in a single place.



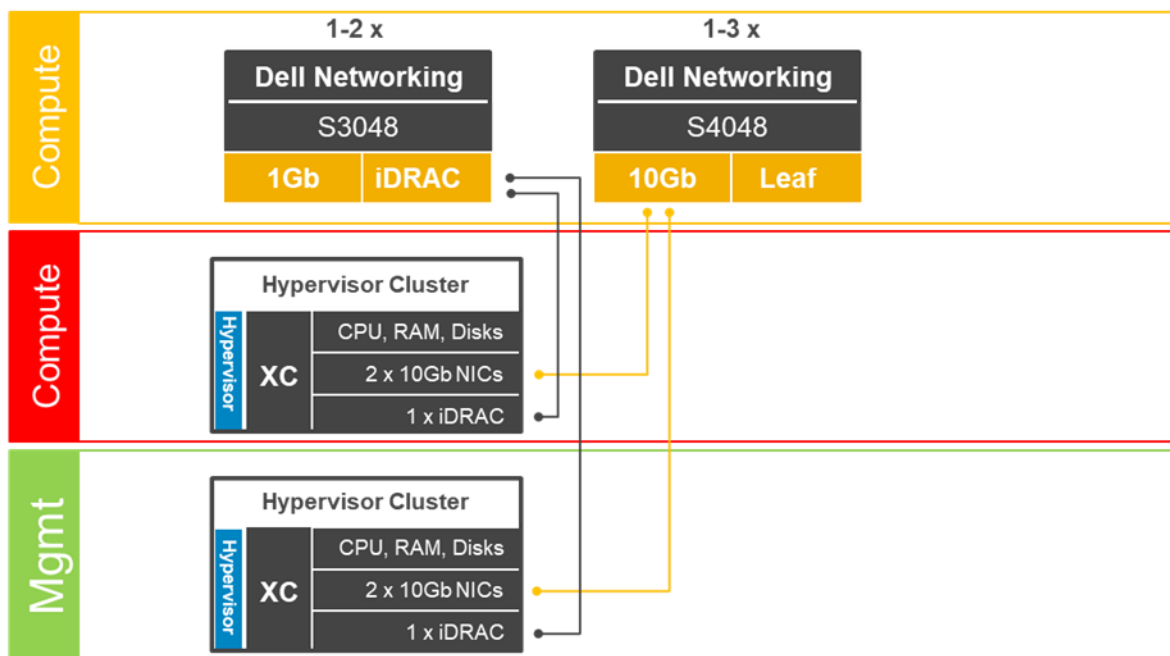
2.6 Nutanix all-flash

Nutanix supports an all-flash configuration where all local disks are SSDs and therefore, the storage pool is fully comprised of SSDs for both capacity and performance. The previously described features and functionality for management, data optimization and protection, and disaster recovery are still present. With all-flash, hot data is stored on SSDs local to each VM. If capacity needs exceed the local SSD storage, capacity on other nodes is automatically and transparently utilized. Compared to traditional all-flash shared storage arrays, XC Series all-flash clusters won't have the typical performance limitations due to network and storage controller bottlenecks. Benefits for VDI include faster provisioning times, low latency, ability to handle extremely high application I/O needs, and accommodating bursts of activity such as boot storms and anti-virus scans.

2.7 Dell EMC XC Series - VDI solution architecture

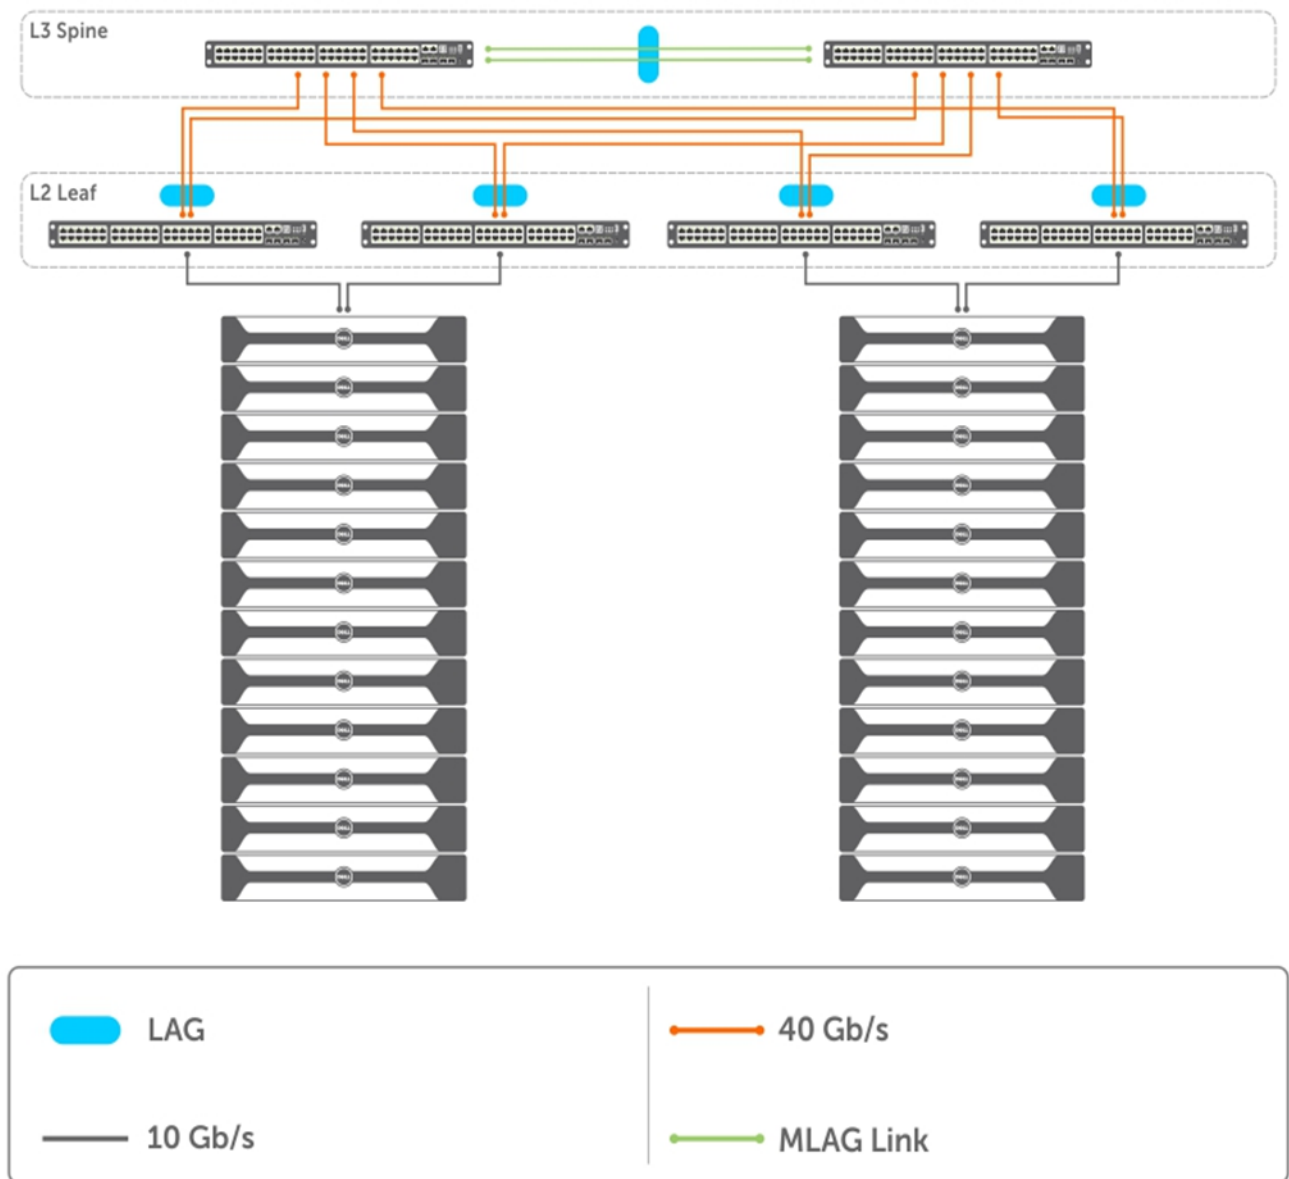
2.7.1 Networking

The networking layer consists of the 10Gb Dell Networking S4048 utilized to build a leaf/spine architecture with robust 1Gb switching in the S3048 for iDRAC connectivity.



Designed for true linear scaling, XC Series leverages a Leaf-Spine network architecture. A Leaf-Spine architecture consists of two network tiers: a 10Gb layer-2 (L2) Leaf segment and a layer-3 (L3) Spine segment based on 40GbE and non-blocking switches. This architecture maintains consistent performance without any throughput reduction due to a static maximum of three hops from any node in the network.

The following figure shows a design of a scale-out Leaf-Spine network architecture that provides 20Gb active throughput from each node to its Leaf and scalable 80Gb active throughput from each Leaf to Spine switch providing scale from 3 XC Series nodes to thousands without any impact to available bandwidth:



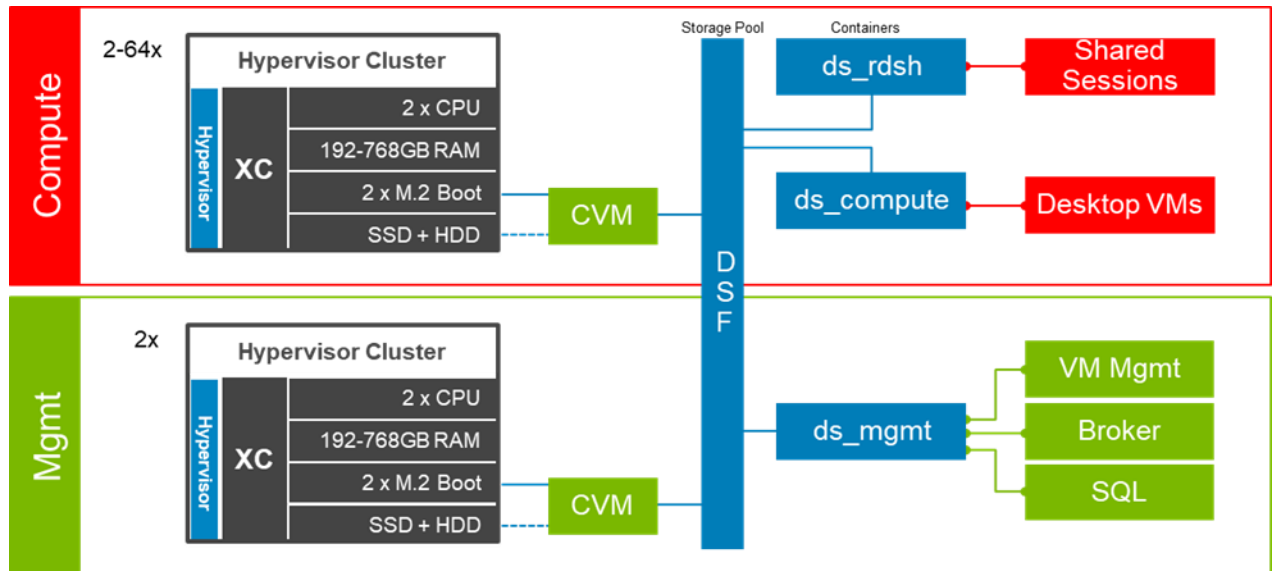
2.7.2 XC Series – Enterprise solution pods

The compute, management and storage layers are converged into each XC Series node in the cluster, hosting VMware vSphere. The recommended boundaries of an individual pod are based on the number of nodes supported within a given hypervisor cluster, 64 nodes for vSphere 6, although the Nutanix ADFS cluster can scale much larger, well beyond the boundaries of the hypervisor in use.

