Dell XC Web-Scale Converged Appliance for VMware Horizon®

Dell Engineering
September 2016

A Dell Appliance Architecture

Revisions

Date	Description
March 2015	Initial release
Sep 2015	Document overhaul, XC730 + vGPU Addition
April 2016	New platforms, new art, new GRID 2.0 architecture
September 2016	Updates to Endpoints section

Contents

1 Introduction

1.1 Purpose

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.

1.2 Scope

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.

1.3 What's New

- Introduce support for NVIDIA GRID 2.0
- Introduce support for additional XC platforms
- Introduce support for Intel E5-2600v4 CPUs (Broadwell)
- Updated Acropolis and appliance architectures

2 Solution Architecture Overview

2.1 Introduction

The Dell XC series delivers an out-of-the-box infrastructure solution for virtual desktops that eliminates the high cost, variable performance, and extensive risk of conventional solutions. The Nutanix™ web-scale converged infrastructure is a turnkey solution that comes ready to run your VDI solution of choice. The Nutanix platform's unique architecture allows enterprises to scale their virtual desktops from 50 to tens of thousands of desktops in a linear fashion, providing customers with a simple path to enterprise deployment with the agility of public cloud providers.

2.2 Nutanix Enterprise Cloud Platform Overview

Nutanix delivers a hyperconverged infrastructure solution purpose-built for virtualization and cloud environments. This solution brings the performance and economic benefits of web-scale architecture to the enterprise through the Nutanix enterprise cloud platform, which is composed of two product families—Nutanix Acropolis and Nutanix Prism.

Attributes of this solution include:

- Storage and compute resources hyperconverged on x86 servers.
- System intelligence located in software.
- Data, metadata, and operations fully distributed across entire cluster of x86 servers.
- Self-healing to tolerate and adjust to component failures.
- API-based automation and rich analytics.

Nutanix Acropolis can be broken down into three foundational components: the Distributed Storage Fabric (DSF), the App Mobility Fabric (AMF), and AHV. Prism provides one-click infrastructure management for virtual environments running on Acropolis. Acropolis is hypervisor agnostic, supporting two third-party hypervisors—ESXi and Hyper-V—in addition to the native Nutanix hypervisor, AHV.

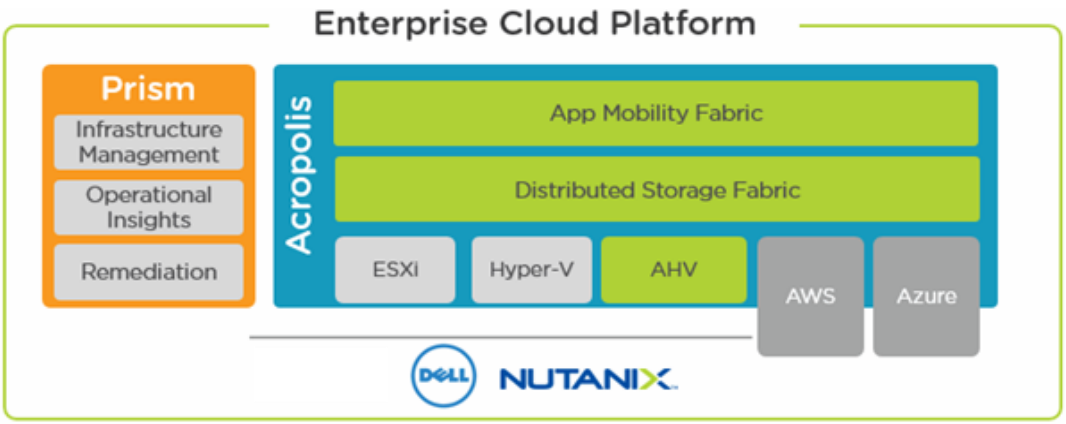

2.3 Distributed Storage Fabric

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.

2.4 Nutanix Acropolis Architecture

Acropolis does not rely on traditional SAN or NAS storage or expensive storage network interconnects. It combines highly dense storage and server compute (CPU and RAM) into a single platform building block. Each building block is based on industry-standard Intel processor technology and delivers a unified, scale-out, shared-nothing architecture with no single points of failure.

The Nutanix solution has no LUNs to manage, no RAID groups to configure, and no complicated storage multipathing to set up. 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 the fastest possible performance.

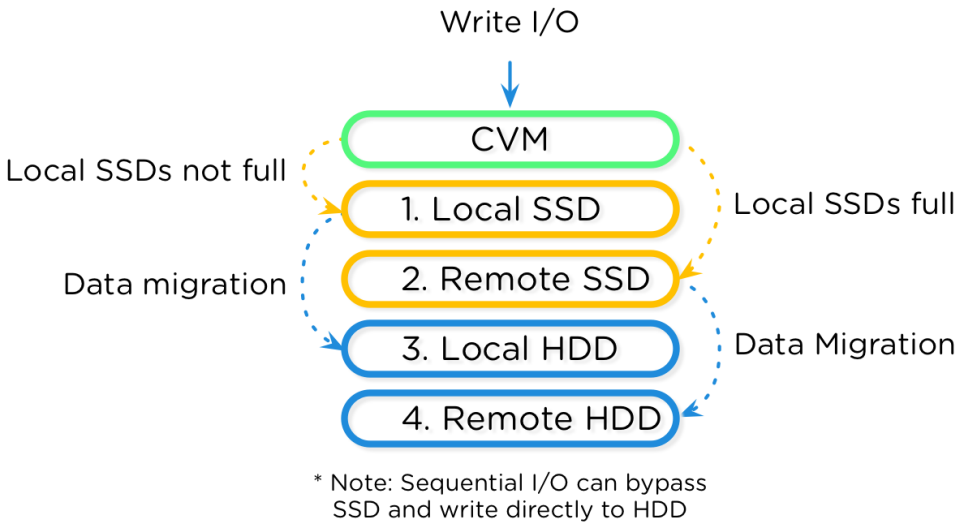

With the DSF, a CVM writes data to local flash memory for fast acknowledgment; the CVM also handles read operations locally for reduced latency and fast data delivery.

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 Nutanix 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. Having a local storage controller on each node ensures that storage performance as well as storage capacity increase linearly with each node addition.

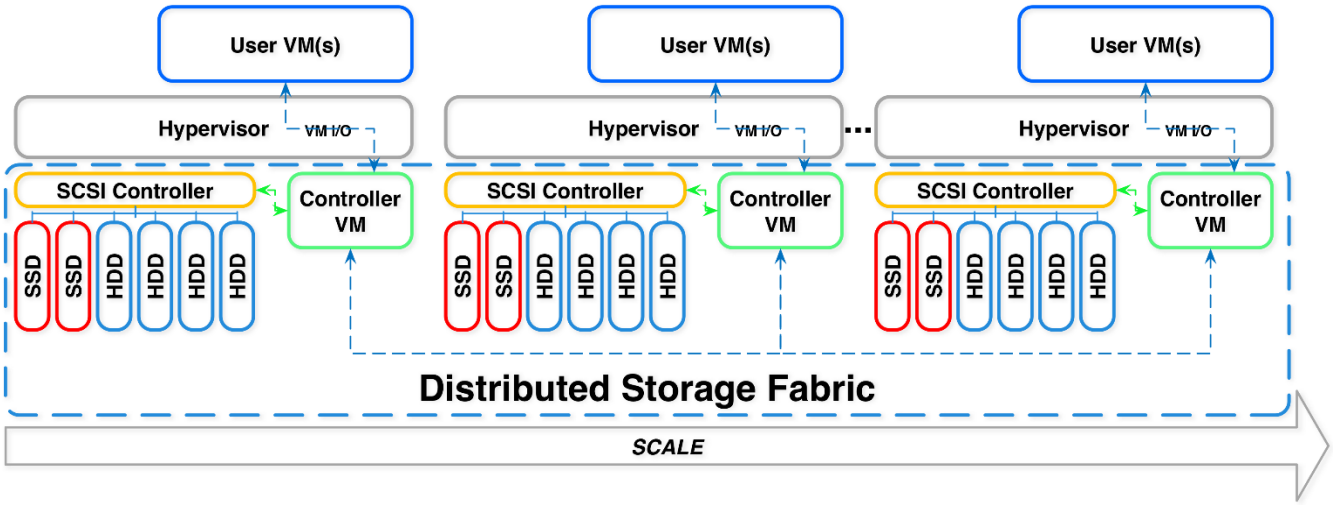

Local storage for each Nutanix node in the architecture appears to the hypervisor as one large pool of shared storage. This allows the DSF to support all key virtualization features. Data localization maintains performance and quality of service (QoS) on each host, minimizing the effect noisy VMs have 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 and high availability, the now local CVM serves a newly migrated VM's data. When reading old data (stored on the now remote CVM) the local CVM forwards the I/O request to the remote CVM. All write I/O occurs locally. The DSF detects that I/O is occurring from a different node and migrates the data to the local node in the background, allowing for all read I/O to now be served locally. The data only migrates when there have been enough reads and writes from the remote node to minimize network utilization.

The next figure shows how data follows the VM as it moves between hypervisor nodes.

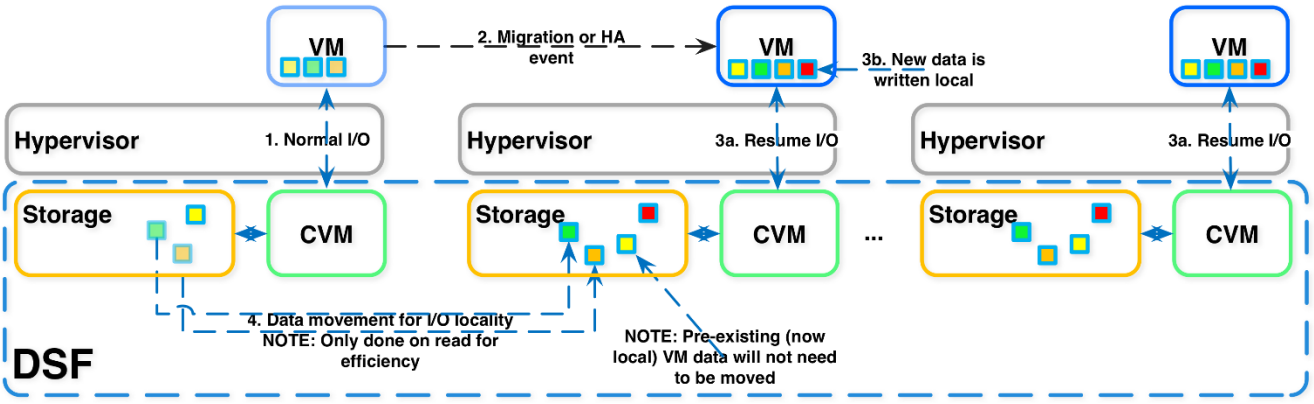

Nutanix Shadow Clones delivers distributed localized caching of virtual disks performance in multi-reader scenarios, such as desktop virtualization using VMware Horizon or RDSH. With Shadow Clones, the CVM actively monitors virtual disk access trends. If there are requests originating from more than two remote CVMs, as well as the local CVM, and all of the requests are read I/O and the virtual disk will be marked as immutable. Once the disk has been marked immutable, the virtual disk is then cached locally by each CVM, so read operations are now satisfied locally by local storage.