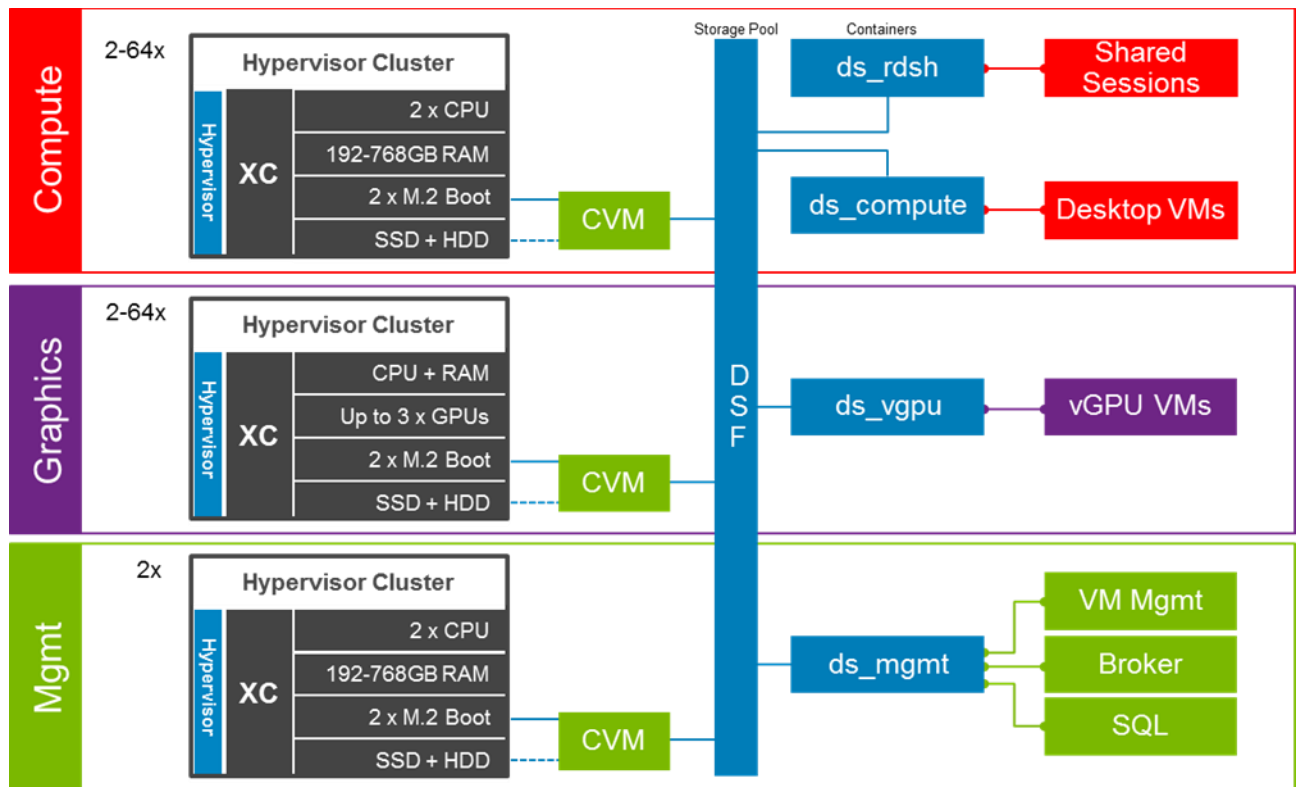
Dell EMC recommends that the VDI management infrastructure nodes be separated from the compute resources onto their own appliance cluster with a common DSF namespace shared between them based on NFS for vSphere. One node for VDI management is required, minimally, and expanded based on size of the

pod. The designations ds_rdsh, ds_compute, ds_vgpu and ds_mgmt as seen below are logical DSF containers used to group VMs of a particular type.

Using distinct containers allows features and attributes, such as compression and deduplication, to be applied to groups of VMs that share similar characteristics. Compute hosts can be used interchangeably for Horizon or RDSH as required. Distinct clusters should be built for management and compute hosts for HA, respectively, to plan predictable failover, scale and load across the pod. The DSF namespace can be shared across multiple hypervisor clusters adding disk capacity and performance for each distinct cluster.



High-performance graphics capabilities compliment the solution and can be added at any time to any new or existing XC Series vSphere-based deployment. Simply add the appropriate number of XC740xd appliances to your DSF cluster and provide a superior user experience with vSphere 6 and NVIDIA GRID vGPU technology. Any XC Series appliance can be utilized for the non-graphics compute or management portions of this solution and vSphere will provide HA accordingly based on the type of VM.



3 Hardware components


3.1 Network

The following sections contain the core network components for the solution. General uplink cabling guidance to consider in all cases is that TwinAx or CAT6 is very cost effective for short 10Gb runs and for longer runs use fiber with SFPs.

3.1.1 Dell Networking S3048 (1Gb ToR switch)

Accelerate applications in high-performance environments with a low-latency top-of-rack (ToR) switch that features 48 x 1GbE and 4 x 10Gb SFP+ ports, a dense 1U design and up to 260Gbps performance. The S3048-ON also supports Open Network Installation Environment (ONIE) for zero-touch installation of alternate network operating systems.


Model	Features	Options	Uses
<ul style="list-style-type: none">Dell Networking S3048-ON	<ul style="list-style-type: none">48 x 1000BaseT4 x 10Gb SFP+Non-blocking, line-rate performance260Gbps full-duplex bandwidth131 Mpps forwarding rate	<ul style="list-style-type: none">Redundant hot-swap PSUs & fans	<ul style="list-style-type: none">1Gb connectivity
		<ul style="list-style-type: none">VRF-lite, Routed VLT, VLT Proxy Gateway	
		<ul style="list-style-type: none">User port stacking (up to 6 switches)	
		<ul style="list-style-type: none">Open Networking Install Environment (ONIE)	



3.1.2 Dell Networking S4048 (10Gb ToR switch)

Optimize your network for virtualization with a high-density, ultra-low-latency ToR switch that features 48 x 10GbE SFP+ and 6 x 40GbE ports (or 72 x 10GbE ports in breakout mode) and up to 720Gbps performance. The S4048-ON also supports ONIE for zero-touch installation of alternate network operating systems.

Model	Features	Options	Uses
• Dell Networking S4048-ON	• 48 x 10Gb SFP+ • 6 x 40Gb QSFP+ • Non-blocking, line-rate performance • 1.44Tbps bandwidth • 720 Gbps forwarding rate • VXLAN gateway support	• Redundant hot-swap PSUs & fans	• 10Gb connectivity
		• 72 x 10Gb SFP+ ports with breakout cables	
		• User port stacking (up to 6 switches)	
		• Open Networking Install Environment (ONIE)	



For more information on the S3048, S4048 switches and Dell Networking, please visit: [LINK](#)

3.2 Dell EMC XC Series Hyper-Converged appliances

Dell EMC XC Series is based on the award-winning 14th generation of Dell EMC PowerEdge servers which offer a number of performance and feature enhancements. The table below outlines the hardware changes between generations.

	XC630	XC640	XC730xd	XC740xd	XC730 to XC740 Increase
CPU and chipset	Broadwell-EP	Skylake	Broadwell-EP	Skylake	
Front side bus	Intel QuickPath Interconnect @ 9.6 GT/s	Intel UltraPath Interconnect @ 11.2 GT/s	Intel QuickPath Interconnect @ 9.6 GT/s	Intel UltraPath Interconnect @ 10.4 GT/s	8%
Cores (max)	18 cores	28 cores	22 cores	28 cores	27%
TDP (max)	145 W	205 W	145 W	205 W	
Instruction set	AVX2	AVX2/ AVX-512	AVX2	AVX2/ AVX-512	
Max DP FLOPS / CLK	16 per core (w /AVX2)	32 per core (w / AVX-512)	16 per core (w /AVX2)	32 per core (w / AVX-512)	100%
Memory channels per socket	4 channels, DDR4	6 channels, DDR4	4 channels, DDR4	6 channels, DDR4	50%
Memory (max)	384 GB/ socket (768 GB total)	768 GB/ socket (1.5 TB total)	768 GB/ socket (1.5 TB total)	1.5 TB / socket (3 TB total)	100%
Memory speed (max)	2133 MT/s	2667 MT/s	2400 MT/s	2667 MT/s	11%
PCIe Lanes	40	48	40	48	20%

Consolidate compute and storage into a single chassis with XC Series Hyper-converged appliances, powered by Nutanix software. XC Series appliances install quickly, integrate easily into any data center, and can be deployed for multiple virtualized workloads including desktop virtualization, test and development, and private cloud projects. For general purpose virtual desktop and virtual application solutions, Dell EMC recommends the XC640 and XC740xd. For workloads requiring graphics the XC740xd with NVIDIA GRID vGPU can be integrated into any environment running any other XC Series appliance. For small Remote Office – Branch Office scenarios we offer the XC640. For more information on the Dell EMC XC Series, please visit: [Link](#)

The XC Series portfolio, optimized for VDI, has been designed and arranged in three top-level optimized configurations which apply to the available physical platforms showcased below.

- **A3** configuration is perfect for small scale, POC or low-density cost-conscious environments. Available on all standard hybrid platform configurations.
- **B5** configuration is geared toward larger scale general purpose workloads, balancing performance and cost-effectiveness. Available on all XC Series platforms.
- **C7** is the premium configuration offering an abundance of high performance and tiered capacity where user density is maximized. Available on all XC Series platforms.

A3	B5	C7
<ul style="list-style-type: none"> • CPU: 2 x 10-Core (4114) • RAM: 192GB (12 x 16GB @ 2400MHz) • Disk: 4TB+ (T2*) • GPU (Optional**): Up to 3 x FLDW 	<ul style="list-style-type: none"> • CPU: 2 x 14-Core (5120) • RAM: 384GB (12 x 32GB @ 2400MHz) • Disk: 6TB+ (T2*) • GPU (Optional**): Up to 3 x FLDW 	<ul style="list-style-type: none"> • CPU: 2 x 20-Core (6138) • RAM: 768GB (24 x 32GB @ 2667MHz) • Disk: 8TB+ (T2*) • GPU (Optional**): Up to 3 x FLDW


*Raw disk capacity target

**Available on XC740xd

3.2.1 Dell EMC XC640


The Dell EMC XC640 is a 10-disk 1U platform with a broad range of configuration options. Each appliance comes equipped with dual CPUs, 10 to 28cores, and up to 1.5TB of high-performance RAM. For the hybrid disk configuration, a minimum of six disks is required in each host, 2 x SSD for the performance tier (Tier1) and 4 x HDD for the capacity tier (Tier2) which can be expanded up to eight HDDs as required. For the all-flash disk configuration, the chassis must be populated with a minimum of 4 x SSDs. The M.2-based BOSS module boots the hypervisor and Nutanix Controller VM while the PERC HBA330 connects the CVM to the SSDs and HDDs. 64GB is consumed on each of the first two SSDs for the Nutanix “home”. All HDD/SSD disks are presented to the Nutanix CVM running locally on each host which contributes to the clustered DSF storage pool. Each platform can be outfitted with SFP+ or BaseT NICs.

3.2.1.1 XC640 hybrid disk storage



XC640 Hybrid	A3	B5	C7
CPU	2 x Intel Xeon Silver 4114 (10C, 2.2GHz)	2 x Intel Xeon Gold 5120 (14C, 2.2GHz)	2 x Intel Gold 6138 (20C, 2.0GHz)
Memory	12 x 16GB 2667MT/s RDIMMs Effective speed: 2400MT/s @ 192GB	12 x 32GB 2667MT/s RDIMMs Effective speed: 2400MT/s @ 384GB	24 x 32GB 2667MT/s RDIMMs Effective speed: 2667MT/s @ 768GB
Storage Ctrl	HBA330 LP		
Storage	CVM/ OS: 2 x 120GB M.2 RAID1 Tools/ Recovery: 32GB SD T1: 2 x 480GB SSD 2.5" T2: 4 x 1.8TB 2.5"/ 2TB HDD 3.5"	CVM/ OS: 2 x 120GB M.2 RAID1 Tools/ Recovery: 32GB SD T1: 2 x 960GB SSD 2.5" T2: 4 x 1.8TB 2.5"/ 2TB HDD 3.5"	CVM/ OS: 2 x 120GB M.2 RAID1 Tools/ Recovery: 32GB SD T1: 2 x 960GB SSD 2.5" T2: 6 x 1.8TB 2.5"/ 2TB HDD 3.5"
Network	2 x 10Gb, 2 x 1Gb SFP+/ BT		
iDRAC	iDRAC9 Enterprise		
Power	2 x 1100W PSUs		

3.2.1.2 XC640 all-flash disk storage



XC640		
All-Flash	B5-AF	C7-AF
CPU	2 x Intel Xeon Gold 5120 (14C, 2.2GHz)	2 x Intel Gold 6138 (20C, 2.0GHz)
Memory	12 x 32GB 2667MT/s RDIMMs Effective speed: 2400MT/s @ 384GB	24 x 32GB 2667MT/s RDIMMs Effective speed: 2667MT/s @ 768GB
Storage Ctrl	HBA330 LP	
Storage	CVM/ OS: 2 x 120GB M.2 RAID1 Tools/ Recovery: 32GB SD T1/ T2: 6 x 960GB SSD 2.5"	CVM/ OS: 2 x 120GB M.2 RAID1 Tools/ Recovery: 32GB SD T1/ T2: 10 x 960GB SSD 2.5"
Network	2 x 10Gb, 2 x 1Gb SFP+ or BaseT	
iDRAC	iDRAC9 Enterprise	
Power	2 x 1100W PSUs	


All-flash configuration requirements:

* 4 x SSD min: XC640-4

** 6 x SSD min XC640-10

3.2.2 Dell EMC XC740xd

The Dell EMC XC740xd is a 2U platform that can be configured with 24 x 2.5" disks or 12 x 3.5" disks to serve a broad range of capacity requirements. Each appliance comes equipped with dual CPUs, 10 to 28 cores, and up to 1.5TB of high-performance RAM. A minimum of six disks is required in each host, 2 x SSD for the performance tier (Tier1) and 4 x HDD for the capacity tier (Tier2) which can be expanded as required up to a possible 45TB+ per node raw. The M.2-based BOSS module boots the hypervisor and Nutanix Controller VM while the PERC HBA330 connects the CVM to the SSDs and HDDs. 64GB is consumed on each of the first two SSDs for the Nutanix "home". All HDD/SSD disks are presented to the Nutanix CVM running locally on each host which contributes to the clustered DSF pool. Each platform can be outfitted with SFP+ or BaseT NICs. The 24-disk XC740xd can support up to 3 NVIDIA M60 or 2 x M10 GPU cards. Please note that higher wattage power supplies will also be required when GPUs are in use, up to 2000W per PSU.



XC740xd	A3	B5	C7
Hybrid			
CPU	2 x Intel Xeon Silver 4114 (10C, 2.2GHz)	2 x Intel Xeon Gold 5120 (14C, 2.2GHz)	2 x Intel Gold 6138 (20C, 2.0GHz)
Memory	12 x 16GB 2667MT/s RDIMMs Effective speed: 2400MT/s @ 192GB	12 x 32GB 2667MT/s RDIMMs Effective speed: 2400MT/s @ 384GB	24 x 32GB 2667MT/s RDIMMs Effective speed: 2667MT/s @ 768GB
Storage Ctrl	HBA330		
Storage	CVM/ OS: 2 x 120GB M.2 RAID1 Tools/ Recovery: 32GB SD T1: 2 x 480GB SSD 2.5" T2: 4 x 1.8TB 2.5"/ 2TB HDD 3.5"	CVM/ OS: 2 x 120GB M.2 RAID1 Tools/ Recovery: 32GB SD T1: 2 x 960GB SSD 2.5" T2: 4 x 1.8TB 2.5"/ 2TB HDD 3.5"	CVM/ OS: 2 x 120GB M.2 RAID1 Tools/ Recovery: 32GB SD T1: 2 x 960GB SSD 2.5" T2: 6 x 1.8TB 2.5"/ 2TB HDD 3.5"

GPU	2 x Tesla M10 or 3 x Tesla M60
Network	2 x 10Gb, 2 x 1Gb SFP+ or BaseT
iDRAC	iDRAC9 Enterprise
Power	2 x 1100W PSUs (2 x 2000w PSUs for GPU)

All-flash configuration requirements:

** 6 x SSD min: XC740, XC740xd-12

*** 12 X SSD min: XC740xd-24

3.3 NVIDIA Tesla GPUs

Accelerate your most demanding enterprise data center workloads with NVIDIA® Tesla® GPU accelerators. Scientists can now crunch through petabytes of data up to 10x faster than with CPUs in applications ranging from energy exploration to deep learning. Plus, Tesla accelerators deliver the horsepower needed to run bigger simulations faster than ever before. For enterprises deploying VDI, Tesla accelerators are perfect for accelerating virtual desktops. GPUs can only be used with the Dell EMC XC730 platform.



3.3.1 NVIDIA Tesla M10

The NVIDIA® Tesla® M10 is a dual-slot 10.5 inch PCI Express Gen3 graphics card featuring four mid-range NVIDIA Maxwell™ GPUs and a total of 32GB GDDR5 memory per card (8GB per GPU). The Tesla® M10 doubles the number of H.264 encoders over the NVIDIA® Kepler™ GPUs and improves encoding quality, which enables richer colors, preserves more details after video encoding, and results in a high-quality user experience.

The NVIDIA® Tesla® M10 GPU accelerator works with NVIDIA GRID™ software to deliver the industry's highest user density for virtualized desktops and applications. It supports up to 64 desktops per GPU card using a 1GB framebuffer (up to 128 desktops per server) and gives businesses the power to deliver great graphics experiences to all of their employees at an affordable cost.

Specs	Tesla M10
Number of GPUs/ card	4 x NVIDIA Maxwell™ GPUs
Total CUDA cores	2560 (640 per GPU)
GPU Clock	Idle: 405MHz / Base: 1033MHz
Total memory size	32GB GDDR5 (8GB per GPU)
Max power	225W
Form Factors	Dual slot (4.4" x 10.5")
Aux power	8-pin connector
PCIe	x16 (Gen3)
Cooling solution	Passive

3.3.2 NVIDIA Tesla M60

The NVIDIA® Tesla® M60 is a dual-slot 10.5 inch PCI Express Gen3 graphics card featuring two high-end NVIDIA Maxwell™ GPUs and a total of 16GB GDDR5 memory per card. This card utilizes NVIDIA GPU Boost™ technology which dynamically adjusts the GPU clock to achieve maximum performance. Additionally, the Tesla® M60 doubles the number of H.264 encoders over the NVIDIA® Kepler™ GPUs.



The NVIDIA® Tesla® M60 GPU accelerator works with NVIDIA GRID™ software to provide the industry's highest user performance for virtualized workstations, desktops, and applications. It allows enterprises to virtualize almost any application (including professional graphics applications) and deliver them to any device, anywhere. M60 can support 3 cards in the XC740xd providing 48 x Windows10 users assigned a 1GB framebuffer each.

Specs	Tesla M60
Number of GPUs/ card	2 x NVIDIA Maxwell™ GPUs
Total CUDA cores	4096 (2048 per GPU)
Base Clock	899 MHz (Max: 1178 MHz)
Total memory size	16GB GDDR5 (8GB per GPU)
Max power	300W
Form Factors	Dual slot (4.4" x 10.5")
Aux power	8-pin connector
PCIe	x16 (Gen3)
Cooling solution	Passive/ Active

3.4 Dell Wyse Endpoints

The following Dell Wyse clients will deliver a superior user experience for VMware Horizon and are the recommended choices for this solution.

3.4.1 Wyse 3040 Thin Client (ThinOS, ThinLinux)

The Wyse 3040 is the industry's first entry-level Intel x86 quad-core thin client, powered by a quad-core Intel Atom 1.44GHz processor, delivering robust connectivity options with a choice of Wyse ThinOS or ThinLinux operating systems. The Wyse 3040 is Dell's lightest, smallest and most power-efficient thin client – it consumes 3.3 Watts in idle state – and offers superb performance and manageability for task and basic productivity users. Despite its small size, the 3040 includes all typical interfaces such as four USB ports including USB 3.1, two DisplayPort interfaces and wired and wireless options. It is highly manageable as it can be monitored, maintained, and serviced remotely via Wyse Device Manager (WDM) or Wyse Management Suite. For more information, please visit: [Link](#)



3.4.2 Wyse 5040 AIO Thin Client (ThinOS)



The Dell Wyse 5040 AIO all-in-one (AIO) thin client runs ThinOS (with or without PCoIP), has a 21.5" Full HD display and offers versatile connectivity options for use in a wide range of industries. With four USB 2.0 ports, Gigabit Ethernet and integrated dual band Wi-Fi options, users can link to their peripherals and quickly connect to the network while working with processing-intensive, graphics-rich applications. Built-in speakers, a camera and a microphone make video conferencing and desktop communication simple and easy. It even supports a second attached display for those who need a dual monitor configuration. A simple one-cord design and out-of-box automatic setup makes deployment effortless while

remote management from a simple file server, Wyse Device Manager (WDM), or Wyse Management Suite can help lower your total cost of ownership as you grow from just a few thin clients to tens of thousands. For more information, please visit: [Link](#)

3.4.3 Wyse 5060 Thin Client (ThinOS, ThinLinux, WES7P, WIE10)

The Wyse 5060 offers high performance and reliability, featuring all the security and management benefits of Dell thin clients. It comes with flexible OS options: ThinOS (with or without PCoIP), ThinLinux, Windows Embedded Standard 7P (WES7P) or Windows 10 IoT Enterprise (WIE10). Designed for knowledge workers demanding powerful virtual desktop performance, and support for unified communications solutions like Skype for Business, the Wyse 5060 thin client delivers the flexibility, efficiency and security organizations require for their cloud environments. It is powered by a quad-core AMD 2.4GHz processor, supports dual 4K (3840x2160) monitors and provides multiple connectivity options with six USB ports, two of which are USB 3.0 for high-speed peripherals, as well as two DisplayPort connectors, wired networking or wireless 802.11 a/b/g/n/ac. The Wyse 5060 can be monitored, maintained, and serviced remotely via Wyse Device Manager (WDM), cloud-based Wyse Management Suite or Microsoft SCCM (5060 with Windows versions). For more information, please visit: [Link](#)



3.4.4 Wyse 7020 Thin Client (WES 7/7P/8, WIE10, ThinLinux)

The versatile Dell Wyse 7020 thin client is a powerful endpoint platform for virtual desktop environments. It is available with Windows Embedded Standard 7/7P/8 (WES), Windows 10 IoT Enterprise (WIE10), Wyse ThinLinux operating systems and it supports a broad range of fast, flexible connectivity options so that users can connect their favorite peripherals while working with processing-intensive, graphics-rich applications. This 64-bit thin client delivers a great user experience and support for local applications while ensuring security. Designed to provide a superior user experience, ThinLinux features broad broker support including Citrix Receiver, VMware Horizon and Amazon Workspace, and support for unified communication platforms including Skype for Business, Lync 2013 and Lync 2010. For additional security, ThinLinux also supports single sign-on and VPN. With a powerful quad core AMD G Series APU in a compact chassis with dual-HD monitor support, the Wyse 7020 thin client delivers stunning performance and display capabilities across 2D, 3D and HD video applications. Its silent diskless and fan less design helps reduce power usage to just a fraction (it only consumes about 15 watts) of that used in traditional desktops. Wyse Device Manager (WDM) helps lower the total cost of ownership for large deployments and offers remote enterprise-wide management that scales from just a few to tens of thousands of cloud clients. For more information, please visit [Link](#)



3.4.5 Wyse 7040 Thin Client (WES7P, WIE10)

The Wyse 7040 is a high-powered, ultra-secure thin client running Windows Embedded Standard 7P (WES7P) or Windows 10 IoT Enterprise (WIE10) operating systems. Equipped with an Intel i5/i7 processors, it delivers extremely high graphical display performance (up to three displays via display-port daisy-chaining, with 4K resolution available on a single monitor) for seamless access to the most demanding applications. The Wyse 7040 is compatible with both data center hosted and client-side virtual desktop environments and is compliant with all relevant U.S. Federal security certifications including OPAL compliant hard-drive options, VPAT/Section 508, NIST BIOS, Energy-Star and EPEAT. Wyse enhanced WES7P OS provides additional security features such as BitLocker. The Wyse 7040 offers a high level of connectivity including dual NIC, 6 x USB3.0 ports and an optional second network port, with either copper or fiber SFP interface. Wyse 7040 devices are highly manageable through Intel vPRO, Wyse Device Manager (WDM), Microsoft System Center Configuration Manager (SCCM) and Dell Command Configure (DCC). For more information, please visit: [Link](#)



Enhanced Security

Note that all the above thin clients running Windows Embedded Standard 7 or Windows 10 IoT can be protected against viruses, ransomware and zero-day threats by installing [Dell Threat Defense](#), a revolutionary anti-malware software solution using artificial intelligence and mathematical modeling and is not signature-based. Threat Defense prevents 99% of executable malware, far above the average 50% of threats identified by the top anti-virus solutions. It doesn't need a constant internet connection nor frequent updates (only about twice a year), it only uses 1-3% CPU and has only a ~40MB memory footprint, making it an ideal choice to protect thin clients without impacting the end user productivity.

If you also want to protect virtual desktops against such malware and threats with a similar success, Dell recommends using [Dell Endpoint Security Suite Enterprise](#), a full suite featuring advanced threat prevention and data-centric encryption using an on-premise management console. This suite can also be used to protect physical PCs, MAC OS X systems and Windows Server.

4 Software components

4.1 VMware

4.1.1 VMware Horizon 7

The solution is based on VMware Horizon which provides a complete end-to-end solution delivering Microsoft Windows virtual desktops to users on a wide variety of endpoint devices. Virtual desktops are dynamically assembled on demand, providing users with pristine, yet personalized, desktops each time they log on.

VMware Horizon provides a complete virtual desktop delivery system by integrating several distributed components with advanced configuration tools that simplify the creation and real-time management of the virtual desktop infrastructure. For the complete set of details, please see the Horizon View resources page at <http://www.vmware.com/products/horizon-view/resources.html>.


The core Horizon components include:

- **Connection Server (VCS)** – Installed on servers in the data center and brokers client connections, The VCS authenticates users, entitles users by mapping them to desktops and/or pools, establishes secure connections from clients to desktops, support single sign-on, sets and applies policies, acts as a DMZ security server for outside corporate firewall connections and more.
- **Client** – Installed on endpoints. Is software for creating connections to View desktops that can be run from tablets, Windows, Linux, or Mac PCs or laptops, thin clients and other devices.
- **Portal** – A web portal to access links for downloading full View clients. With HTML Access Feature enabled enablement for running a View desktop inside a supported browser is enabled.
- **Agent** – Installed on all VMs, physical machines and Terminal Service servers that are used as a source for View desktops. On VMs the agent is used to communicate with the View client to provide services such as USB redirection, printer support and more.
- **Horizon Administrator** – A web portal that provides admin functions such as deploy and management of View desktops and pools, set and control user authentication and more.
- **Composer** – This software service can be installed standalone or on the vCenter server and provides enablement to deploy and create linked clone desktop pools (not required if using Instant Clones).
- **vCenter Server** – This is a server that provides centralized management and configuration to entire virtual desktop and host infrastructure. It facilitates configuration, provision, management services. It is installed on a Windows Server host (can be a VM).
- **Transfer Server** – Manages data transfers between the data center and the View desktops that are checked out on the end users' desktops in offline mode. This Server is required to support desktops that run the View client with Local Mode options. Replications and syncing are the functions it will perform with offline images.