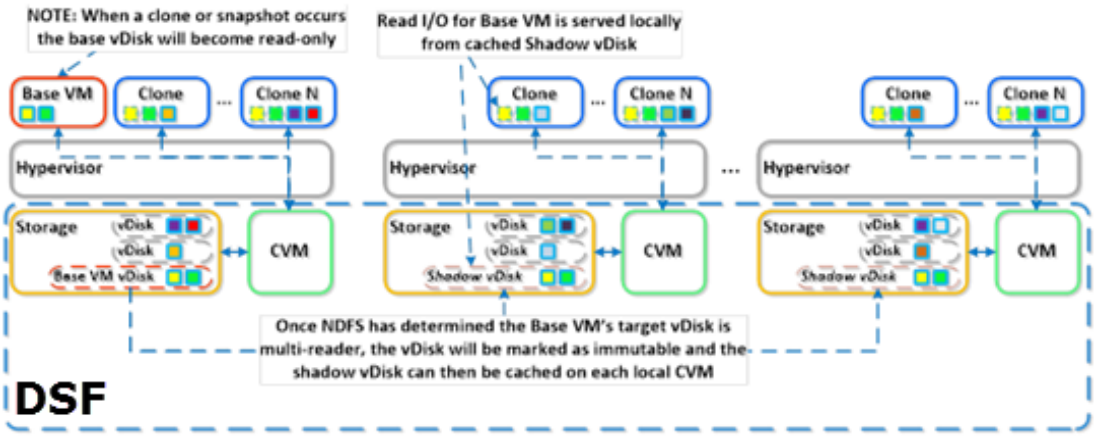

2.5 Nutanix Web-scale Converged Infrastructure

The Nutanix web-scale 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 contains zero reconfiguration of or customization to the Nutanix product to optimize for this use case.

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

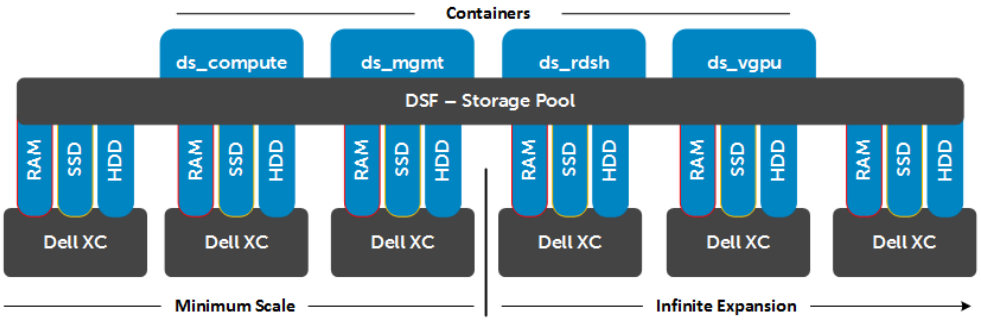

Dell XC Web Scale allows organizations to deliver virtualized or remote desktops and applications through a single platform and support end users with access to all of their desktops and applications in a single place.

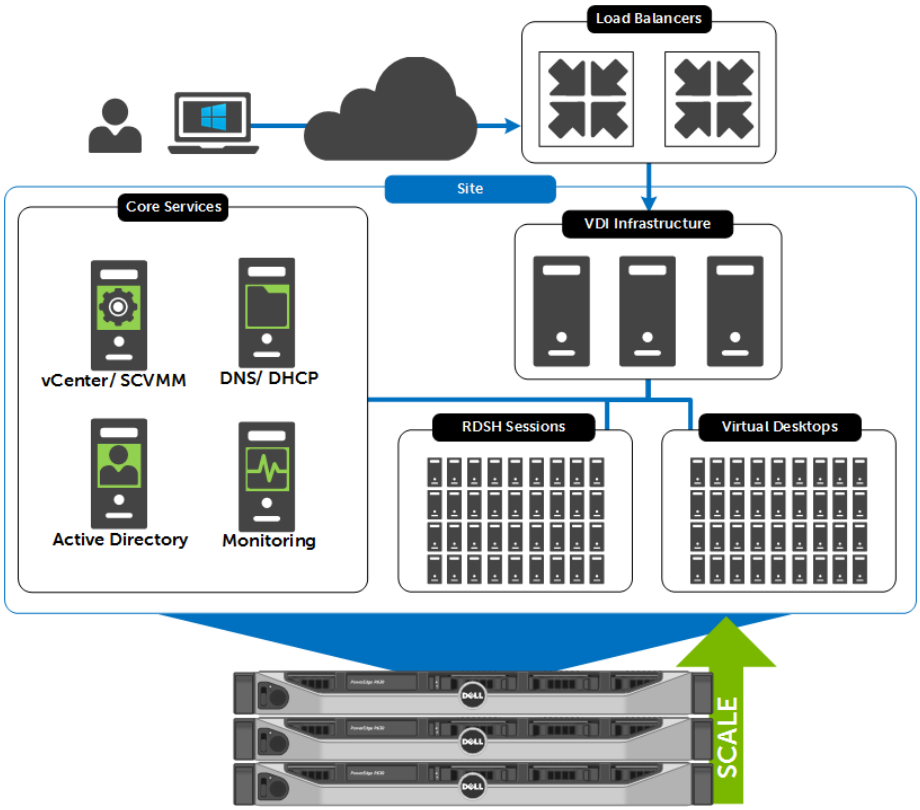

2.6 Dell XC Web Scale – Solution Pods

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

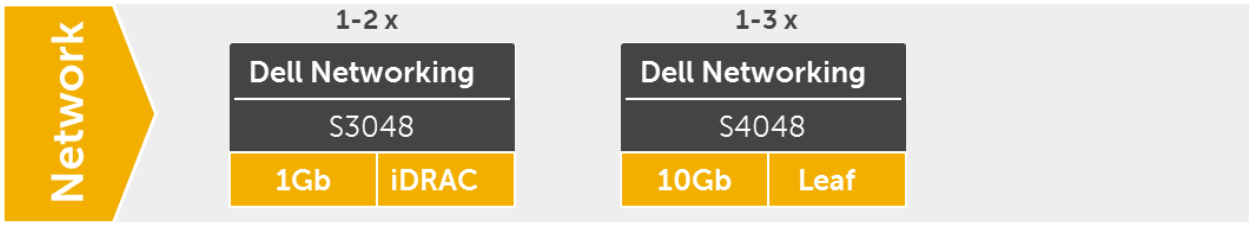

The compute, management and storage layers are converged into a single server XC Series appliance cluster, hosting VMware vSphere. The recommended boundaries of an individual pod are based on number of nodes supported within a given hypervisor cluster, 64 nodes for vSphere 6, although the Nutanix DFS cluster can scale much larger.

Dell recommends that the VDI management infrastructure nodes be separated from the compute resources onto their own appliance cluster with a common storage 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 NFS namespace can be shared across multiple hypervisor clusters adding disk capacity and performance for each distinct cluster.

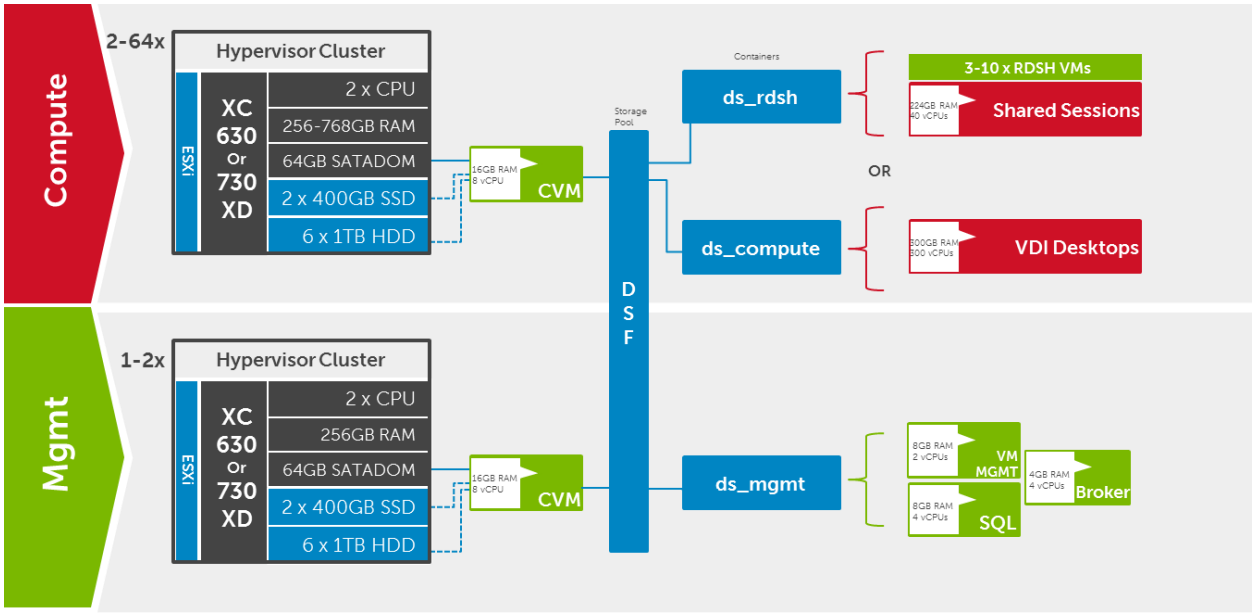

High-performance graphics capabilities compliment the solution and can easily be added at any time to any new or existing Dell XC vSphere deployment. Simply add the appropriate number of XC730 appliances to your DSF cluster and provide a superior user experience with vSphere 6 and NVIDIA GRID vGPU technology. Any XC appliance can be utilized for the compute or mgmt portions of this solution.

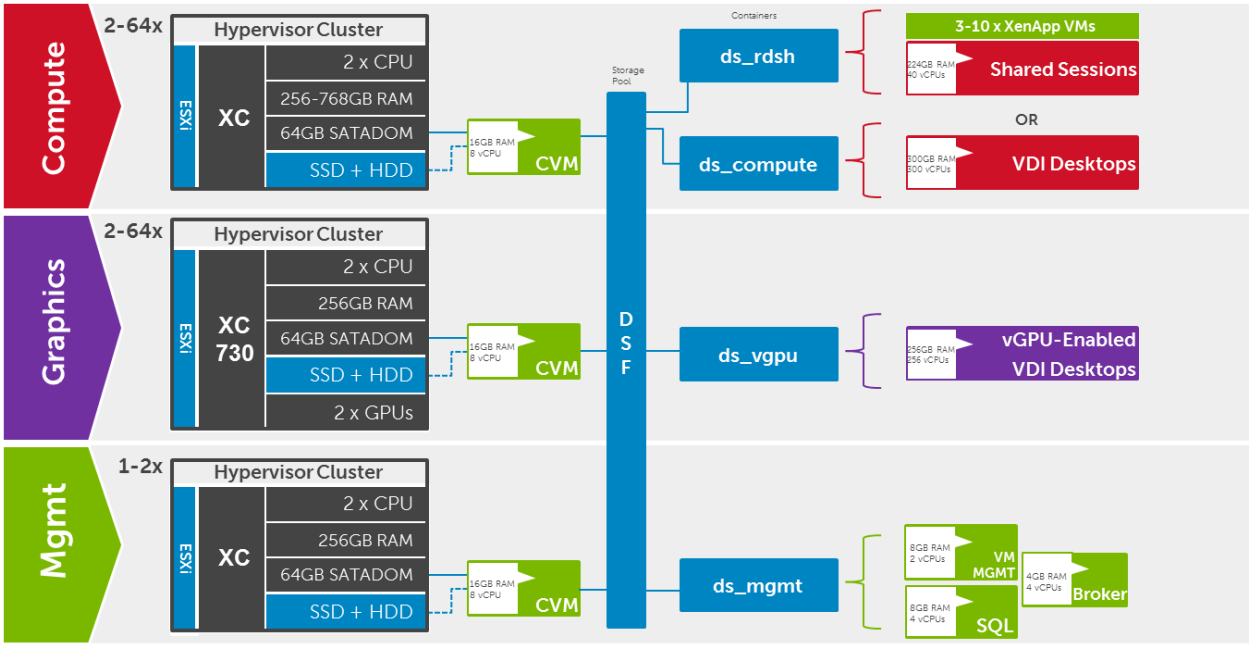

2.6.1 Network Architecture

Designed for true linear scaling, Dell XC series leverages a Leaf-Spine network architecture. A Leaf-Spine architecture consists of two network tiers: an L2 Leaf and an L3 Spine 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 nodes to thousands without any impact to available bandwidth:

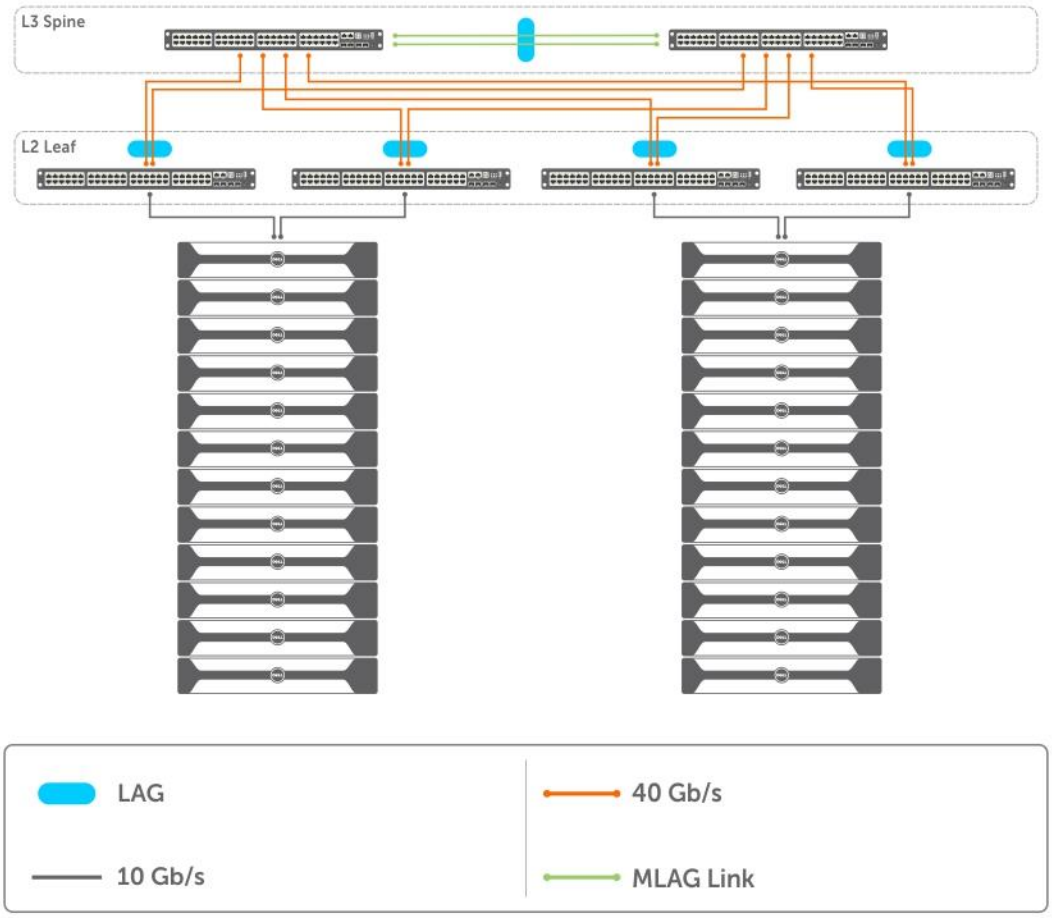
L3 Spine
L2 Leaf
LAG
40 Gb/s
10 Gb/s
MLAG Link

3 Hardware Components

3.1 Network

The following sections contain the core network components for the Dell Wyse Datacenter solutions. General uplink cabling guidance to consider in all cases is that TwinAx 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 10GbE 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
Dell Networking S3048-ON	48 x 1000BaseT 4 x 10Gb SFP+	Redundant hot-swap PSUs & fans	1Gb connectivity
	Non-blocking, line-rate performance	VRF-lite, Routed VLT, VLT Proxy Gateway	
	260Gbps full-duplex bandwidth	User port stacking (up to 6 switches)	
	131 Mpps forwarding rate	Open Networking Install Environment (ONIE)	

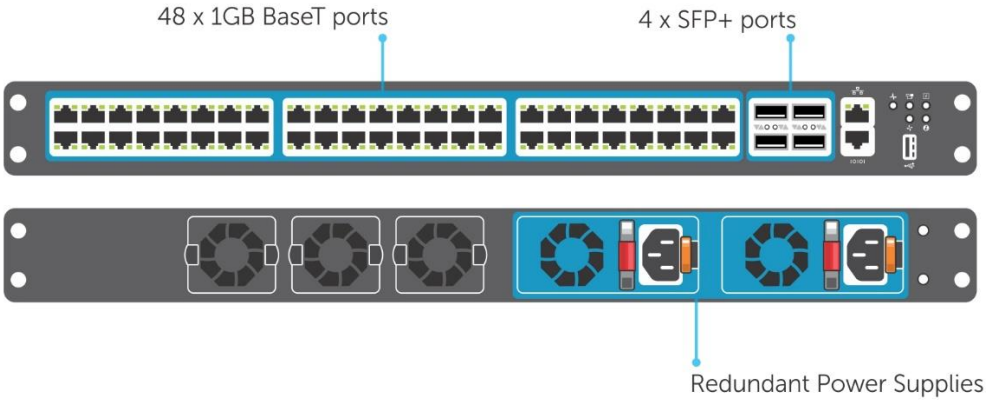

3.1.2 Dell Networking S4048 (10Gb ToR Switch)

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 72 x 10Gb SFP+ ports with breakout cables User port stacking (up to 6 switches) Open Networking Install Environment (ONIE)	10Gb connectivity

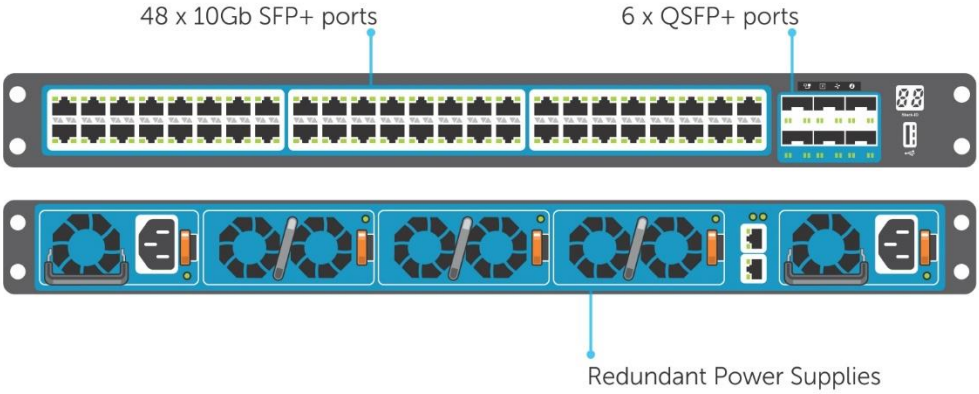

For more information on the S3048, S4048 switches and Dell Networking, please visit: LINK

3.2 Dell XC Converged Appliance

Consolidate compute and storage into a single chassis with Dell XC Series web-scale 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 recommends the XC630 or XC730XD. For workloads requiring graphics the XC730 with

NVIDIA GRID can be integrated into any environment running any other XC appliance. For small Remote Office – Branch Office scenarios we offer the XC430 and for high density requirements the 4-node in 2U XC6320. For more information on the Dell XC Series, please visit: Link

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

- **A5** configuration is perfect for small scale, POC or low density cost-conscience environments. Available in the XC630, XC730XD, XC430 and XC6320.
- **B5** configuration is geared toward larger scale general purpose workloads, balancing performance and cost-effectiveness. Available in the XC630, XC730XD, XC430 and XC6320.
- **B7** is the premium configuration offering an abundance of high performance and tiered capacity where user density is maximized. Available in the XC630, XC730XD, XC430 and XC6320.
- **B7-GFX** for high-performance graphical workloads is available in the XC730.

A5	B5	B7	B7-GFX
• **CPU:** 2 x 10-Core • **RAM:** 256GB • **Disk:** 2 x SSD 4TB+ HDD	• **CPU:** 2 x 14-Core • **RAM:** 384GB • **Disk:** 2 x SSD 4TB+ HDD	• **CPU:** 2 x 20-Core • **RAM:** 384GB • **Disk:** 2 x SSD 6TB+ HDD	• **CPU:** 2 x 18-Core • **RAM:** 256GB • **Disk:** 2 x SSD 4TB+ HDD • **GPU:** 2 x NVIDIA GRID

3.2.1 Dell XC630

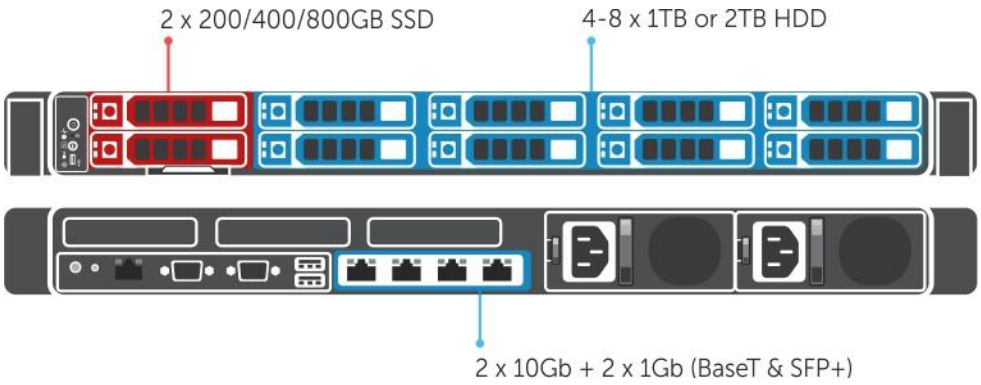

The Dell XC630 is a 1U platform with a broad range of configuration options. Each appliance comes equipped with dual CPUs, 10 to 20 cores, and up to 384GB of high-performance RAM by default. A minimum of six disks is required in each host, 2 x SSD for the hot tier (Tier1) and 4 x HDD for the cold tier (Tier2) which can be expanded up to 8 HDDs as required. The 64GB SATADOM boots the hypervisor and Nutanix Controller VM while the PERC H330 is configured in pass-through mode connecting to the SSDs and HDDs. 64GB is consumed on each of the first two SSDs for the Nutanix "home". These six 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.