4.1.2 VMware vSphere 6

The vSphere hypervisor also known as ESXi is a bare-metal hypervisor that installs directly on top of your physical server and partitions it into multiple virtual machines. Each virtual machine shares the same physical resources as the other virtual machines and they can all run at the same time. Unlike other hypervisors, all management functionality of vSphere is done through remote management tools. There is no underlying operating system, reducing the install footprint to less than 150MB.

VMware vSphere 6 includes three major layers: Virtualization, Management and Interface. The Virtualization layer includes infrastructure and application services. The Management layer is central for configuring, provisioning and managing virtualized environments. The Interface layer includes the vSphere web client.



Throughout this Dell EMC solution, all VMware and Microsoft best practices and prerequisites for core services are adhered to (NTP, DNS, Active Directory, etc.). The vCenter 6 VM used in the solution is a single Windows Server 2012 R2 VM or vCenter 6 virtual appliance, residing on a host in the management layer. SQL server is a core component of the Windows version of vCenter and is hosted on another VM also residing in the management layer. It is recommended that all additional Horizon components be installed in a distributed architecture, one role per server VM.

4.2 Microsoft RDSH

The RDSH servers can exist as physical or virtualized instances of Windows Server 2012 R2. A minimum of one, up to a maximum of ten virtual servers are installed per physical compute host. Since RDSH instances are easily added to an existing Horizon stack, the only additional components required are one or more Windows Server OS instances added to the Horizon site

The total number of required virtual RDSH servers is dependent on application type, quantity and user load. Deploying RDSH virtually and in a multi-server farm configuration increases overall farm performance, application load balancing as well as farm redundancy and resiliency.


4.2.1 NUMA architecture considerations

Best practices and testing has showed that aligning RDSH design to the physical Non-Uniform Memory Access (NUMA) architecture of the server CPUs results in increased and optimal performance. NUMA alignment ensures that a CPU can access its own directly-connected RAM banks faster than those banks of the adjacent processor which are accessed via the Quick Path Interconnect (QPI). The same is true of VMs with large vCPU assignments, best performance will be achieved if your VMs receive their vCPU allotment from a single physical NUMA node. Ensuring that your virtual RDSH servers do not span physical NUMA nodes will ensure the greatest possible performance benefit.

The general guidance for RDSH NUMA-alignment on the Dell EMC XC Series is as follows:

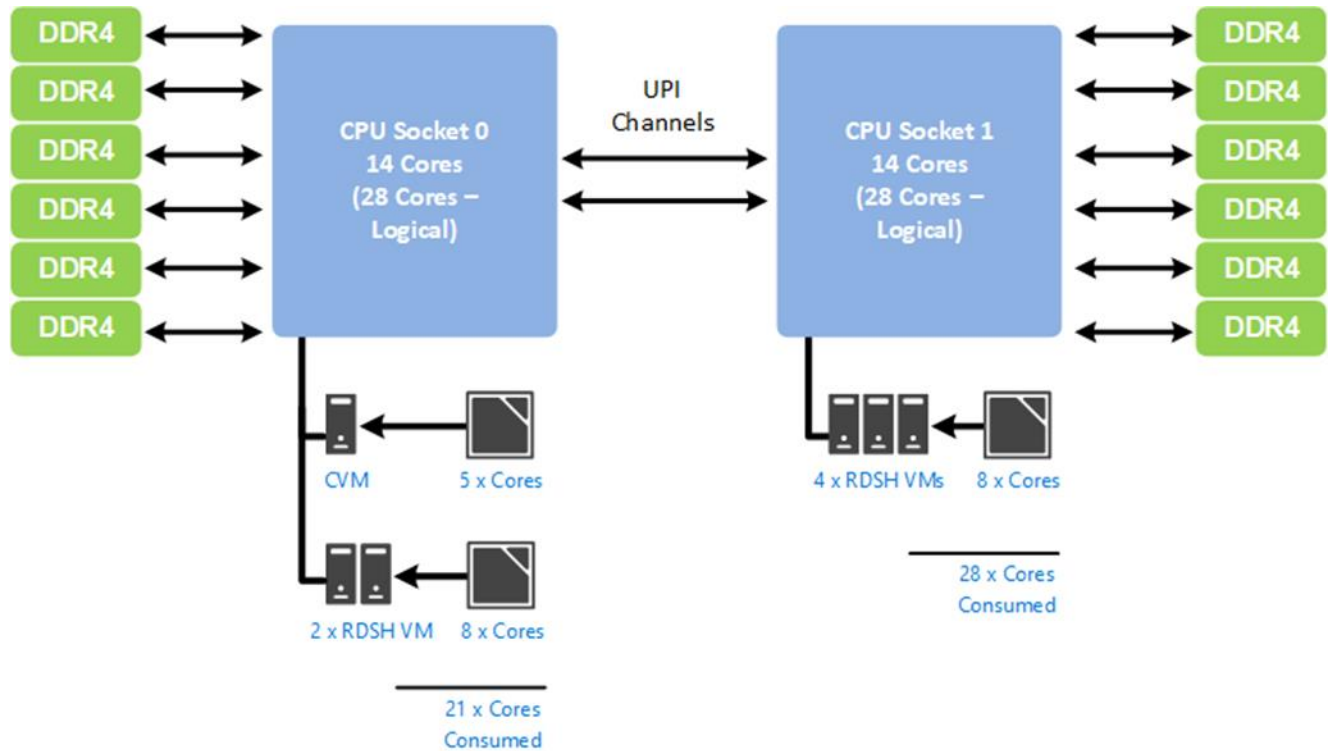
4.2.1.1 A3 NUMA alignment

10 physical cores per CPU in the A3 configuration, 20 logical with Hyper-threading active, gives a total of 40 consumable cores per appliance. The Nutanix CVM will receive its vCPU allotment from the first physical CPU and by configuring the RDSH VMs as shown below will ensure that no NUMA spanning occurs which could lower performance. Per the example below, we have three total RDSH VMs configured with 8 vCPUs each, along with the Nutanix CVM configured with 8 or 10 vCPUs (automatically detects the number of cores on the host and sizes accordingly). Note that the CVM actually reserves 10,000MHz which equates to roughly 5 cores reserved. This leaves both sockets with some additional headroom for workload assignment and burst scheduling. Please note that the A3 and B5 configs have 2 x UPI channels, versus 3 channels on the C7.



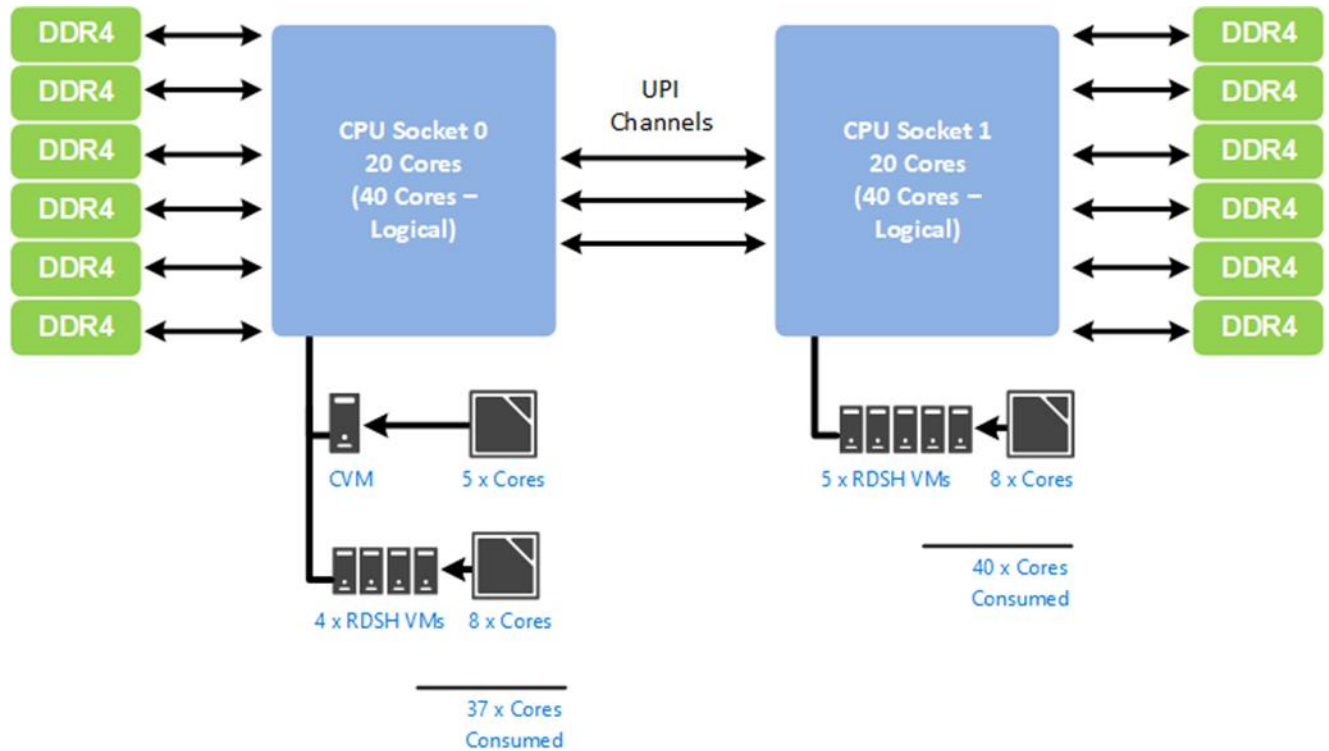
4.2.1.2 B5 NUMA alignment

14 physical cores per CPU in the B5 configuration, 28 logical with Hyper-threading active, gives a total of 56 consumable cores per node. The Nutanix CVM will receive its vCPU allotment from the first physical CPU and by configuring the RDSH VMs as shown below will ensure that no NUMA spanning occurs, which could lower performance. Per the example below, we have six RDSH VMs configured with 8 vCPUs each, along with the Nutanix CVM with 12 vCPUs. Note that the CVM actually reserves 10,000MHz which equates to roughly 5 cores reserved. This leaves some additional headroom on socket 0 for workload assignment and burst scheduling. Please note that the A3 and B5 configs have 2 x UPI channels, versus 3 channels on the C7.



4.2.1.3 C7 NUMA alignment

20 physical cores per CPU in the C7 configuration, 40 logical with Hyper-threading active, gives us a total of 80 consumable cores per node. The Nutanix CVM will receive its vCPU allotment from the first physical CPU and by configuring the RDSH VMs as shown below will ensure that no NUMA spanning occurs which could lower performance. Per the example below, we have eight RDSH VMs configured with 8 vCPUs each, along with the Nutanix CVM with 12 vCPUs. Note that the CVM actually reserves 10,000MHz which equates to 5 cores reserved. This leaves some additional headroom on socket 0 for workload assignment and burst scheduling.



4.3 NVIDIA GRID vGPU

NVIDIA GRID™ vGPU™ brings the full benefit of NVIDIA hardware-accelerated graphics to virtualized solutions. This technology provides exceptional graphics performance for virtual desktops equivalent to local PCs when sharing a GPU among multiple users.

GRID vGPU is the industry's most advanced technology for sharing true GPU hardware acceleration between multiple virtual desktops—without compromising the graphics experience. Application features and compatibility are exactly the same as they would be at the user's desk.

With GRID vGPU technology, the graphics commands of each virtual machine are passed directly to the GPU, without translation by the hypervisor. This allows the GPU hardware to be time-sliced to deliver outstanding shared virtualized graphics performance.




Image provided courtesy of NVIDIA Corporation, Copyright NVIDIA Corporation

4.3.1 vGPU profiles

Virtual Graphics Processing Unit, or GRID vGPU™, is technology developed by NVIDIA® that enables hardware sharing of graphics processing for virtual desktops. This solution provides a hybrid shared mode allowing the GPU to be virtualized while the virtual machines run the native NVIDIA video drivers for better performance. Thanks to OpenGL support, VMs have access to more graphics applications. When utilizing vGPU, the graphics commands from virtual machines are passed directly to the GPU without any hypervisor translation. Every virtual desktop has dedicated graphics memory so they always have the resources they need to launch and run their applications at full performance. All this is done without sacrificing server performance and so is truly cutting edge.

The combination of Dell EMC servers, NVIDIA GRID vGPU™ technology and NVIDIA Tesla™ cards enable high-end graphics users to experience high fidelity graphics quality and performance, for their favorite applications at a reasonable cost.

For more information about NVIDIA GRID vGPU, please visit: [LINK](#)

The number of users per appliance is determined by the number of GPU cards in the system (max 2 x M10 or 3 x M60), vGPU profiles used for each GPU in a card, and GRID license type. The same profile must be used on a single GPU but profiles can differ across GPUs within a single card.

NVIDIA® Tesla® M10 GRID vGPU Profiles:

Card	vGPU Profile	Graphics Memory (Frame Buffer)	Virtual Display Heads	Maximum Resolution	Maximum Graphics-Enabled VMs		
					Per GPU	Per Card	Per Server (2 cards)
Tesla M10	M10-8Q	8GB	4	4096x2160	1	4	8
	M10-4Q	4GB	4	4096x2160	2	8	16
	M10-2Q	2GB	4	4096x2160	4	16	32
	M10-1Q	1GB	2	4096x2160	8	32	64
	M10-0Q	512MB	2	2560x1600	16	64	128
	M10-1B	1GB	4	2560x1600	8	32	64
	M10-0B	512MB	2	2560x1600	16	64	128
	M10-8A	8GB	1	1280x1024	1	4	8
	M10-4A	4GB			2	8	16
	M10-2A	2GB			4	16	32
	M10-1A	1GB			8	32	64

Card	vGPU Profile	Guest VM OS Supported*		License Required
		Win	64bit Linux	
Tesla M10	M10-8Q	•	•	NVIDIA® Quadro® Virtual Data Center Workstation
	M10-4Q	•	•	
	M10-2Q	•	•	
	M10-1Q	•	•	
	M10-0Q	•	•	
	M10-1B	•		GRID Virtual PC
	M10-0B	•		
	M10-8A	•		GRID Virtual Application
	M10-4A	•		
	M10-2A	•		
	M10-1A	•		

Supported Guest VM Operating Systems*	
Windows	Linux
Windows 7 (32/64-bit)	RHEL 6.6 & 7
Windows 8.x (32/64-bit)	CentOS 6.6 & 7
Windows 10 (32/64-bit)	Ubuntu 12.04 & 14.04 LTS
Windows Server 2008 R2	
Windows Server 2012 R2	
Windows Server 2016	

*NOTE: Supported guest operating systems listed as of the time of this writing. Please refer to NVIDIA's documentation for latest supported operating systems.

NVIDIA® Tesla® M60 GRID vGPU Profiles:

Card	vGPU Profile	Graphics Memory (Frame Buffer)	Virtual Display Heads	Maximum Resolution	Maximum Graphics-Enabled VMs		
					Per GPU	Per Card	Per Server (3 cards)
Tesla M60	M60-8Q	8GB	4	4096x2160	1	2	6
	M60-4Q	4GB	4	4096x2160	2	4	12
	M60-2Q	2GB	4	4096x2160	4	8	24
	M60-1Q	1GB	2	4096x2160	8	16	48
	M60-0Q	512MB	2	2560x1600	16	32	96
	M60-1B	1GB	4	2560x1600	8	16	48
	M60-0B	512MB	2	2560x1600	16	32	96
	M60-8A	8GB	1	1280x1024	1	2	6
	M60-4A	4GB			2	4	12
	M60-2A	2GB			4	8	24
	M60-1A	1GB			8	16	48

Card	vGPU Profile	Guest VM OS Supported*		License Required
		Win	64bit Linux	
Tesla M60	M60-8Q	•	•	NVIDIA® Quadro® Virtual Data Center Workstation
	M60-4Q	•	•	
	M60-2Q	•	•	
	M60-1Q	•	•	
	M60-0Q	•	•	
	M60-1B	•		GRID Virtual PC
	M60-0B	•		
	M60-8A	•		GRID Virtual Application
	M60-4A	•		
	M60-2A	•		
	M60-1A	•		

Supported Guest VM Operating Systems*	
Windows	Linux
Windows 7 (32/64-bit)	RHEL 6.6 & 7
Windows 8.x (32/64-bit)	CentOS 6.6 & 7
Windows 10 (32/64-bit)	Ubuntu 12.04 & 14.04 LTS
Windows Server 2008 R2	
Windows Server 2012 R2	
Windows Server 2016	

*NOTE: Supported guest operating systems listed as of the time of this writing. Please refer to NVIDIA's documentation for latest supported operating systems.

4.3.1.1 GRID vGPU licensing and architecture

NVIDIA® GRID vGPU™ is offered as a licensable feature on Tesla® GPUs. vGPU can be licensed and entitled using one of the three following software editions.




NVIDIA® GRID® Virtual Applications	NVIDIA® GRID® Virtual PC	NVIDIA® Quadro® Virtual Data Center Workstation
For organizations deploying or other RDSH solutions. Designed to deliver Windows applications at full performance.	For users who want a virtual desktop, but also need a great user experience leveraging PC applications, browsers, and high-definition video.	For users who need to use professional graphics applications with full performance on any device, anywhere.
Up to 2 displays @ 1280x1024 resolution supporting virtualized Windows applications	Up to 4 displays @ 2560x1600 resolution supporting Windows desktops, and NVIDIA Quadro features	Up to 4 displays @ 4096x2160* resolution supporting Windows or Linux desktops, NVIDIA Quadro, CUDA**, OpenCL** & GPU pass-through

*OQ profiles only support up to 2560x1600 resolution

**CUDA and OpenCL only supported with M10-8Q, M10-8A, M60-8Q, or M60-8A profiles

The GRID vGPU Manager, running on the hypervisor installed via the VIB, controls the vGPUs that can be assigned to guest VMs. A properly configured VM obtains a license from the GRID license server during the boot operation for a specified license level. The NVIDIA graphics driver running on the guest VM provides direct access to the assigned GPU. When the VM is shut down, it releases the license back to the server. If a vGPU enabled VM is unable to obtain a license, it will run at full capability without the license but users will be warned each time it tries and fails to obtain a license.



5 Solution architecture for Horizon

5.1 Management role configuration

The Management role recommendations for the base solution are summarized below. Use data disks for role-specific application files such as data, logs and IIS web files in the Management volume.

5.1.1 VMware Horizon management role requirements

Role	vCPU	vRAM (GB)	vNIC	OS vDisk	
				Size (GB)	Location
Nutanix CVM	8	16	2	-	(BOSS)
Connection Server	4	8	1	40	DSF: ds_mgmt
Primary SQL	4	8	1	40 + 200	DSF: ds_mgmt
vCenter Appliance	2	8	1	125	DSF: ds_mgmt
Total	18	40	5	405	-

5.1.2 RDSH on vSphere

When using NVIDIA Tesla cards, graphics enabled VMs must obtain a license from a GRID License server on your network to be entitled for vGPU. To configure, a virtual machine with the following specifications must be added to a management host in addition to the management role VMs.

Role	HW Config	VMs per host	vCPUs per VM	RAM (GB)	vNIC	OS vDisk	
						Size (GB)	Location
RDSH VM	A3	3	8	32	1	80	DSF: ds_rdsh
RDSH VM	B5	6	8	32	1	80	DSF: ds_rdsh
RDSH VM	C7	8	8	32	1	80	DSF: ds_rdsh

5.1.3 NVIDIA GRID license server requirements

When using NVIDIA Tesla cards, graphics enabled VMs must obtain a license from a GRID License server on your network to be entitled for vGPU. To configure, a virtual machine with the following specifications must be added to a management host in addition to the management role VMs.

Role	vCPU	vRAM (GB)	NIC	OS vDisk	
				Size (GB)	Location
NVIDIA GRID License Srv	2	4	1	40 + 5	DSF: ds_mgmt

GRID License server software can be installed on a system running the following operating systems:

- Windows 7 (x32/x64)
- Windows 8.x (x32/x64)
- Windows 10 x64
- Windows Server 2008 R2
- Windows Server 2012 R2
- Red Hat Enterprise 7.1 x64
- CentOS 7.1 x64

Additional license server requirements:

- A fixed (unchanging) IP address. The IP address may be assigned dynamically via DHCP or statically configured, but must be constant.
- At least one unchanging Ethernet MAC address, to be used as a unique identifier when registering the server and generating licenses in NVIDIA's licensing portal.
- The date/time must be set accurately (all hosts on the same network should be time synchronized).

5.1.4 SQL databases

The VMware databases are hosted by a single dedicated SQL 2012 R2 Server VM in the Management layer. Use caution during database setup to ensure that SQL data, logs, and TempDB are properly separated onto their respective volumes. Create all Databases that are required for:

- VMware Horizon
- vCenter (if using Windows version)

Initial placement of all databases into a single SQL instance is fine unless performance becomes an issue, in which case database need to be separated into separate named instances. Enable auto-growth for each DB.

Best practices defined by Microsoft and VMware are to be adhered to, to ensure optimal database performance.

Align all disks to be used by SQL Server with a 1024K offset and then formatted with a 64K file allocation unit size (data, logs, and TempDB).

5.1.5 DNS




DNS plays a crucial role in the environment not only as the basis for Active Directory but is used to control access to the various VMware and Microsoft software components. All hosts, VMs, and consumable software components need to have a presence in DNS, preferably via a dynamic and AD-integrated namespace. Microsoft best practices and organizational requirements are to be adhered to.

Pay consideration for eventual scaling, access to components that may live on one or more servers (SQL databases, VMware Horizon services) during the initial deployment. Use CNAMEs and the round robin DNS mechanism to provide a front-end “mask” to the back-end server actually hosting the service or data source.

5.1.5.1 DNS for SQL

To access the SQL data sources, either directly or via ODBC, a connection to the server name\instance name must be used. To simplify this process, as well as protect for future scaling (HA), instead of connecting to server names directly, alias these connections in the form of DNS CNAMEs. So instead of connecting to SQLServer1\<instance name> for every device that needs access to SQL, the preferred approach is to connect to <CNAME>\<instance name>.

For example, the CNAME “VDisQL” is created to point to SQLServer1. If a failure scenario was to occur and SQLServer2 would need to start serving data, we would simply change the CNAME in DNS to point to SQLServer2. No infrastructure SQL client connections would need to be touched.

 SQLServer1	Host (A)	10.1.1.28
 SQLServer2	Host (A)	10.1.1.29
 SQLVDI	Alias (CNAME)	SQLServer1.fcs.local

5.2 Storage architecture overview

All Dell EMC XC Series appliances come with two tiers of storage by default, SSD for performance and HDD for capacity. Additionally, all-flash configurations are available utilizing only SSD disks. A single common Software Defined Storage namespace is created across the Nutanix cluster and presented as either NFS or SMB to the hypervisor of each host. This constitutes a storage pool and one should be sufficient per cluster. Within this common namespace, logical containers are created to group VM files as well as control the specific storage-related features that are desired to be enabled such as deduplication and compression.

5.2.1 Nutanix containers

The following table outlines the recommended containers, their purpose and settings given the use case. Best practices suggest using as few features as possible, only enable what is absolutely required. For example, if you are not experiencing disk capacity pressure then there is no need to enable Capacity Tier Deduplication. Enabling unnecessary services increases the resource demands of the Controller VMs. Capacity tier deduplication requires that CVMs be configured with 32GB RAM. Erasure Coding (EC-X) is recommended to increase usable capacity of the cluster.

Container	Purpose	Replication Factor	EC-X	Perf Tier Deduplication	Capacity Tier Deduplication	Compression
Ds_compute	Desktop VMs	2	Enabled	Enabled	Disabled	Disabled
Ds_mgmt	Mgmt Infra VMs	2	Enabled	Enabled	Disabled	Disabled
Ds_rdsh	RDSH VMs	2	Enabled	Enabled	Disabled	Disabled
Ds_vgpu	vGPU VMs	2	Enabled	Enabled	Disabled	Disabled

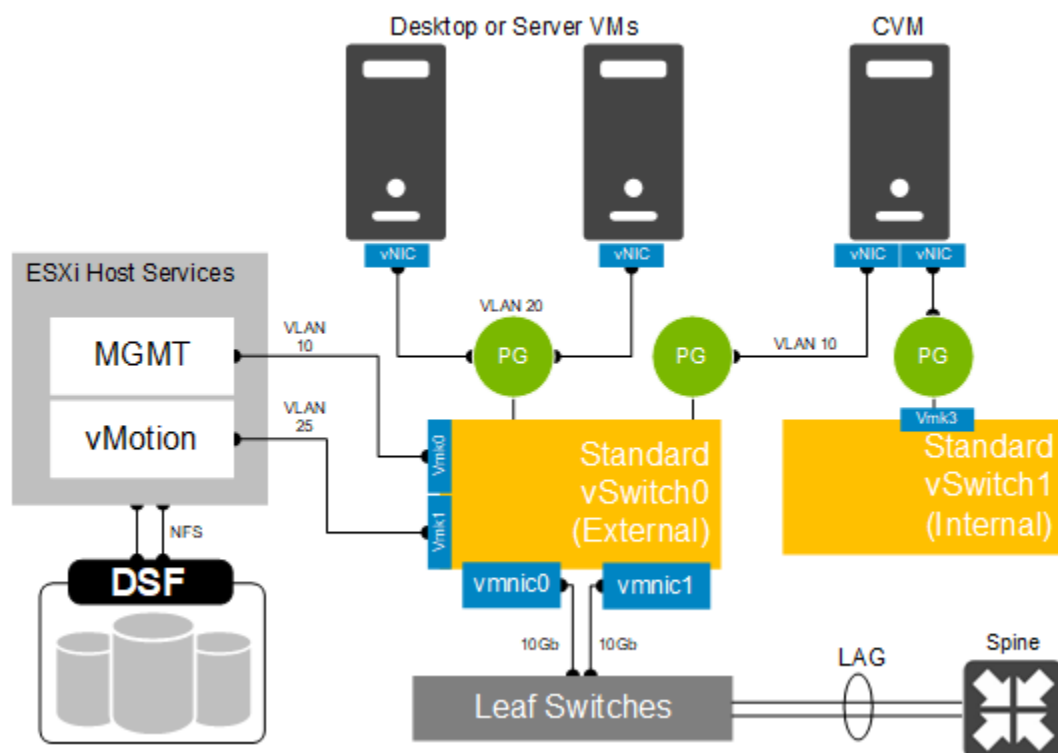
5.3 Virtual networking

The network configuration for the Dell EMC XC Series appliances utilizes a 10Gb converged infrastructure model. All required VLANs will traverse 2 x 10Gb NICs configured in an active/active team. For larger scaling it is recommended to separate the infrastructure management VMs from the compute VMs to aid in predictable compute host scaling. The following outlines the suggested VLAN requirements for the Compute and Management hosts in this solution model:

- Compute hosts
 - Management VLAN: Configured for hypervisor infrastructure traffic – L3 routed via spine layer
 - Live Migration VLAN: Configured for Live Migration traffic – L2 switched via leaf layer
 - VDI VLAN: Configured for VDI session traffic – L3 routed via spine layer
- Management hosts
 - Management VLAN: Configured for hypervisor Management traffic – L3 routed via spine layer
 - Live Migration VLAN: Configured for Live Migration traffic – L2 switched via leaf layer
 - VDI Management VLAN: Configured for VDI infrastructure traffic – L3 routed via spine layer
- An iDRAC VLAN is configured for all hardware management traffic – L3 routed via spine layer

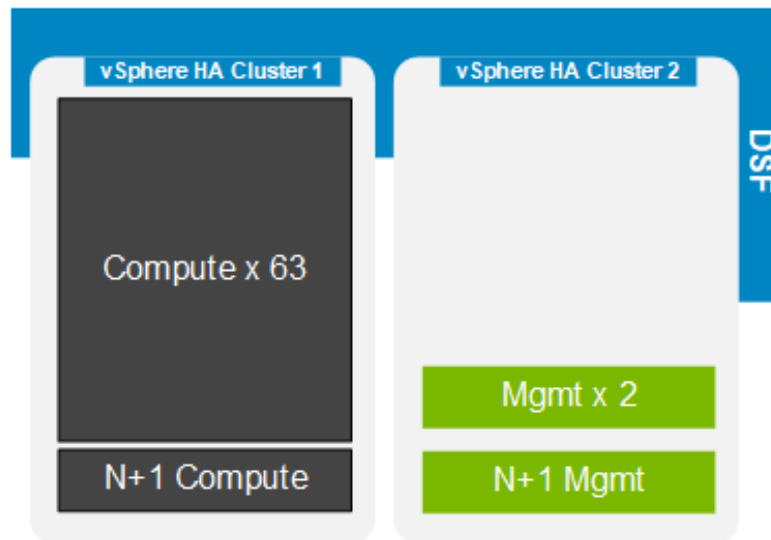
5.3.1 vSphere

Both the compute and management host network configuration consists of a standard vSwitch teamed with 2 x 10Gb physical adapters assigned to VMNICS. The CVM connects to a private internal vSwitch to communicate directly with the hypervisor as well as the standard external vSwitch to communicate with other CVMs in the cluster. All VDI infrastructure VMs connect through the primary port group on the external vSwitch.

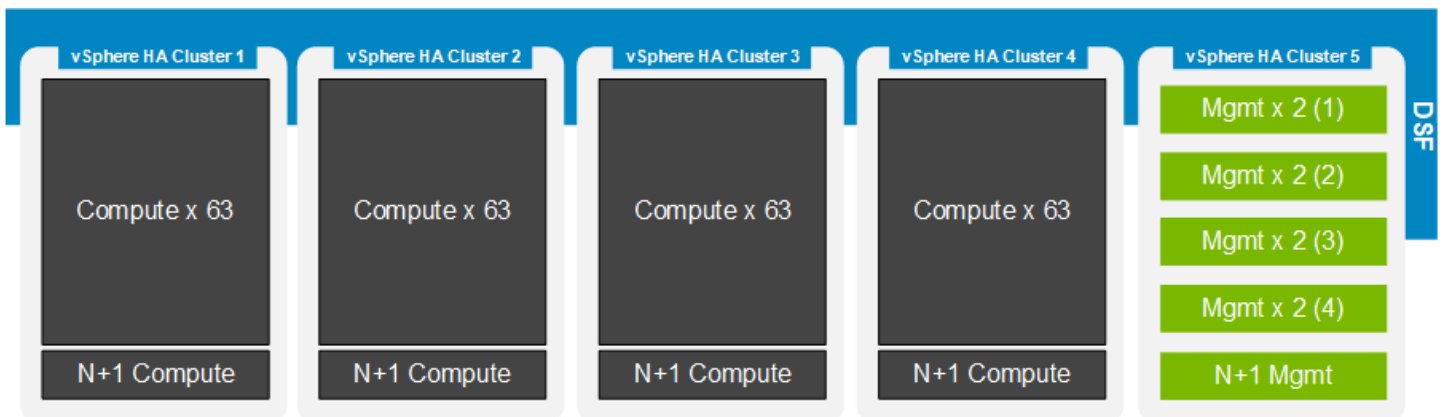


5.4 Scaling guidance

Each component of the solution architecture scales independently according to the desired number of supported users. Additional appliance nodes can be added at any time to expand the Nutanix SDS pool in a modular fashion. While there is no scaling limit of the Nutanix architecture itself, practicality might suggest scaling pods based on the limits of hypervisor clusters (64 nodes for vSphere). Isolating management and compute to their own HA clusters provides more flexibility with regard to scaling and functional layer protection while stretching the DSF cluster namespace between them.



Another option is to design a large single contiguous NDFS namespace with multiple hypervisor clusters within to provide single pane of glass management. For example, portrayed below is a large-scale user environment segmented by vSphere HA cluster and broker farm. Each farm compute instance is segmented into an HA cluster with a hot standby node providing N+1, served by a dedicated pair of management nodes in a separate HA cluster. This provides multiple broker farms with separated HA protection while maintaining a single NDFS cluster across all nodes.




- The components are scaled either horizontally (by adding additional physical and virtual servers to the server pools) or vertically (by adding virtual resources to the infrastructure)
- Eliminate bandwidth and performance bottlenecks as much as possible
- Allow future horizontal and vertical scaling with the objective of reducing the future cost of ownership of the infrastructure.

Component	Metric	Horizontal scalability	Vertical scalability
Virtual Desktop Host/Compute Servers	VMs per physical host	Additional hosts and clusters added as necessary	Additional RAM or CPU compute power
View Composer	Desktops per instance	Additional physical servers added to the Management cluster to deal with additional management VMs.	Additional network and I/O capacity added to the servers
View Connection Servers	Desktops per instance	Additional physical servers added to the Management cluster to deal with additional management VMs.	Additional VCS Management VMs.
RDSH Servers	Desktops per instance	Additional virtual servers added as necessary	Additional physical servers to host virtual RDSH servers.
VMware vCenter	VMs per physical host and/or ESX hosts per vCenter instance	Deploy additional servers and use linked mode to optimize management	Additional vCenter Management VMs.
Database Services	Concurrent connections, responsiveness of reads/writes	Migrate databases to a dedicated SQL server and increase the number of management nodes	Additional RAM and CPU for the management nodes
File Services	Concurrent connections, responsiveness of reads/writes	Split user profiles and home directories between multiple file servers in the cluster. File services can also be migrated to the optional NAS device to provide high availability.	Additional RAM and CPU for the management nodes

5.5 Solution high availability

High availability (HA) is offered to protect each architecture solution layer, individually if desired. Following the N+1 model, additional ToR switches are added to the Network layer and stacked to provide redundancy as required, additional compute and management hosts are added to their respective layers, vSphere clustering is introduced in both the management and compute layers, SQL is configured for AlwaysOn or clustered and NetScaler is leveraged for load balancing.


The HA options provide redundancy for all critical components in the stack while improving the performance and efficiency of the solution as a whole.



- Additional switches added to the existing thereby equally spreading each host's network connections across multiple switches.
- Additional ESXi hosts added in the compute or management layers to provide N+1 protection.
- Applicable VMware infrastructure server roles are duplicated and spread amongst management host instances where connections to each are load balanced via the addition of virtual NetScaler appliances.
- SQL Server databases also are protected through the addition and configuration of an "AlwaysOn" Failover Cluster Instance or Availability Group.

Please refer to these links for more information: [SQL Server AlwaysOn Availability Groups](#) and [Windows Server Failover Clustering with SQL Server](#)

5.6 Communication flow for Horizon



6 Solution performance and testing

6.1 Summary

At the time of publication, these are the available density recommendations per appliance/node. Please refer to the Platform Configurations section for hardware specifications.

NOTE: All-flash configurations yield the same user densities with our test methodology since processor and memory resources are exhausted before storage resources are impacted.

Standard user density summary

Host Config*	Hypervisor	Broker & Provisioning	Workload	Template	User Density*
XC Series C7	ESXi 6.5 U1	Horizon 7 Linked Clones	Knowledge Worker	Windows 10 & Office 2016	145/150
XC Series C7	ESXi 6.5 U1	Horizon 7 Linked Clones	Power Worker	Windows 10 & Office 2016	123/125
XC Series C7	ESXi 6.5 U1	Horizon 7 Linked Clones	Power Worker	Windows 10 & Office 2016	48 + 105**

The User Density column shows the Management Host's desktop density and the Compute Hosts' density. This optional configuration is good for POCs or small deployments looking to maximize user density. The first value in the User Density column shows the Management Host's density if desktop VMs were to be deployed in addition to the mgmt. VMs, the second value represents the Compute Hosts density only. (ex. 145/150). Proper HA design should be part of any production deployment!

The Nutanix CVM reserved 10000 MHz of CPU on each host, and uses 12 cores with high priority. This was a factor in reducing density to a level well below the usual target CPU and Memory usage for a vSphere cluster.

*User density values are reported here as management + desktop VMs/ desktop VMs only. In smaller 3-4 node clusters, this shows the impact of running desktop VMs in addition to mgmt. infra VMs on the mgmt. node to maximize user density. For example, the first result of 190/200 indicates 190 desktop VMs on the mgmt. host and 200 desktop VMs per node on the compute hosts, totaling 790 users on this 3-node cluster.

**For graphics acceleration testing we performed assigned 1GB framebuffers both with vGPUs VMs only as well as vGPU VMs + standard non-accelerated VMs. Here 48 represents the vGPU VMs, 105 represents the non-accelerated VMs also running on the same server node.

6.2 Test and performance analysis methodology

6.2.1 Testing process

In order to ensure the optimal combination of end-user experience (EUE) and cost-per-user, performance analysis and characterization (PAAC) on Dell EMC VDI solutions is carried out using a carefully designed,

holistic methodology that monitors both hardware resource utilization parameters and EUE during load-testing.

Login VSI is currently the load-generation tool used during PAAC of Dell EMC solutions. Each user load is tested against multiple runs. First, a pilot run to validate that the infrastructure is functioning and valid data can be captured, and then, subsequent runs allowing correlation of data.

At different times during testing, the testing team will complete some manual “User Experience” Testing while the environment is under load. This will involve a team member logging into a session during the run and completing tasks similar to the User Workload description. While this experience will be subjective, it will help provide a better understanding of the end user experience of the desktop sessions, particularly under high load, and ensure that the data gathered is reliable.

6.2.1.1 Load generation

Login VSI by Login Consultants is the de-facto industry standard tool for testing VDI environments and server-based computing (RDSH environments). It installs a standard collection of desktop application software (e.g. 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. Once the user is connected, the workload is started via a logon script which starts the test script once the user environment is configured by the login script. Each launcher system can launch connections to a number of ‘target’ machines (i.e. VDI desktops). The launchers and Login VSI environment are configured and managed by a centralized management console.

Additionally, the following login and boot paradigm is used:

- Users are logged in within a login timeframe of 1 hour. Exception to this login timeframe occurs when testing low density solutions such as GPU/graphics based configurations. With those configurations, users are logged on every 10-15 seconds.
- All desktops are pre-booted in advance of logins commencing.
- All desktops run an industry-standard anti-virus solution. Windows Defender is used for Windows 10 due to issues implementing McAfee.

6.2.1.2 Profiles and workloads

It's important to understand user workloads and profiles when designing a desktop virtualization solution in order to understand the density numbers that the solution can support. At Dell EMC, we use five workload / profile levels, each of which is bound by specific metrics and capabilities with two targeted at graphics-intensive use cases. We will present more detailed information in relation to these workloads and profiles below but first it is useful to define the terms “profile” and “workload” as they are used in this document.

- **Profile:** This is the configuration of the virtual desktop - number of vCPUs and amount of RAM configured on the desktop (i.e. available to the user).
- **Workload:** This is the set of applications used by performance analysis and characterization (PAAC) of Dell EMC VDI solutions (e.g. Microsoft Office applications, PDF Reader, Internet Explorer etc.)

Load-testing on each profile is carried out using an appropriate workload that is representative of the relevant use case and summarized in the table below:

Profile to workload mapping

Profile Name	Workload
Task Worker	Login VSI Task worker
Knowledge Worker	Login VSI Knowledge worker
Power Worker	Login VSI Power worker

Login VSI workloads are summarized in the sections below. Further information for each workload can be found on Login VSI's [website](#).

Login VSI Task Worker Workload

The Task Worker workload runs fewer applications than the other workloads (mainly Excel and Internet Explorer with some minimal Word activity, Outlook, Adobe, copy and zip actions) and starts/stops the applications less frequently. This results in lower CPU, memory and disk IO usage.