XC630	A5	B5	B7
CPU	2 x E5-2630v4 (10C, 2.2GHz)	2 x E5-2660v4 (14C, 2.0GHz)	2 x E5-2698v4 (20C, 2.2GHz)
Memory	16 x 16GB 2400MT/s RDIMMs **Effective speed: 2400MT/s @ 256GB**	24 x 16GB 2400MT/s RDIMMs **Effective speed: 1866MT/s @ 384GB**	24 x 16GB 2400MT/s RDIMMs **Effective speed: 1866MT/s @ 384GB**
Storage Ctrls	PERC H330 – no RAID	PERC H330 – no RAID	PERC H330 – no RAID
Storage	64GB SATADOM **(CVM/ Hypervisor)** 16GB SD Module (Tools/ Recovery) 2 x 200GB SSD 2.5" (T1) 4 x 1TB HDD 2.5" (T2)	64GB SATADOM **(CVM/ Hypervisor)** 16GB SD Module (Tools/ Recovery) 2 x 400GB SSD 2.5" (T1) 4 x 1TB HDD 2.5" (T2)	64GB SATADOM **(CVM/ Hypervisor)** 16GB SD Module (Tools/ Recovery) 2 x 400GB SSD 2.5" (T1) 6 x 1TB HDD 2.5" (T2)
Network	2 x 10Gb, 2 x 1Gb SFP+/ BT	2 x 10Gb, 2 x 1Gb SFP+/ BT	2 x 10Gb, 2 x 1Gb SFP+/ BT
iDRAC	iDRAC8 Ent w/ vFlash, 8GB SD	iDRAC8 Ent w/ vFlash, 8GB SD	iDRAC8 Ent w/ vFlash, 8GB SD
Power	2 x 750W PSUs	2 x 750W PSUs	2 x 750W PSUs

3.2.2 Dell XC730XD (High capacity)

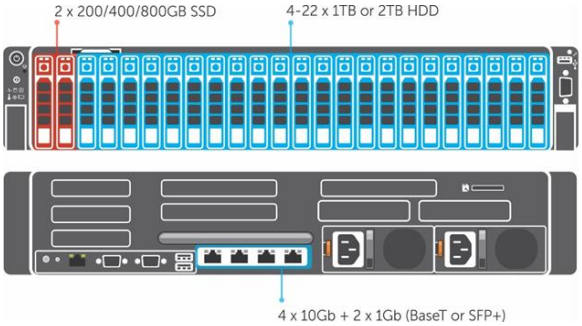

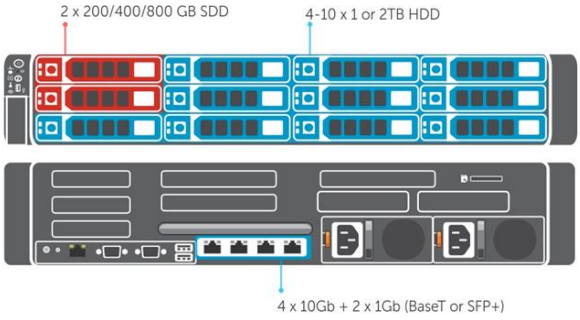

The Dell XC730XD is a 2U platform that can be configured with 24 x 2.5" disks or 12 x 3.5" disks to serve a broad range of capacity requirements. Each appliance comes equipped with dual CPUs, 10 to 20 cores, and up to 384GB of high-performance RAM by default. A minimum of six disks is required in each host, 2 x SSD for the hot tier (Tier1) and 4 x HDD for the cold tier (Tier2) which can be expanded as required up to a possible 45TB per node raw. The 64GB SATADOM boots the hypervisor and Nutanix Controller VM while the PERC H330 is configured in pass-through mode connecting to the SSDs and HDDs. 64GB is consumed on each of the first two SSDs for the Nutanix "home". These six 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.

XC730xd	A5	B5	B7
CPU	2 x E5-2630v4 (10C, 2.2GHz)	2 x E5-2660v4 (14C, 2.0GHz)	2 x E5-2698v4 (20C, 2.2GHz)
Memory	16 x 16GB 2400MT/s RDIMMs **Effective speed: 2400MT/s @ 256GB**	24 x 16GB 2400MT/s RDIMMs **Effective speed: 1866MT/s @ 384GB**	24 x 16GB 2400MT/s RDIMMs **Effective speed: 1866MT/s @ 384GB**
Storage Ctrls	PERC H330 – no RAID	PERC H330 – no RAID	PERC H330 – no RAID
Storage*	64GB SATADOM **(CVM/ Hypervisor)** 16GB SD Module (Tools/ Recovery) 2 x 200GB SSD 2.5"/ 3.5" (T1) 4 x 1TB HDD 2.5"/ 3.5" (T2)	64GB SATADOM **(CVM/ Hypervisor)** 16GB SD Module (Tools/ Recovery) 2 x 400GB SSD 2.5"/ 3.5" (T1) 4 x 1TB HDD 2.5"/ 3.5" (T2)	64GB SATADOM **(CVM/ Hypervisor)** 16GB SD Module (Tools/ Recovery) 2 x 400GB SSD 2.5"/ 3.5" (T1) 6 x 1TB HDD 2.5"/ 3.5" (T2)
Network	2 x 10Gb, 2 x 1Gb SFP+/ BT	2 x 10Gb, 2 x 1Gb SFP+/ BT	2 x 10Gb, 2 x 1Gb SFP+/ BT
iDRAC	iDRAC8 Ent w/ vFlash, 8GB SD	iDRAC8 Ent w/ vFlash, 8GB SD	iDRAC8 Ent w/ vFlash, 8GB SD
Power	2 x 750W PSUs	2 x 750W PSUs	2 x 750W PSUs

*Available in 24 x 2.5" or 12 x 3.5" configurations

3.2.3 Dell XC730 (Graphics)

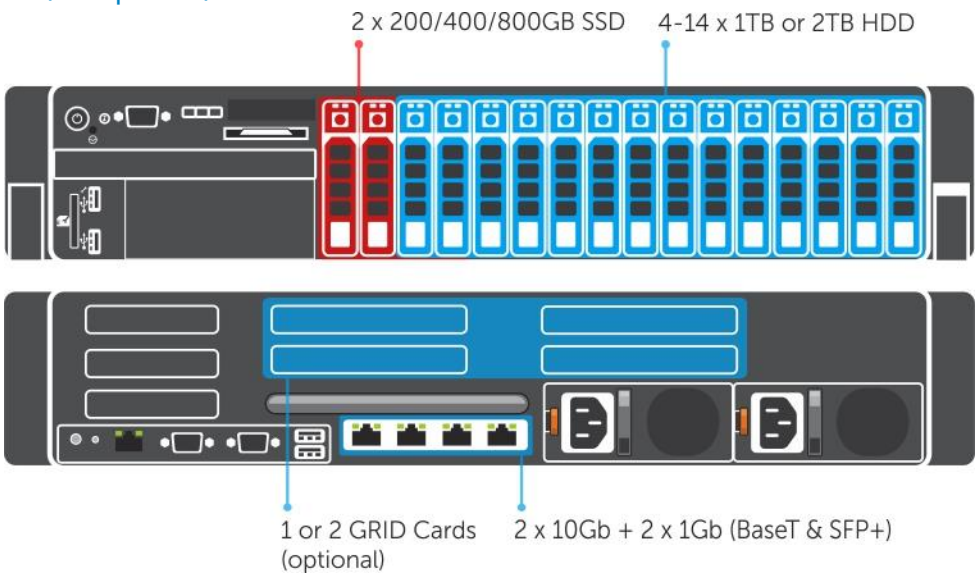

The Dell XC730 is a 2U platform that can be configured with dual NVIDIA GRID cards using vGPU to supply high-performance virtualized graphics. Each appliance comes equipped with dual 18core CPUs and 256GB of high-performance RAM by default supporting up to 64 users per node. A minimum of six disks is required in each host, 2 x SSD for the hot tier (Tier1) and 4 x HDD for the cold tier (Tier2) which can be expanded as required. The 64GB SATADOM boots the hypervisor and Nutanix Controller VM while the PERC H330 is configured in pass-through mode connecting to the SSDs and HDDs. 64GB is consumed on each of the first two SSDs for the Nutanix "home". These six 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. Solutions can be designed around the XC730 entirely which can be purchased

with or without GRID cards. Additionally, the XC730 can be used to augment other non-graphics enabled deployments based on a differing XC platform such as with the XC630 or XC730XD.

XC730	B7-GFX
CPU	2 x E5-2695v4 (18C, 2.1GHz)
Memory	16 x 16GB 2400MT/s RDIMMs **Effective speed: 2400MT/s @ 256GB**
GPU Boards	2 x NVIDIA GRID K2 or M60 vGPU Profile: K2xxQ or M60-xQ
Storage Ctrls	PERC H330 – no RAID
Storage	64GB SATADOM **(CVM/ Hypervisor)** 16GB SD Module (Tools/ Recovery) 2 x 400GB SSD 2.5" (T1) 4 x 1TB HDD 2.5" (T2)
Network	2 x 10Gb, 2 x 1Gb SFP+/ BT
iDRAC	iDRAC8 Ent w/ vFlash, 8GB SD
Power	2 x 1100W PSUs

3.2.4 Dell XC430 (ROBO)

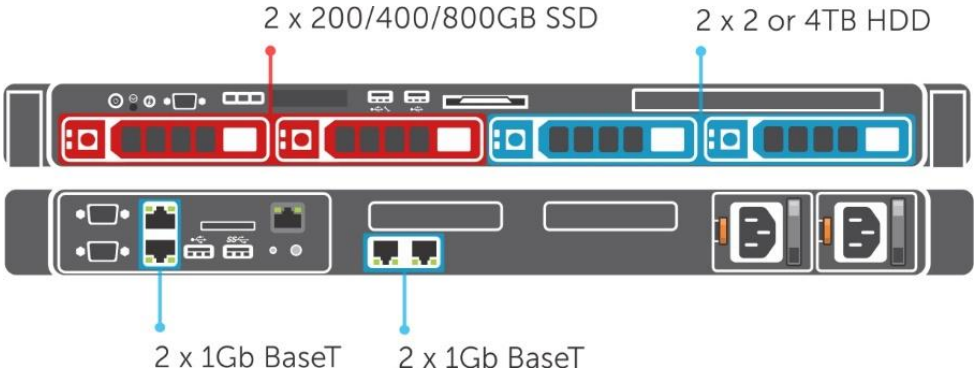

The Dell XC430 is a 1U platform that offers short depth (24") space savings perfect for the Remote Office/ Branch Office use case. Each appliance comes equipped with single or dual CPUs, 10 to 14 cores, and up to 384GB of high-performance RAM by default. A minimum of four disks is required in each host, 2 x SSD for the hot tier (Tier1) and 2 x HDD for the cold tier (Tier2) which can be configured to provide 4 or 8TB raw per node. The 64GB SATADOM boots the hypervisor and Nutanix Controller VM while the PERC H330 is configured in pass-through mode connecting to the SSDs and HDDs. 64GB is consumed on each of the first two SSDs for the Nutanix "home". These six 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.

XC430	A5	B5
CPU	2 x E5-2630v4 (10C, 2.2GHz)*	2 x E5-2660v4 (14C, 2.0GHz)*
Memory	8 x 32GB 2400MT/s RDIMMs **Effective speed: 2400MT/s @ 256GB**	12 x 32GB 2400MT/s RDIMMs **Effective speed: 2400MT/s @ 384GB**
Storage Ctrls	PERC H330 – no RAID	PERC H330 – no RAID
Storage	64GB SATADOM **(CVM/ Hypervisor)** 16GB SD Module (Tools/ Recovery) 2 x 200GB SSD 3.5" (T1) 2 x 2TB HDD 3.5" (T2)	64GB SATADOM **(CVM/ Hypervisor)** 16GB SD Module (Tools/ Recovery) 2 x 400GB SSD 3.5" (T1) 2 x 2TB HDD 3.5" (T2)
Network	4 x 1Gb BaseT *(10Gb optional)*	4 x 1Gb BaseT *(10Gb optional)*
iDRAC	iDRAC8 Enterprise	iDRAC8 Enterprise
Power	2 x 550W PSUs	2 x 550W PSUs

*Can be sold with single CPU, 8 DIMMs max

3.2.5 Dell XC6320 (High Density)

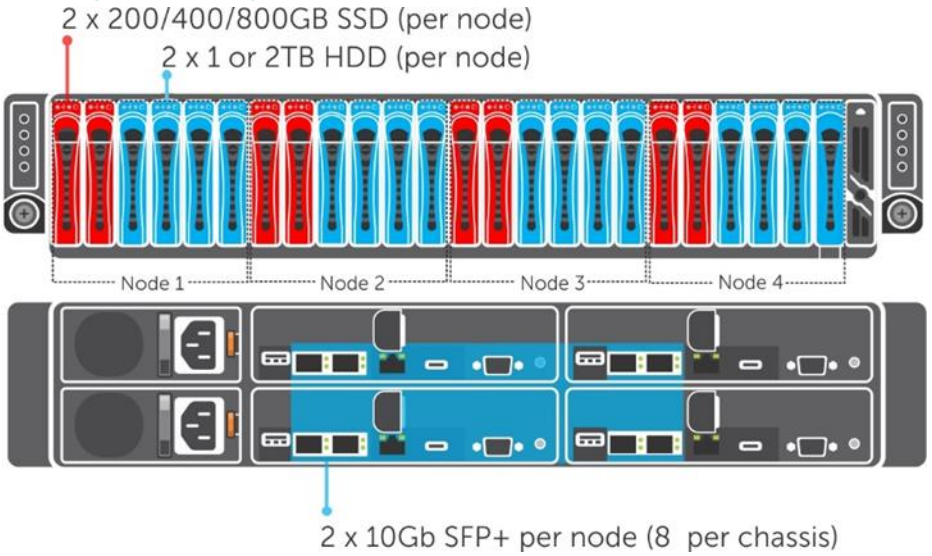

The Dell XC6320 is a 4-node in 2U platform offering maximum user density per rack unit. Each of the fours nodes within a single 2U appliance comes equipped with dual CPUs, 10 to 14 cores, and up to 512GB of high-performance RAM by default. Each node is equipped with six disks, 2 x SSD for the hot tier (Tier1) and 4 x HDD for the cold tier (Tier2). The 64GB SATADOM boots the hypervisor and Nutanix Controller VM while the LSI2008 HBA connects the SSDs and HDDs. 64GB is consumed on each of the first two SSDs for

the Nutanix "home". These six disks are presented to the Nutanix CVM running locally on each host which contributes to the clustered DSF pool. Each platform is outfitted with SFP+ NICs.

XC6320	XC6300*	A5	B5
CPU		2 x E5-2630v4 (10C, 2.2GHz)	2 x E5-2660v4 (14C, 2.0GHz)
Memory		16 x 16GB 2400MT/s RDIMMs **Effective speed: 2400MT/s @ 256GB**	16 x 32GB 2400MT/s RDIMMs **Effective speed: 1866MT/s @ 512GB**
Storage Ctrls		LSI2008– no RAID	LSI2008– no RAID
Storage	**24 x 2.0" disks in 2 tiers (SSD + HDD)**	64GB SATADOM **(CVM/ Hypervisor)** 16GB SD Module (Tools/ Recovery) 2 x 200GB SSD 2.5" (T1) 4 x 1TB HDD 2.5" (T2)	64GB SATADOM **(CVM/ Hypervisor)** 16GB SD Module (Tools/ Recovery) 2 x 400GB SSD 2.5" (T1) 4 x 1TB HDD 2.5" (T2)
Network		2 x 10Gb SFP+	2 x 10Gb SFP+
iDRAC		iDRAC8 Enterprise	iDRAC8 Enterprise
Power	2 x 1600W PSUs		

*4 nodes required per chassis

3.3 GPUs

3.3.1 NVIDIA GRID K1 and K2

NVIDIA GRID™ technology offers the ability to offload graphics processing from the CPU to the GPU in virtualized environments, allowing the data center manager to deliver true PC graphics-rich experiences to more users for the first time. NVIDIA's Kepler™-based GRID K1 and K2 boards are specifically designed to enable rich graphics in virtualized environments.

GPU Virtualization

GRID boards allow hardware virtualization of the GPU. This means multiple users can share a single GPU, improving user density while providing true PC performance and compatibility.

Low-Latency Remote Display

NVIDIA's patented low-latency remote display technology greatly improves the user experience by reducing the lag that users feel when interacting with their virtual machine. With this technology, the virtual desktop screen is pushed directly to the remoting protocol.

Maximum User Density

NVIDIA GRID boards have an optimized multi-GPU design that helps to maximize user density. GRID K1 boards, which include four Kepler-based GPUs and 16GB of memory, are designed to host the maximum number of concurrent users. GRID K2 boards, which include two higher end Kepler GPUs and 8GB of memory, deliver maximum density for users of graphics-intensive applications.

Specs	Grid K1	Grid K2
Number of GPUs	4 x Kepler GPUs (GK107)	2 x high-end Kepler GPUs (GK104)
Total CUDA cores	768 (192 per GPU)	3072 (1536 per GPU)
Core Clock	850 MHz	745 MHz
Total memory size	16 GB DDR3	8 GB GDDR5
Max power	130 W	225 W
Form Factors	Dual slot (4.4" x 10.5")	Dual slot (4.4" x 10.5")
Aux power	6-pin connector	8-pin connector
PCIe	x16 (Gen3)	x16 (Gen3)
Cooling solution	Passive	Passive/ Active

3.3.2 NVIDIA Tesla M60

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 to any user, anywhere.

Specs	Tesla M60
Number of GPUs	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

For more information on NVIDIA Grid, please visit: Link

3.4 Dell Wyse Endpoints

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

3.4.1 Wyse 3030 LT Thin Client with ThinOS

The Wyse 3030 LT thin client offers an excellent user experience within a cost-effective offering, and features the virus resistant and extremely efficient Wyse ThinOS, for environments in which security is critical—there's no attack surface to put your data at risk. The 3030 LT delivers outstanding performance based on its dual core processor design, and delivers smooth multimedia, bi-directional audio and flash playback. Boot up in just seconds and log in securely to almost any network. In addition, the Wyse 3030 LT is designed for smooth playback of high bit-rate HD video and graphics within a very compact form factor, with very efficient energy consumption and low heat emissions. Using less than 7 watts of electricity, the Wyse 3030 LT's small size enables discrete mounting options: under desks, to walls, and behind monitors, creating cool workspaces in every respect. For more information, please visit: Link

3.4.2 Wyse 5010 PCoIP Thin Client with ThinOS

Designed for knowledge workers and power users, the Wyse 5010 thin client with ThinOS and PCoIP is a high-performance thin client based on Dell Wyse ThinOS, the virus-resistant firmware base designed for optimal thin client security, performance, and ease-of-use. Highly secure, compact and powerful, the Wyse 5010 combines Dell Wyse ThinOS with a dual-core AMD processor and a revolutionary unified graphics engine for an outstanding user experience. The Wyse 5010 (ThinOS) addresses the performance challenges of processing-intensive applications like computer-aided design, multimedia, HD video and 3D modeling. Scalable enterprise-wide on-premise or cloud-based management provides simple deployment, patching and updates. Take a unit from box to productivity in minutes with auto configuration. Delivering outstanding processing speed and power, security and display performance, the Wyse 5010 (ThinOS) offers a unique combination of performance, efficiency, and affordability. The Wyse 5010 (ThinOS) is Citrix HDX, Microsoft® RemoteFX, and VMware® Horizon View certified. It also supports legacy peripherals via an optional USB adapter. For more information, please visit: Link

3.4.3 Wyse 5010 Thin Client with WES8

In addition to Wyse ThinOS, the Dell Wyse 5010 thin client is available with Windows Embedded Standard 8 and packs dual-core processing power into a compact form factor for knowledge workers who need performance for demanding virtual Windows® desktops and cloud applications. It's also great for kiosks, and multi-touch displays in a wide variety of environments, including manufacturing, hospitality, retail, and healthcare. It features dual-core processing power and an integrated graphics engine for a fulfilling Windows® 8 user experience. Knowledge workers will enjoy rich content creation and consumption as

well as everyday multimedia. Kiosk displays will look great on a thin client that is Microsoft RemoteFX®, Citrix® HDX, VMware PCoIP, and HD video-enabled. Operating with less than 9 watts of energy, the Dell Wyse 5010 (Windows) offers cool, quiet operations, potentially lowering your overall carbon footprint. For more information, please visit: Link

3.4.4 Wyse 5030 PCoIP Zero Client

Uncompromising computing with the benefits of secure, centralized management. The Wyse 5030 PCoIP zero client for VMware Horizon and Amazon WorkSpaces is a secure, easily managed zero client that provides outstanding graphics performance for advanced applications such as CAD, 3D solids modeling, video editing and advanced worker-level office productivity applications. Smaller than a typical notebook, this dedicated zero client is designed specifically for VMware Horizon and Amazon WorkSpaces. It features the latest processor technology from Teradici to process the PCoIP protocol in silicon and includes client-side content caching to deliver the highest level of performance available over 2 HD displays in an extremely compact, energy-efficient form factor. The Wyse 5030 delivers a rich user experience while resolving the challenges of provisioning, managing, maintaining and securing enterprise desktops. For more information, please visit: Link

3.4.5 Wyse 5040 AIO Thin Client

The Wyse 5040 AIO all-in-one (AIO) 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 by Wyse Device Manager (WDM), or Wyse Cloud Client Manager can help lower your total cost of ownership as you grow from just a few thin clients to tens of thousands.

3.4.6 Wyse 5050 AIO PCoIP Zero Client

The Wyse 5050 All-in-One (AIO) PCoIP zero client combines the security and performance of the Wyse 5030 PCoIP zero client for VMware with the elegant design of Dell's best-selling P24 LED monitor. The Wyse 5050 AIO provides a best-in-class virtual experience with superior manageability – at a better value than purchasing a zero client and high resolution monitor separately. A dedicated hardware PCoIP engine delivers the highest level of display performance available for advanced applications, including CAD, 3D solids modeling, video editing and more. Elegant in appearance and energy efficient, the Wyse 5050 AIO is a fully functional VMware Horizon endpoint that delivers a true PC-like experience. It offers the full benefits of an efficient and secure centralized computing environment, like rich multimedia, high-resolution 3D graphics, HD media, and full USB peripheral interoperability locally (LAN) or remotely (WAN). For more information, please visit: Link

3.4.7 Wyse 7010 thin client with Linux OS

Designed for power users, the Wyse 7010 is the highest performing thin client on the market. Highly secure and ultra-powerful, it combines Dell Wyse-enhanced SUSE Linux Enterprise with a dual-core AMD 1.65 GHz CPU and a revolutionary unified engine for an unprecedented user experience. It eliminates performance constraints for high-end, processing-intensive applications like computer-aided design, multimedia, HD video and 3D modeling. Scalable enterprise-wide management provides simple deployment, patching and updates. Take a unit from box to productivity in minutes with auto configuration. Delivering unmatched processing speed and power, security and display performance, it's no wonder no other thin client can compare. For more information, please visit: Link

3.4.8 Wyse 7020 Thin Client with WES7

The Wyse 7020 is a super high-performance Windows Embedded Standard 7 thin client for virtual desktop environments. Featuring a quad-core AMD CPU, and an integrated graphics engine that significantly boost performance; it achieves exceptional speed and power for the most demanding VDI and embedded Windows applications, rotational 3D graphics, 3D simulation and modeling, unified communications, and multi-screen HD multimedia. Take a unit from box to productivity in minutes. Just select the desired configuration and it does the rest automatically—no need to reboot. Scale to tens of thousands of endpoints with Wyse Device Manager or leverage your existing Microsoft System Center Configuration Manager platform. The Wyse 7020 is the thin client for power users who need workstation-class performance on their desktop or within a desktop virtualization environment (x86 or x64). For more information, please visit: Link

3.4.9 Wyse 7020 Thin Client with WES8

Wyse 7020 is a super high-performance Windows Embedded 8 Standard thin client for virtual desktop environments. Featuring a quad-core AMD CPU, it offers a vibrant Windows 8 experience and achieves exceptional speed and power for the most demanding embedded Windows applications, rich 3D graphics, and HD multimedia. You can scale to tens of thousands of Wyse 7020 endpoints with Wyse Device Manager (WDM) software or leverage your existing Microsoft System Center Configuration Manager platform. With single-touch or multi-touch capable displays, it adds the ease of an intuitive touch user experience. The Wyse 7020 is an ideal thin client for offering a high-performance Windows 8 experience with the most demanding mix of virtual desktop or cloud applications (x86 or x64). For more information please visit: Link

3.4.10 Wyse 7030 PCoIP Zero Client

Uncompromising computing with the benefits of secure, centralized management. The Wyse 7030 PCoIP zero client for VMware Horizon and Amazon WorkSpaces is a secure, easily managed zero client that provides outstanding graphics performance for advanced applications such as CAD, 3D solids modeling, video editing and advanced worker-level office productivity applications. About the size of a notebook, this dedicated zero client designed specifically for VMware Horizon and Amazon WorkSpaces. It features the latest processor technology from Teradici to process the PCoIP protocol in silicon and includes client-side content caching to deliver the highest level of display performance available over 4 HD displays in a compact, energy-efficient form factor. The Dell Wyse 7030 delivers a rich user experience while resolving the challenges of provisioning, managing, maintaining and securing enterprise desktops. For more information, please visit: Link