Login VSI Knowledge Worker Workload

The Knowledge Worker workload is designed for virtual machines with 2vCPUs. This workload contains the following activities:

- Outlook, browse messages.
- Internet Explorer, browse different webpages and a YouTube style video (480p movie trailer) is opened three times in every loop.
- Word, one instance to measure response time, one instance to review and edit a document.
- Doro PDF Printer & Acrobat Reader, the Word document is printed and exported to PDF.
- Excel, a very large randomized sheet is opened.
- PowerPoint, a presentation is reviewed and edited.
- FreeMind, a Java based Mind Mapping application.
- Various copy and zip actions.

Login VSI Power Worker Workload

The Power Worker workload is the most intensive of the standard workloads. The following activities are performed with this workload:

- Begins by opening four instances of Internet Explorer which remain open throughout the workload.
- Begins by opening two instances of Adobe Reader which remain open throughout the workload.
- There are more PDF printer actions in the workload as compared to the other workloads.
- Instead of 480p videos a 720p and a 1080p video are watched.
- The idle time is reduced to two minutes.
- Various copy and zip actions.

6.2.2 Resource monitoring

The following sections explain respective component monitoring used across all Dell EMC solutions where applicable.

6.2.2.1 GPU resources

ESXi hosts

For gathering of GPU related resource usage, a script is executed on the ESXi host before starting the test run and stopped when the test is completed. The script contains NVIDIA System Management Interface commands to query each GPU and log GPU utilization and GPU memory utilization into a .csv file.

ESXi 6.5 and above includes the collection of this data in the vSphere Client/Monitor section. GPU processor utilization, GPU temperature, and GPU memory utilization can be collected the same was as host CPU, host memory, host Network, etc.

6.2.2.2 Microsoft Performance Monitor

Microsoft Performance Monitor is used for Hyper-V based solutions to gather 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 (when multiple are tested). While the report does not include specific performance metrics for the Management host servers, these servers are monitored during testing to ensure they are performing at an expected performance level with no bottlenecks.

6.2.2.3 VMware vCenter

VMware vCenter is used for VMware vSphere-based solutions to gather 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 (when multiple are tested). While the report does not include specific performance metrics for the Management host servers, these servers are monitored during testing to ensure they are performing at an expected performance level with no bottlenecks.

6.2.3 Resource utilization

Poor end-user experience is one of the main risk factors when implementing desktop virtualization but a root cause for poor end-user experience is resource contention: hardware resources at some point in the solution have been exhausted, thus causing the poor end-user experience. In order to ensure that this does not happen, PAAC on Dell EMC solutions monitors the relevant resource utilization parameters and applies relatively conservative thresholds as shown in the table below. Thresholds are carefully selected to deliver an optimal combination of good end-user experience and cost-per-user, while also providing burst capacity for seasonal / intermittent spikes in usage. Utilization within these thresholds is used to determine the number of virtual applications or desktops (density) that are hosted by a specific hardware environment (i.e. combination of server, storage and networking) that forms the basis for a Dell EMC RA.

Parameter	Pass/Fail Threshold
Physical Host CPU Utilization (ESXi hypervisor)*	100%
Physical Host Memory Utilization	85%
Network Throughput	85%
Storage IO Latency	20ms

*Turbo mode is enabled; therefore, the CPU threshold is increased as it will be reported as over 100% utilization when running with turbo.

6.3 Test configuration details

The following components were used to complete the validation testing for the solution:

XC Series hardware and software test components

Component	Description/Version
Hardware platform(s)	Dell EMC XC740xd C7
Hypervisor(s)	VMware vSphere ESXi 6.5 U1
Broker technology	Horizon 7.3.2
Broker database	Microsoft SQL 2014
Management VM OS	Microsoft Windows Server 2012 R2
Virtual desktop OS	Microsoft Windows 10 Enterprise 64-bit
Office application suite	Microsoft Office 2016 Professional Plus
Login VSI test suite	4.1.25

6.3.1 Compute VM configurations

The following table summarizes the compute VM configurations for the various profiles/workloads tested.

ESXi Desktop VM specifications

User Profile	vCPUs	ESXi Memory Configured	ESXi Memory Reservation	Screen Resolution	Operating System
Task Worker	2	2GB	1GB	1280 X 720	Windows 10 Enterprise 64-bit
Knowledge Worker	2	3GB	1.5GB	1920 X 1080	Windows 10 Enterprise 64-bit
Power Worker	2	4GB	2GB	1920 X 1080	Windows 10 Enterprise 64-bit
Graphics Density Configuration	2	4 GB	4GB	4GB	Windows 10 Enterprise 64-bit
Graphics Performance Configuration	4	8GB	8GB	8GB	Windows 10 Enterprise 64-bit

Profile to workload mapping

Profile Name	Workload	OS Images
Standard	Login VSI Task worker	Shared
Enhanced	Login VSI Knowledge worker	Shared
Professional	Login VSI Power worker	Shared + Profile Virtualization
Graphics Density Configuration	Login VSI Power worker with ProLibrary	Persistent
Graphics Performance Configuration	Custom specific	Persistent

6.4 Standard VDI test results and analysis

The following table summarizes the test results for the compute hosts using the various workloads and configurations.

In typical VDI density tests we attempt to approach “full loading” of the compute hosts at some predefined threshold of CPU or Memory. For vSphere the threshold for CPU is defined as 100%. In this case it was not feasible to achieve full loading at that threshold due to the resource reservations imposed on each host by the Nutanix CVM. Each CVM reserved 10000 MHz on each host, out of a total of 80000 MHz nominally available according to the hypervisor (2x sockets, 20x Cores/socket, 2000 MHz/core). This left 70000 MHz available for scheduling desktop and management VM CPU activity. Also, the CVM is assigned 12 vCPUs that all must be scheduled together with high priority. Between the reservation and the core-scheduling, it was necessary to respect a lower CPU threshold to avoid a higher-than-desired CPU “readiness” on the desktop and management VMs, which caused excessive application latency and very bad user experience despite not fully loading the host.

These problems were evident in the VSI results in the chart named “VSIMax per computer” where we could see that a small subset of user sessions were experiencing extreme latency on the order of 24000ms in many cases. While the overall performance of the system appeared to be well within constraints, the number of VMs contending for CPU time saturated the CPUs such that many desktop VMs had to wait behind the CVM and other VMs. These issues required us to adjust our threshold for “full loading”.

According to this new threshold, the available CPU for desktops is 70000 MHz, or 87.5% of a host’s total. If we restrict to 95% of this, we can use as much as 83% of the host CPU for desktops before causing a steep decline in user experience. We found that the CVM’s themselves consume 7-8% of the host CPU during steady state usage.

The number of desktops resident on the Management host was slightly less than on the Compute hosts. The relative densities are shown in the “Density Per Host” column below.

Test result summary

Platform Config	Hypervisor	Broker & Provisioning	Login VSI Workload	Density Per Host	Avg CPU	Avg Mem Consumed	Avg IOPS / User	Avg Net Mbps / User
XC Series C7	ESXi 6.5U1	Horizon 7 Linked Clones	Knowledge Worker	145/150	76%	483GB	2.11	5.30
XC Series C7	ESXi 6.5U1	Horizon 7 Linked Clones	Power Worker	123/125	69%	483GB	2.26	6.71

Density per Host: Density reflects number of users per compute host that successfully completed the workload test within the acceptable resource limits for the host. For clusters, this reflects the average of the density achieved for all compute hosts in the cluster.

Avg CPU: This is the average CPU usage over the steady state period. For clusters, this represents the combined average CPU usage of all compute hosts.

NOTE: On the latest Intel series processors, the ESXi host CPU metrics will exceed the reported 100% utilization for the host if Turbo Boost is enabled on the CPU (by default). Up to an additional 35% of CPU is available from the Turbo Boost feature but this additional CPU headroom is not reflected in the metrics where the performance data is gathered. Therefore, CPU usage for ESXi hosts is adjusted and a line indicating the potential performance headroom provided by Turbo boost is included in each CPU graph. Hyper-V CPU percentages take into account this additional capability.

Avg Consumed Memory: ESXi consumed memory is the amount of host physical memory granted within a host. For clusters, this is the average consumed memory across all compute hosts over the steady state period.

Avg Mem Active: For ESXi hosts, active memory is the amount of memory that is actively used, as estimated by VMKernel based on recently touched memory pages. For clusters, this is the average amount of guest “physical” memory actively used across all compute hosts over the steady state period.

Avg IOPS/User: IOPS calculated from the average Disk IOPS figure over the steady state period divided by the number of users.

Avg Net Mbps/User: Amount of network usage over the steady state period divided by the number of users. For clusters, this is the combined average of all compute hosts over the steady state period divided by the number of users on a host. The unit is Mega-bits per second, not Mega-Bytes per second (MBps).

6.4.1 XC740xd-C7

Refer to section 3 for hardware configuration details. Please note that the performance results reported below can be expected to apply to any platform with a C7 configuration.

The hardware configurations that were tested are summarized in the table(s) below.

XC Series hardware configuration for standard VDI

Enterprise Platform	Platform Config	CPU	Memory	RAID Ctlr	Drive Config	Network
XC740xd-24	C7	6138 Gold (20- Core 2.0 GHz)	768GB @2666 MT/s	HBA 330	2 x 120GB M.2 2x 960 GB SSD 4x 1.8TB HDD	4x Intel X710 rNDC

Compute and Management resources were split out with the following configuration across a three node Nutanix cluster and all test runs were completed with this configuration.


- Node 1 – Dedicated Management: vCenter Appliance, SQL Server, View Connection Server, View Composer, and Nutanix CVM
- Node 2 – Dedicated Compute, Nutanix CVM and User VMs only.
- Node 3 – Dedicated Compute, Nutanix CVM and User VMs only.

1GB networking was used for the deployment of the XC Series appliances only while 10GB networking is required for standard cluster operation.


Instead of dedicated nodes, Nutanix CVMs and VDI management roles were deployed on the cluster with desktop VMs, which reduced maximum density on that node. Each compute node was loaded with desktops to its maximum density; no failover capacity was reserved.

6.4.1.1 Knowledge Worker, 445 Total Users, ESXi 6.5, Horizon 7.3.2


In this workload test, the compute hosts each had hosted 150 desktop VMs, while the management host had 145 sessions in addition to the Horizon management VMs. The peak CPU Usage was 84% on one host during logon phase, while the Steady State average was 76% across all hosts. The relatively low steady state CPU usage was essential to avoid CPU scheduling problems that would have placed many of the desktops into the Ready-waiting state, causing VSI errors to a large number of sessions. The average CVM CPU usage during steady state was 7.9%, while the management VMs altogether used 1.2%.




The memory consumption averaged 483 GB in steady state across all hosts, and the peak usage on any host was 499 GB during the steady state phase. The peak usage was 65%, well below the 85% threshold, while the Steady State average usage for all hosts was 63%. There was no swapping or ballooning during the test run and each desktop consumed 3.0 GB, after accounting for CVM and management VM memory consumption. The CVM on each host consumed its full 32GB of reserved memory throughout the test run.




Active memory during the test run peaked at 330GB on the management host during the boot storm, and averaged 162GB during Steady state. Each desktop accounted for 0.86 GB of active memory usage after deducting CVM and management VM active memory. The CVM on each host used a full 32 GB of active memory throughout the test run.




Network usage peaked at 874 Mbps during Boot storm on one host, and the average network usage for all hosts was 786 Mbps during Steady State. Each desktop produced network throughput of 5.30 Mbps in steady state.



The peak Cluster IOPS for the test run was 9090 IOPS during the Boot Storm phase, while the average in Steady State was 939 IOPS. Based on these numbers each user session generated 2.11 IOPS in steady state.




The peak Cluster IO Latency was 0.6 ms during the Steady State. The average Cluster IO latency during steady state was 0.5 ms. The highest IO Latency on any host was 0.7 ms during the Boot Storm. The chart clearly shows a very steady and very low level of IO Latency throughout the test run.



The baseline performance of 731 indicated that the user experience for this test run was Very Good. The Index average reached 1122, well below the VSIMax threshold of 1732. Although there is considerable space left over for additional sessions according to the VSI graph, adding more session would have resulted in excessive session errors. The irregular shape of the VSI Index Average curve below shows the effect of the excess session latency that occurred on a subset of user sessions.


Login VSI Baseline	VSI Index Average	VSIMax Reached	VSI Threshold
731	1122	NO	1732




6.4.1.2 Power Worker, 375 Total Users, ESXi 6.5, Horizon 7.3.2

In this workload test run each compute host had 125 user sessions while the designated management host had the full set of management VMs plus 125 desktops. The peak CPU Usage was 86% on one host during Steady State phase, while the Steady State average was 69% across all hosts. The relatively low steady state CPU usage was essential to avoid session latency problems that would have caused excessive VSI errors.


The CVM's on each host averaged 7.3% CPU usage during steady state. The management VMs used only 1.2% CPU on the management host in steady state.




The memory consumption averaged 483 GB in steady state across all hosts, and the peak usage on any host was 554 GB during the Logoff phase. The peak usage was 72%, well below the 85% threshold, while the Steady State average usage for all hosts was 63%. There was no swapping or ballooning during the test run and each desktop consumed 3.54 GB after accounting for CVM and management VM memory consumption.