3.4.11 Wyse 7040 Thin Client with WES7P

The Wyse 7040 is a high-powered, ultra-secure thin client. Equipped with 6th generation Intel Core 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 Windows Embedded Standard 7P 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

3.4.12 Dell Venue 11 Pro 5000

Meet the ultimate in productivity, connectivity and collaboration. Enjoy full laptop performance in an ultra-portable tablet that has unmatched flexibility for a business in motion. This dual purpose device works as a tablet when you're out in the field but also enables you to work on your desktop in the office thanks to an optional dock. For more information, please visit: Link

3.4.13 Dell Chromebook 13

The lightweight, easy-to-use Dell Chromebook 13 helps turn education into exploration - without the worries of safety or security. Priced to make 1:1 computing affordable today, Chromebook 13 is backed by Dell support services to make the most of your budget for years to come. The Chrome OS and Chrome browser get students online in an instant and loads web pages in seconds. A high-density battery supported by a 5th Gen Intel® CPU provides up to 12 hours of power. Encourage creativity with the Chromebook 13 and its multimedia features that include a 13.3" screen, stereo sound and webcam. For more information, please visit: Link

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:

- **View 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.
- **View 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.
- **View 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.
- **View 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.
- **View 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.
- **View Composer –** This software service can be installed standalone or on the vCenter server and provides enablement to deploy and create linked clone desktop pools (also called non-persistent desktops).
- **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 2008 host (can be a VM).
- **View 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

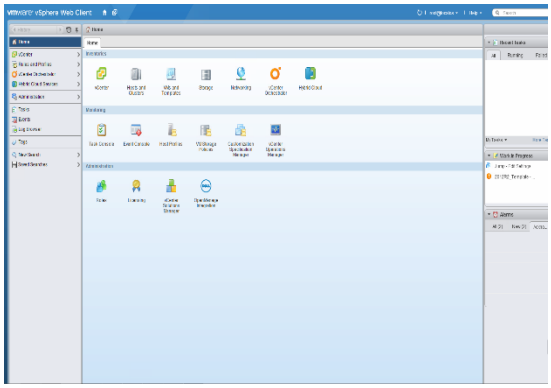

The vSphere hypervisor also known as ESXi is a bare-metal hypervisor that installs directly on top of your physical server and partitions it into multiple virtual machines. Each virtual machine shares the same physical resources as the other virtual machines and they can all run at the same time. Unlike 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 the 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.

4.2.1.1 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 XC appliance is as follows:

4.2.1.2 A5 NUMA Alignment

10 physical cores per CPU in the A5 configuration, 20 logical with Hyperthreading active, gives a total of 40 consumable cores per appliance. The Nutanix CVM will receive its vCPU allotment from the first physical CPU and so configuring the RDSH VMs as shown below will ensure that no NUMA spanning occurs which could lower performance. Per the example below, we have 3 RDSH VMs configured along with the Nutanix CVM all receiving 8 vCPUs, this results in a total oversubscription rate of 1.25x per host.

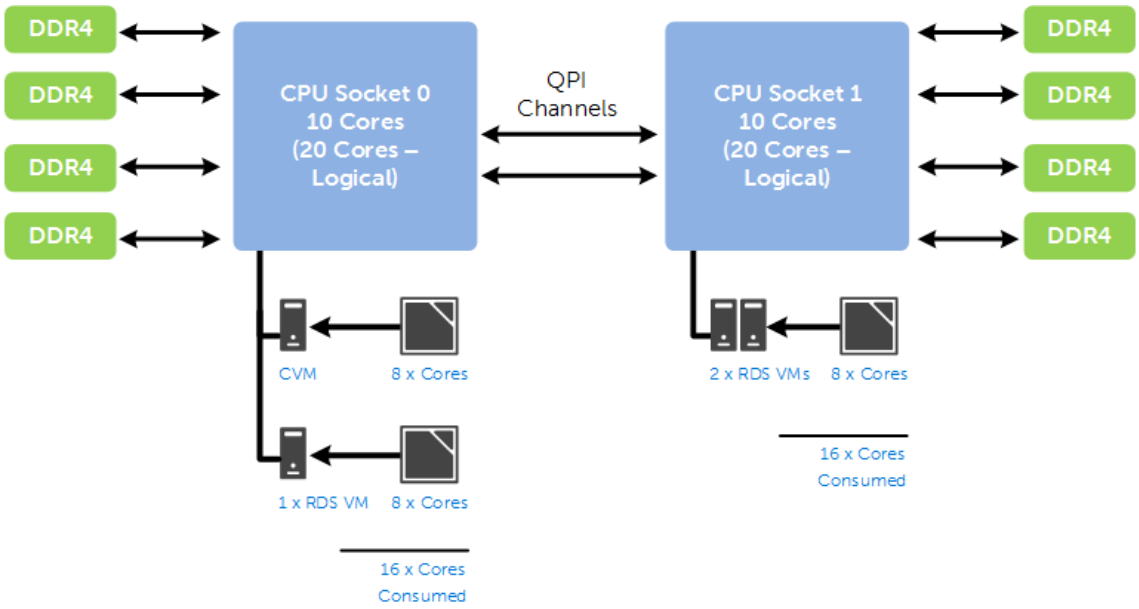

4.2.1.3 B5 NUMA Alignment

14 physical cores per CPU in the B5 configuration, 28 logical with Hyper-threading active, gives a total of 56 consumable cores per appliance. The Nutanix CVM will receive its vCPU allotment from the first physical CPU and so configuring the RDSH VMs as shown below will ensure that no NUMA spanning occurs which could lower performance. Per the example below, we have 5 RDSH VMs configured along with the Nutanix CVM all receiving 8 vCPUs, this results in a total oversubscription rate of 1.17x per host.

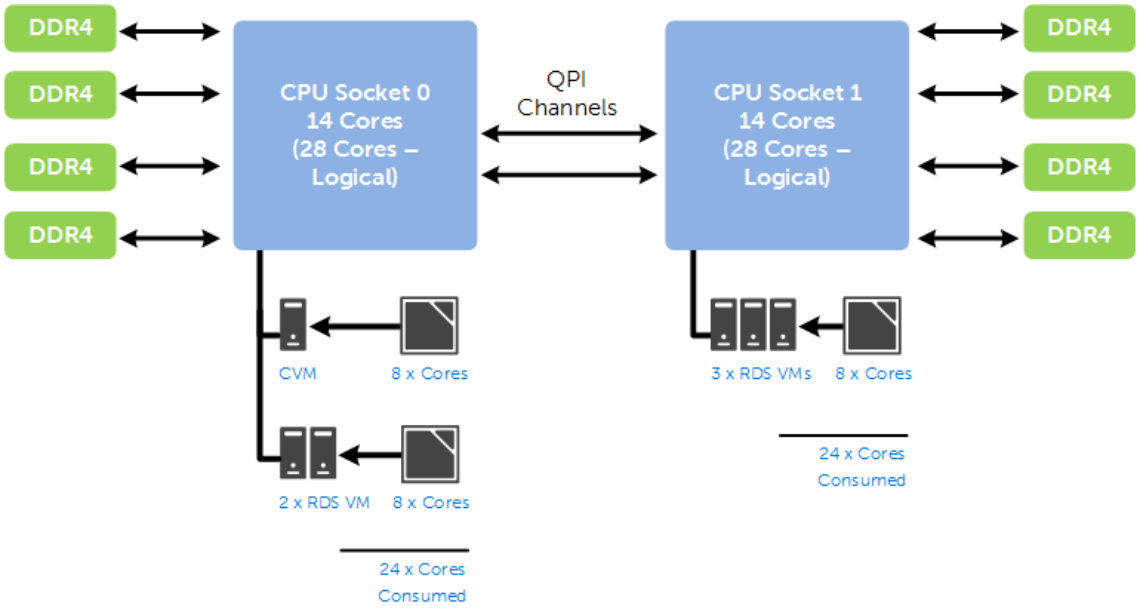

4.2.1.4 B7 NUMA Alignment

20 physical cores per CPU in the B7 configuration, 40 logical with Hyper-threading active, gives us a total of 80 consumable cores per appliance. The Nutanix CVM will receive its vCPU allotment from the first physical CPU and so configuring the RDSH VMs as shown below will ensure that no NUMA spanning occurs which could lower performance. Per the example below, we have 9 RDSH VMs configured along with the Nutanix CVM all receiving 8 vCPUs, this results in a total oversubscription rate of 2x per host.

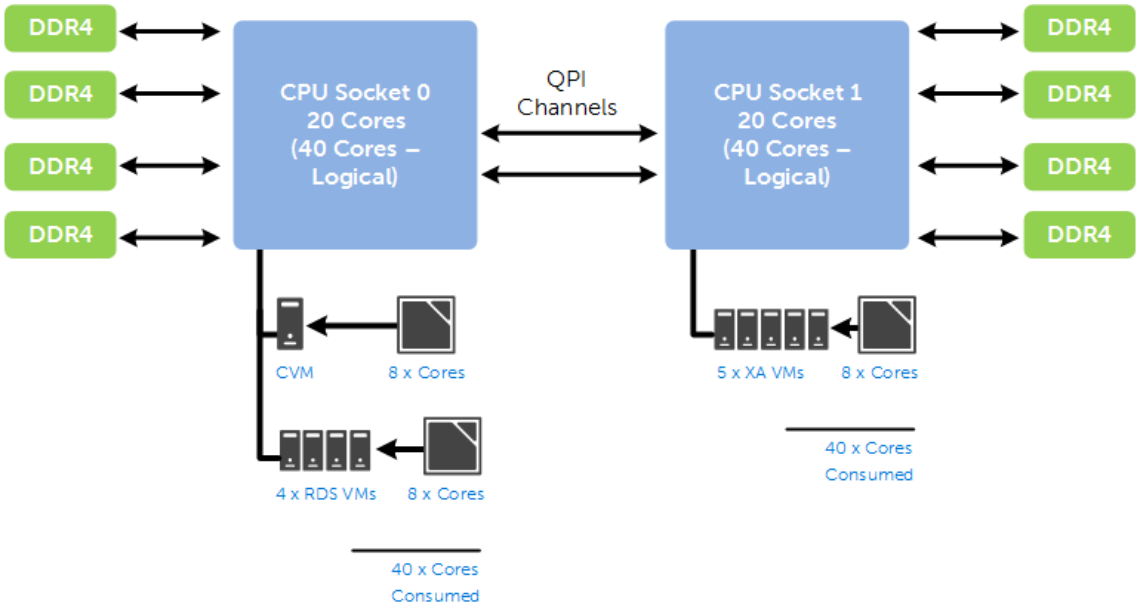

4.3 NVIDIA GRID vGPU

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 the ultimate in shared virtualized graphics performance.