Active memory usage reached a maximum of 416 GB on the Compute B host during the Boot Storm, and the average Steady State usage for all hosts was 154 GB. Each desktop used 0.96 GB of active memory after deducting for CVM and management VM usage. The CVM used its full 32GB of active memory throughout the test run.




Network usage peaked at 1041 Mbps on one host during Steady State phase, and the average network usage for all hosts during Steady State was 839 Mbps. Each desktop accounted for 6.71 Mbps in Steady State.



The peak Cluster IOPS for the test run was 6585 IOPS during the Boot Storm phase, while the average in Steady State was 848 IOPS. Based on these numbers each user session generated 2.26 IOPS during Steady State.




The peak Cluster IO Latency was 0.6 ms during the Boot Storm, while the peak on any host was 0.6 ms. The average Cluster IO latency during steady state was 0.5 ms. The chart clearly shows a very steady and very low level of IO Latency throughout the test run.



The baseline performance of 725 indicated that the user experience for this test run was Very Good. The Index average reached 988, well below the threshold of 1725. Although the difference between the VSIMax and the average would seem to indicate that more desktops could be used, the CPU limitations described above would have reduced the user experience dramatically. The irregular shape of the VSI Index Average curve below shows the effect of the excess session latency that occurred on a subset of user sessions.

Login VSI Baseline	VSI Index Average	VSIMax Reached	VSI Threshold
725	988	NO	1725



6.5 vGPU test results and analysis

All test results graphs include the performance of the platform during the deletion and recreation of the linked clone virtual machines after all users log off when the test run has completed. The different phases of the test cycle are displayed in the test results graphs later in this document as 'Boot Storm', 'Logon', 'Steady State' and 'Logoff'.

We tested three scenarios for graphics acceleration on Windows10 VMs: vGPU compute-only using M60-1Q, vGPU + standard non-vGPU VMs on the same node collocated and 48 standard non-vGPU VMs to compare system performance against the accelerated variants. Please note that all scenarios consist of the minimum three-node cluster with one mgmt. node and only one compute node active for these tests. Since GPUs were only added to a single host, we performed all vGPU testing against this one host. Hence you will see the 2nd compute node marked at "not used" in some of the graphs that follow.

The following table summarizes the test results for the various workloads and configurations.

Hyper-visor	Provisioning	Login VSI Workload	Density Per Host	Remote Display Protocol	Avg' CPU %	Avg' GPU %	Avg' Memory Consumed GB	Avg' Memory Active GB	Avg' IOPS/User	Avg' Net Mbps/User
ESXi	Linked Clones	Power Worker	48 vGPU	PCoIP	41%	40%	239 GB	224 GB	4.8	5 Mbps
ESXi	Linked Clones	Power Worker	105 Std + 48 vGPU	PCoIP	95 %	31%	656 GB	331 GB	3.6	5.2 Mbps
ESXi	Linked Clones	Power Worker	48 Std	PCoIP	32%	-	224 GB	79 GB	4.3	5.5 Mbps

CPU Usage. The figure shown in the table, 'Avg' CPU %', is the combined average CPU usage of all compute hosts over the steady state period.

GPU Usage. The figure shown in the table, 'Avg' GPU %', is the average GPU usage of all hosts containing GPU cards over the steady state period.

Consumed Memory. Consumed memory is the amount of host physical memory consumed by a virtual machine, host, or cluster. The figure 'Avg' Memory Consumed GB' in the table is the average consumed memory across all compute hosts over the steady state period.

Active Memory. Active Memory is the amount of memory that is actively used, as estimated by VMkernel based on recently touched memory pages. The figure 'Avg' Memory Active GB' in the table is the average amount of guest "physical" memory actively used across the compute (and or management) hosts over the steady state period.

Disk IOPS. Disk IOPS are calculated from the Cluster Disk IOPS steady state average divided by the number of users to produce the 'IOPS / User' figure.

Network Usage. Network Usage per User. The figure shown in the table 'Avg' Net Mbps/User' is the Network Usage average of the hosts over the steady state period divided by the number of users per host in Megabits per second.

CPU usage for ESXi hosts is adjusted to account for the fact that on the latest Intel series processors, the ESXi host CPU metrics will exceed the rated 100% for the host if Turbo Boost is enabled (by default). An additional 35% of CPU is available from the Turbo Boost feature when all cores are active, but this additional CPU headroom is not reflected in the VMware vSphere metrics where the performance data is gathered from. As a result, a line indicating the potential performance headroom provided by Turbo boost is included in each CPU graph.

Without the inclusion of the turbo there is a total of 80,000 MHz available for Desktops, with Turbo boost the total available MHz value is 108,000 MHz.

The user virtual machines were created using VMware Horizon Linked Clones. The virtual machine desktops used local logon profiles and each user was assigned the same desktop for all logons. The vGPU enabled VM's used Windows 10 Enterprise and had the NVIDIA GRID drivers for Windows 10 installed and aligned with the Login VSI 4.1 virtual machine configuration. Office 2016 was used as the office suite and each virtual

machine's virtual disk sized at 60 GB. Workload configuration of the user virtual machines and is shown in the table below.

User Workload Profile	vCPUs	Memory GB	Reserved Memory	vGPU Profile	OS Bit Level	HD Size GB	Screen Resolution
Power Worker	2	4GB	4 GB	M60-1Q	Windows 10 64 Bit	60 GB	1920 x 1080

6.5.1 XC740xd-C7 with Tesla M60

Refer to the section 3 for hardware configuration details. GPUs can only be added to the 24-disk variant of the XC740xd.

XC Series hardware configuration for vGPU

Enterprise Platform	CPU	Memory	RAID Ctrlr	Drive Config	Network	GPU
XC740xd-24	6138 Gold (20- Core 2.0 GHz)	768GB @2666 MT/s	Dell HBA 330 Adapter	2 x 120GB M.2 2x 960 GB SSD 4x 1.8TB HDD	Intel 10GbE 4P X710 rNDC	3 X NVIDIA M60 GPU cards installed in one host.

Compute and Management resources were split out with the following configuration across a three-node cluster and all test runs were completed with this configuration.


- Node 1 –XC740xd-24 – Dedicated Management.
- Node 2 – XC740xd-24– Dedicated Compute (Unused for hosting VM's for this testing).
- Node 3 – XC740xd-24– Dedicated Compute with 3 X M60 GPU Cards installed.

10GB networking was used for all PAAC testing. The host containing the M60 GPU cards had the appropriate ESXi drivers for this card installed.


6.5.1.1 Power Worker, 48 vGPU users, ESXi 6.5, Horizon 7.3.2

The GPU enabled Compute Host was populated with 48 vGPU enabled virtual machines and used the NVIDIA M60-1Q profile. With all user virtual machines powered on and before starting test, the CPU usage was approximately 8%.


The below graph shows the performance data for 48 user sessions on the GPU enabled Compute host and also the performance of the dedicated Management host. The CPU reaches a steady state average of 41% during the test cycle when all 48 users are logged on to the GPU enabled Compute host.




The GPU Metrics were gathered from the vSphere web client and the GPU Profiler application was run during a test session on one of the VM's to determine the framebuffer and vGPU usage. The GPU usage during the steady state period averaged approximately 40% and reached a peak usage of 44% with the Power Worker workload.




Taken from a single VM for a representative sampling, the framebuffer in use under load averaged ~50%.




In regards 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 239 GB with active memory usage reaching a max of 224 GB. There were no variations in memory usage throughout the test as all vGPU enabled VM memory was reserved. There was no memory ballooning or swapping on either host.







Network bandwidth is not an issue on this test run with a steady state peak of approximately 286 Mbps on the Compute / GPU Host. The busiest period for network traffic was during the Boot Storm phase during the reboot of all the VM's before testing started. The Compute / GPU host reached a peak of 1,448 Mbps during the Boot Storm.




The IOPS graphs and numbers are taken from the Nutanix Prism web console and they clearly display the boot storm, the initial logon of the desktops then the steady state and finally the logoff phase. The graphs show IOPS data for the individual hosts in the cluster and for the cluster as a whole.

The cluster reached a maximum of 3,669 Disk IOPS during the reboot of all the VM's before test start and 657 IOPS at the start of steady state. The Compute / GPU host reached a peak of 2,325 Disk IOPS during the reboot of all the VM's and 351 at the start of steady state.



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.




Notes:

- As indicated above, the CPU graphs do not take into account the extra 35% of CPU resources available through the Intel Xenon Gold 6138 processors turbo feature.
- 768 GB of memory installed on each node is just about sufficient for the number of desktops. With memory usage going close to maximum, no extra desktops could have been accommodated for this configuration.
- The PCoIP remote display protocol was used during testing.
- There were no disk latency issues during testing.


6.5.1.2 Power Worker, 48 vGPU + 105 standard users, ESXi 6.5, Horizon 7.3.2

The GPU enabled Compute Host was populated with 48 vGPU enabled virtual machines using the NVIDIA M60-1Q profile. In addition, this same host was populated with 105 standard non-vGPU VM's adhering to the standard Power worker profile. With all user virtual machines powered on and before starting test, the CPU usage on the GPU enabled Compute host was approximately 14%.



The below graph shows the performance data for 153 total user sessions on the GPU enabled Compute host and also the performance of the dedicated Management host. The CPU reaches a steady state average of 95% during the test cycle when all 153 users are logged on to the GPU enabled Compute host.




The GPU Metrics were gathered from the vSphere Web client. The GPU usage during the steady state period averaged approximately 31% and reached a peak usage of 41% with the Power worker workload.



In regards 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 663 GB with active memory usage reaching a max of 635 GB. All memory on the vGPU enabled VM's was reserved and half of the memory on the standard VM's was reserved. There was no memory ballooning or swapping on either host.





Network bandwidth is not an issue on this test run with a steady state peak of approximately 934 Mbps on the Compute / GPU Host. The busiest period for network traffic was during the Boot Storm phase during the reboot of all the VM's before testing started. The Compute / GPU host reached a peak of 1,633 Mbps during the Boot Storm.




The IOPS graphs and IOPS numbers are taken from the Nutanix Prism web console and they clearly display the boot storm, the initial logon of the desktops then the steady state and finally the logoff phase. The graphs show IOPS data for the individual hosts in the cluster and for the cluster as a whole.

The cluster reached a maximum of 8,497 Disk IOPS during the reboot of all the VM's before test start and 725 IOPS at the start of steady state. The Compute / GPU host reached a peak of 6,190 Disk IOPS during the reboot of all the VM's and 370 IOPS at the start of steady state.





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.




Notes:

- As indicated above, the CPU graphs do not take into account the extra 35% of CPU resources available through the Intel Xenon Gold 6138 processors turbo feature.
- The PCoIP remote display protocol was used during testing.
- There were no disk latency issues during testing.
- 48 vGPU enabled VM's and 105 standards Power worker VM's were co-located on the same GPU enabled Compute host for this test run.



6.5.1.3 Power Worker, 48 standard users (non-vGPU), ESXi 6.5, Horizon 7.3.2

The GPU enabled Compute Host was populated with 48 standard non-vGPU enabled virtual machines and used to compare the performance of 48 vGPU enabled VM's on the same hardware. With all user virtual machines powered on and before starting test, the Compute host's CPU usage was approximately 6%.


The below graph shows the performance data for 48 user sessions on the Management and Compute hosts. The CPU reaches a steady state average of 32% during the test cycle when all 48 users are logged on to the Compute host.



In regards to memory consumption for this test run there were no constraints on the Management or Compute hosts. Of a total of 768 GB available memory per node, the Compute host reached a maximum memory consumption of 238 GB with active memory usage reaching a max of 222 GB. There was no memory ballooning or swapping on either host.





Network bandwidth is not an issue on this test run with a steady state peak of approximately 331 Mbps on the Compute host. The busiest period for network traffic was during the Boot Storm phase during the reboot of all the VM's before testing started. The Compute host reached a peak of 989 Mbps during the Boot Storm.




The IOPS graphs and IOPS numbers are taken from the Nutanix Prism web console and they clearly display the boot storm, the initial logon of the desktops then the steady state and finally the logoff phase. The graphs show IOPS data for the individual hosts in the cluster and for the cluster as a whole.

The cluster reached a maximum of 5,432 Disk IOPS during the reboot of all the VM's before test start and 410 IOPS at the start of steady state. The Compute host reached a peak of 3,276 Disk IOPS during the reboot of all the VM's and 225 IOPS at the start of steady state.





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.



Notes:

- As indicated above, the CPU graphs do not take into account the extra 35% of CPU resources available through the Intel Xenon Gold 6138 processors turbo feature.
- The PCoIP remote display protocol was used during testing.
- There were no disk latency issues during testing.

A Related resources

See the following referenced or recommended resources:

- The [Dell EMC Cloud-Client Computing Solutions for VMware Tech Center page](#) which includes this RA and other VMware Horizon based RAs.
- Dell EMC Tech Center for XC Series:
<http://en.community.dell.com/techcenter/storage/w/wiki/11454.dell-emc-xc-series-hyper-converged-solution>
- <http://www.dell.com/XCseriesSolutions> for Dell EMC XC Series white papers.
- www.Dell.com/xcseriesmanuals for deployment guides (XC Xpress only), manuals, support info, tools, and videos.