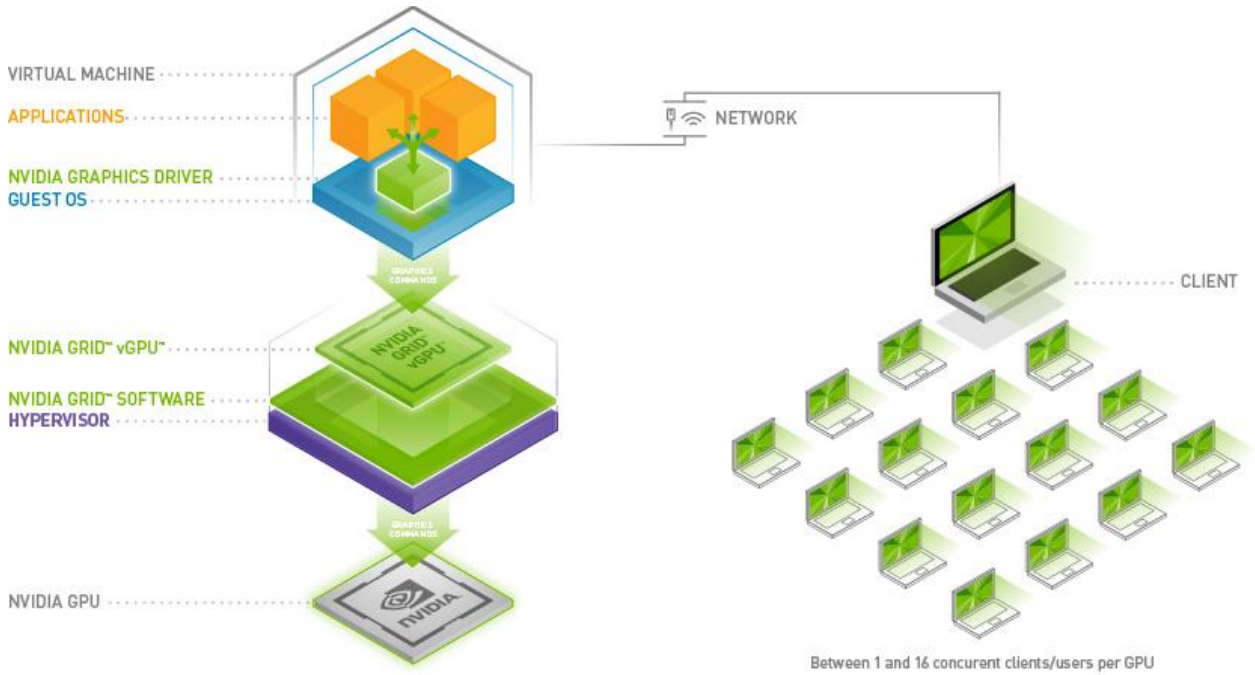

Image provided curtesy of NVIDIA Corporation, Copyright NVIDIA Corporation

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. All this is done without sacrificing server performance and so is truly cutting edge.

The Dell XC graphics platform is the award winning XC730 which can accommodate two NVIDIA K1/K2 or M60 graphics cards. The combination of Dell servers, NVIDIA vGPU technology and NVIDIA GRID cards enable high-end graphics users to experience high fidelity graphics quality and performance, for their favorite applications at a reasonable cost.

NOTE: GRID vGPU™ is supported on NVIDIA K2 cards and is available as a licensed feature on NVIDIA Tesla M60 cards.

GRID K2 vGPU Profiles:

Card	vGPU Profile	Graphics Memory (Frame Buffer)	Virtual Display Heads	Maximum Resolution	Physical GPUs	Maximum vGPUs		Intended Use Case
						Per GPU	Per Card	
GRID K2	K280Q	4GB	4	2560x1600	2	1	2	Designer
	K260Q	2GB	4	2560x1600		2	4	Designer
	K240Q	1GB	2	2560x1600		4	8	Power User
	K220Q	512MB	2	2560x1600		8	16	Power User
GRID K1	K180Q	4GB	4	2560x1600	4	1	4	Entry Designer
	K160Q	2GB	4	2560x1600		2	8	Power User
	K140Q	1GB	2	2560x1600		4	16	Power User
	K120Q	512MB	2	2560x1600		8	32	Power User

NVIDIA® Tesla® M60 GRID vGPU Profiles:

Card	vGPU Profile	Graphics Memory (Frame Buffer)	Virtual Display Heads	Maximum Resolution	Maximum vGPUs		Guest VM OS Support*		GRID License Required
					Per GPU	Per Card	Win	64bit Linux	
Tesla M60	M60-8Q	8GB	4	4096x2160	1	2	•	•	GRID Virtual Workstation
	M60-4Q	4GB	4	4096x2160	2	4	•	•	
	M60-2Q	2GB	4	4096x2160	4	8	•	•	
	M60-1Q	1GB	2	4096x2160	8	16	•	•	
	M60-0Q	512MB	2	2560x1600	16	32	•	•	
	M60-1B	1GB	4	2560x1600	8	16	•		GRID Virtual PC
	M60-0B	512MB	2	2560x1600	16	32	•		
	M60-8A	8GB	1	1280x1024	1	2	•		GRID Virtual Application
	M60-4A	4GB			2	4	•		
	M60-2A	2GB			4	8	•		
	M60-1A	1GB			8	16	•		

***NOTE:** Windows guest operating systems supported with all profiles. 64-Bit Linux operating systems supported are listed in the table on the next page and are only supported by vGPU 2.0 on the M60 cards.

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	

4.3.1.1 GRID vGPU Licensing and Architecture

NVIDIA GRID vGPU™ is offered as a licensable feature on Tesla M60 GPUs. vGPU can be licensed and entitled using one of the three following software editions.

NVIDIA GRID Virtual Applications	NVIDIA GRID Virtual PC	NVIDIA GRID Virtual Workstation
For organizations deploying XenApp or other RDSH solutions. Designed to deliver Windows applications at full performance.	For users who need 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 supporting virtualized Windows applications	Up to 4 displays supporting Windows desktops, and NVIDIA Quadro features	Up to 4 displays supporting Windows or Linux desktops, NVIDIA Quadro, CUDA*, OpenCL* & GPU pass-through

*CUDA and OpenCL only supported with M60-8Q profile

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.

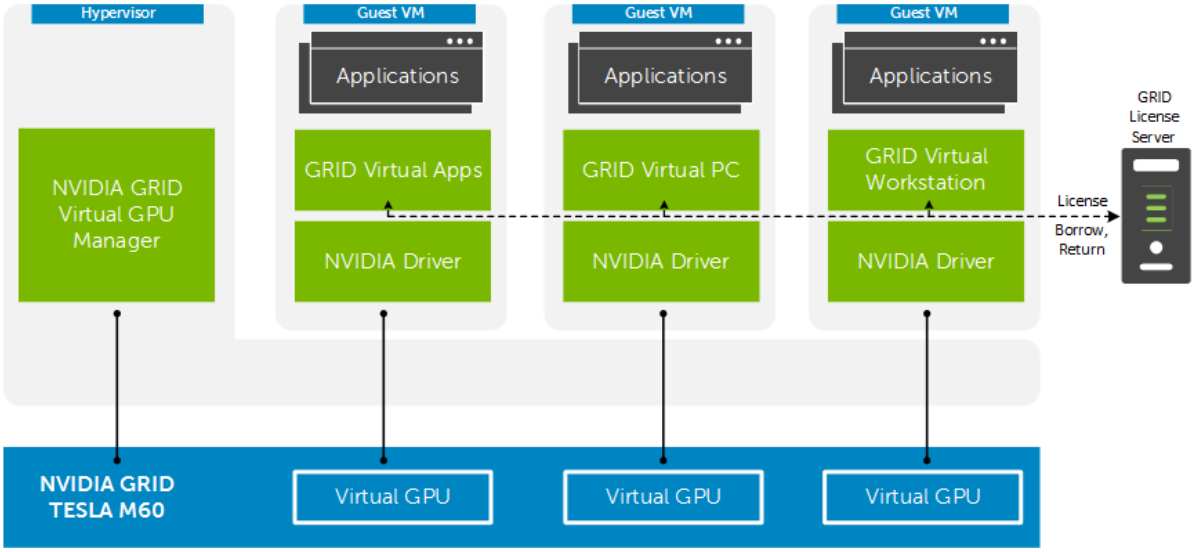

For more information about NVIDIA GRID vGPU, please visit: LINK

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)	NIC	OS vDisk	
				Size (GB)	Location
Nutanix CVM	8	16	2	-	(SATADOM)
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 vCPUs**	**40GB**	**5 vNICs**	**405GB**	**-**

5.1.2 RDSH on vSphere

The recommended number of RDSH VMs and their configurations on vSphere are summarized below based on applicable hardware platform.

Role	HW Config	VMs per host	vCPUs per VM	RAM (GB)	NIC	OS vDisk	
						Size (GB)	Location
RDSH VM	A5	3	8	32	1	80	DSF: ds_rdsh
RDSH VM	B5	5	8	32	1	80	DSF: ds_rdsh
RDSH VM	B7	9	8	32	1	80	DSF: ds_rdsh

5.1.3 SQL Databases

The VMware databases are hosted by a single dedicated SQL Server 2012 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).

5.1.4 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, 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.

5.1.4.1 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.

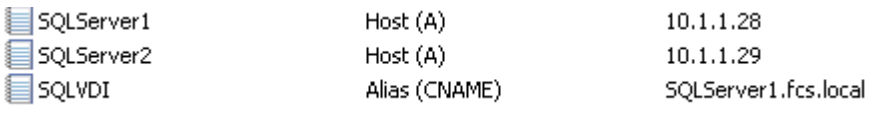

5.2 Storage Architecture Overview

All Dell XC Web Scale appliances come with two tiers of storage by default, SSD for performance and HDD for capacity. A single common Software Defined Storage namespace is created across the Nutanix cluster and presented as NFS 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 Server VMs	2	Enabled	Enabled	Disabled	Disabled
Ds_vgpu	vGPU-enabled VMs	2	Enabled	Enabled	Disabled	Disabled

*Minimum node requirement for Erasure Coding (EC-X):
-RF2 – 4 nodes
-RF3 – 6 nodes

5.3 Virtual Networking

The network configuration for the Dell XC Web Scale 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 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
  - vMotion VLAN: Configured for vMotion 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
  - vMotion VLAN: Configured for vMotion 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

The Management host network configuration consists of a standard vSwitch teamed with 2 x 10Gb physical adapters assigned. The CVM connects to a private internal vSwitch as well as the standard external vSwitch. All VMkernel service ports connect to the standard external vSwitch. All VDI infrastructure VMs connect through the primary port group on the external vSwitch.

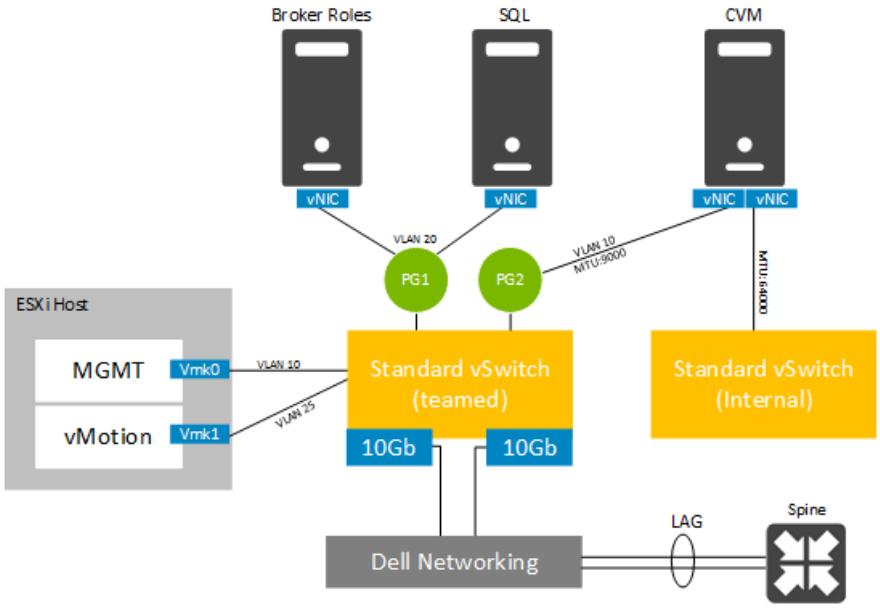

The Compute hosts are configured in the same basic manner with the desktop VMs connecting to the primary port group on the external vSwitch.

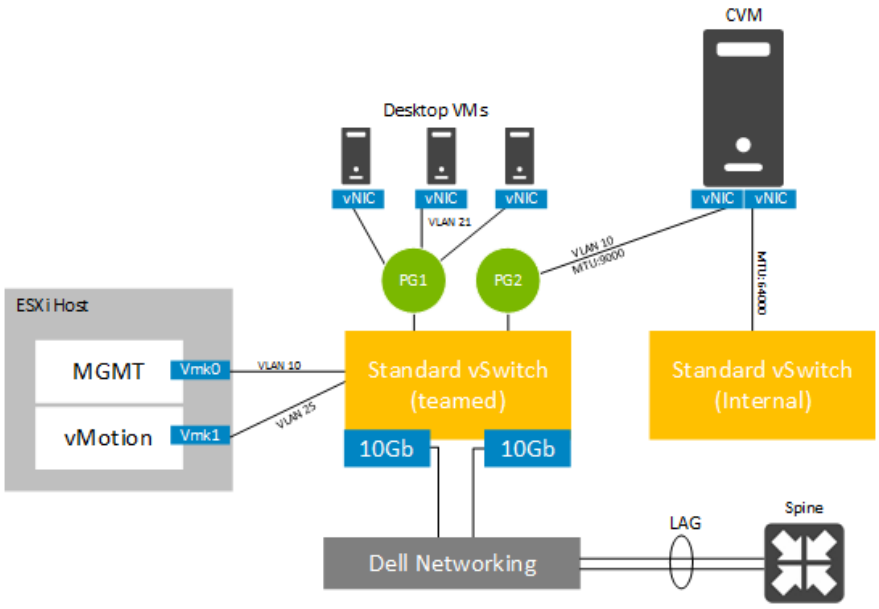

5.4 Scaling Guidance

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 mgmt 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.

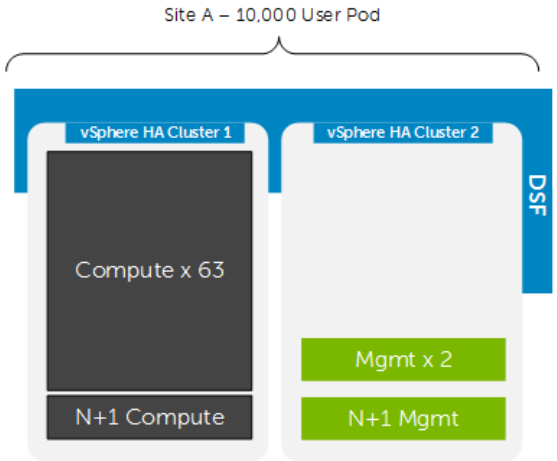

Another option is to design a large single contiguous NDFS namespace with multiple hypervisor clusters within to provide single pane of glass management. For example, portrayed below is a 30,000 professional user environment segmented by vSphere HA cluster and broker farm. Each farm compute instance is segmented into an HA cluster with a hot standby node providing N+1, served by a dedicated pair of mgmt nodes in a separate HA cluster. This provides multiple broker farms with separated HA protection while maintaining a single NDFS cluster across all nodes.

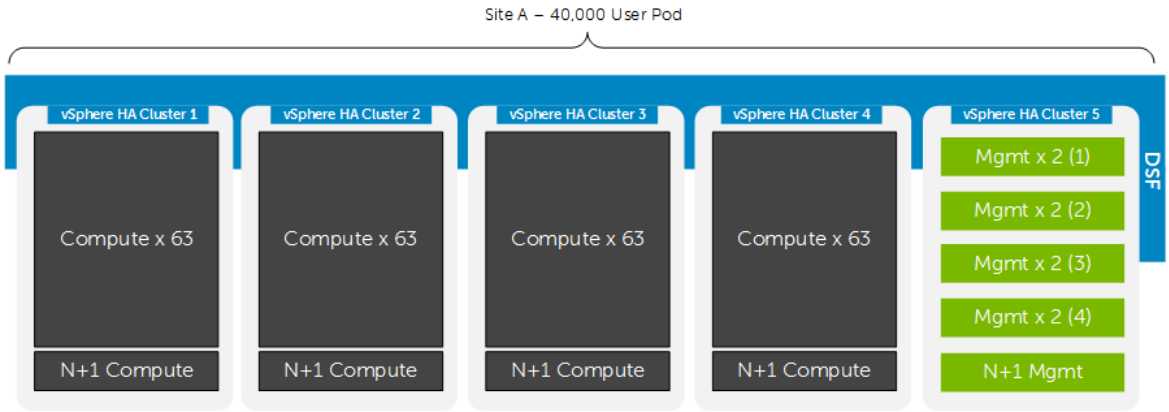

- The components are scaled either horizontally (by adding additional physical and virtual servers to the server pools) or vertically (by adding virtual resources to the infrastructure)
- Eliminate bandwidth and performance bottlenecks as much as possible
- Allow future horizontal and vertical scaling 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 RAM or CPU compute power
View Connection Servers	Desktops per instance	Additional physical servers added to the Management cluster to deal with additional management VMs.	Additional VCS Management VMs.

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 F5 is leveraged for load balancing.

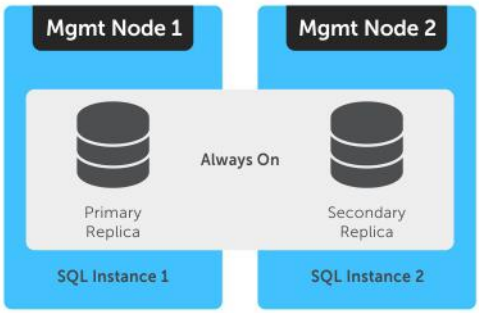

The HA options provide redundancy for all critical components in the stack while improving the performance and efficiency of the solution as a whole.

- 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 mgmt layers to provide N+1 protection.
- Applicable VMware Horizon infrastructure server roles are duplicated and spread amongst mgmt host instances where connections to each are load balanced via the addition of F5 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: LINK1 and LINK2

5.6 Dell Wyse Datacenter for Horizon Communication Flow

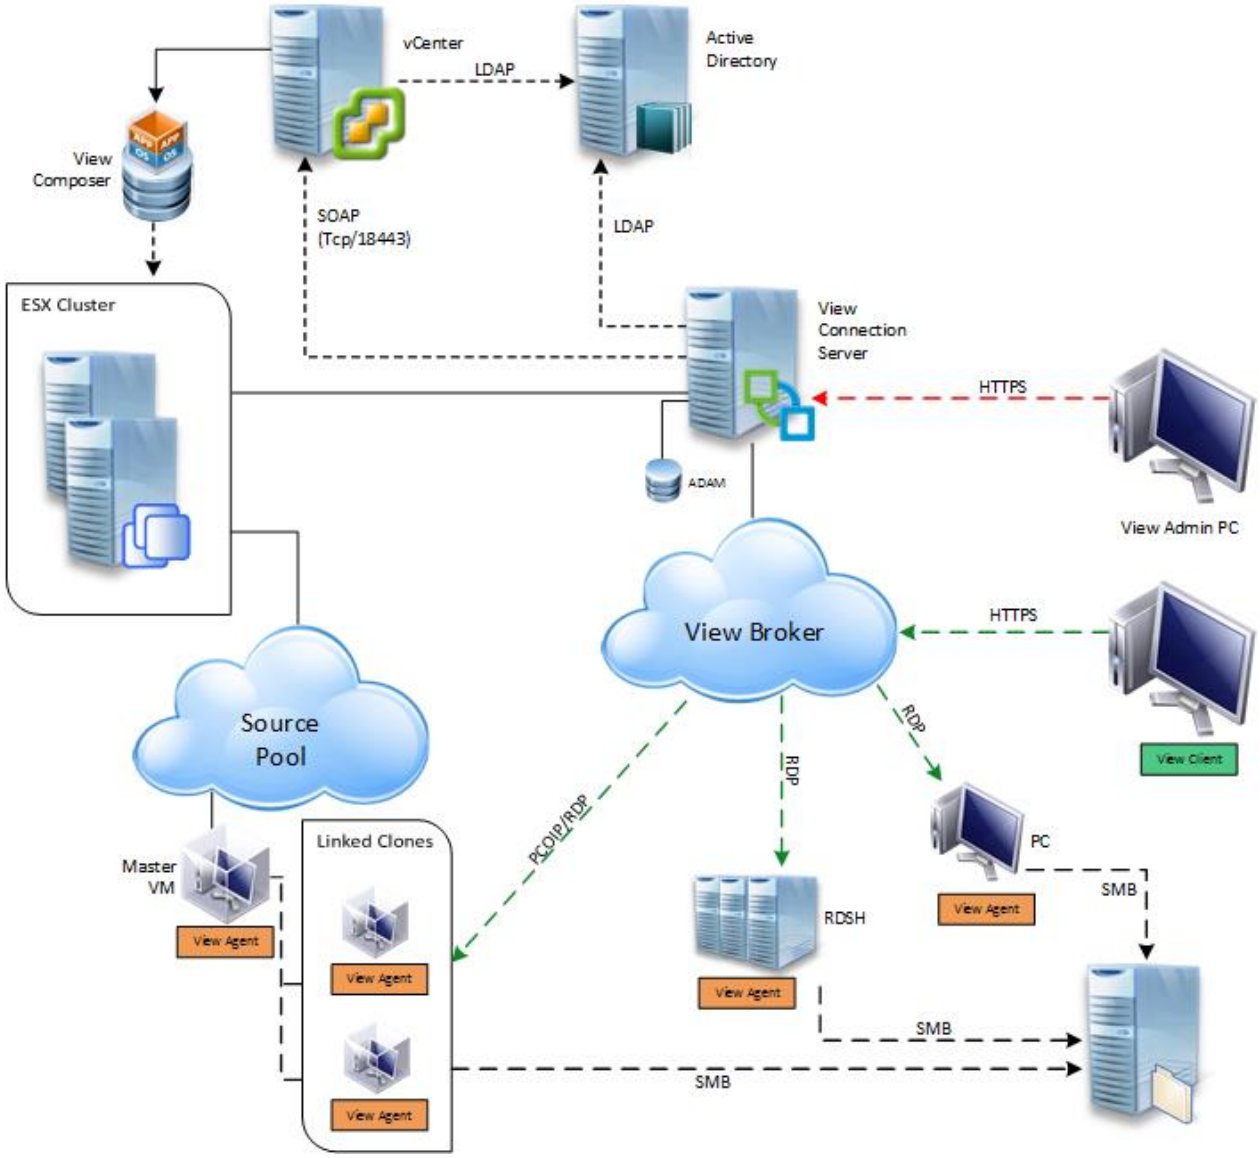

6 Solution Performance and Testing

At the time of publication these are the available per node appliance density recommendations.

Hypervisor	Provisioning	Profile	Template OS	Config	User Density
vSphere	Linked Clone	Enhanced	Windows 8.1	B5	145

For detailed up-to-date validation results, test methodology and analysis of these reference designs and more, please visit: LINK

Acknowledgements

Thank you to the Nutanix Technical Marketing and Solution Engineering teams for the detail presented in section 2.2-2.4 of this document.

Thank you to Manish Chacko of the Dell Wyse Technical Marketing team for his support and assistance with datacenter EUC programs at Dell.

About the Authors

Peter Fine is the Chief Architect and CTO of EUC Enterprise Engineering at Dell. Peter owns the strategy, architecture and leads the engineering of the datacenter EUC product and solutions portfolio. Follow Peter @ExitTheFastLane or www.ExitTheFastLane.com.

Jerry Van Blaricom is a Lead Architect in the Cloud Client Solutions Engineering Group at Dell. Jerry has extensive experience with the design and implementation of a broad range of enterprise systems and is focused on making Dell's virtualization offerings consistently best in class.

Rick Biedler is the Solutions Development Manager for Datacenter appliances at Dell, managing the development and delivery of Enterprise class Desktop virtualization solutions based on Dell datacenter components and core virtualization platforms.

Geoff Dillon is a Solutions Engineer in the Cloud Client Solutions Engineering Group at Dell. Geoff is experienced in enterprise system deployment and administration and is working on making Dell's VDI solutions stand out in the market.