Dell XC Web-Scale Converged Appliance for Wyse vWorkspace™

Dell Wyse Technical Marketing

March 2015

A Dell Appliance Architecture

Revisions

Date	Description
March 2015	Initial release

Contents

1 Introduction

1.1 Purpose

This document addresses the architecture design, configuration and implementation considerations for the key components required to deliver non-persistent FlexCast virtual desktops via Wyse vWorkspace on Microsoft® Windows Server® Hyper-V® 2012 R2 or VMware® vSphere® 5.5.

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 that 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.

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 Architecture

The Nutanix web-scale converged infrastructure is a scale-out cluster of high-performance nodes (or servers), each running a standard hypervisor and containing processors, memory, and local storage (consisting of SSD Flash and high capacity SATA disk drives). Each node runs virtual machines just like a standard virtual machine host.

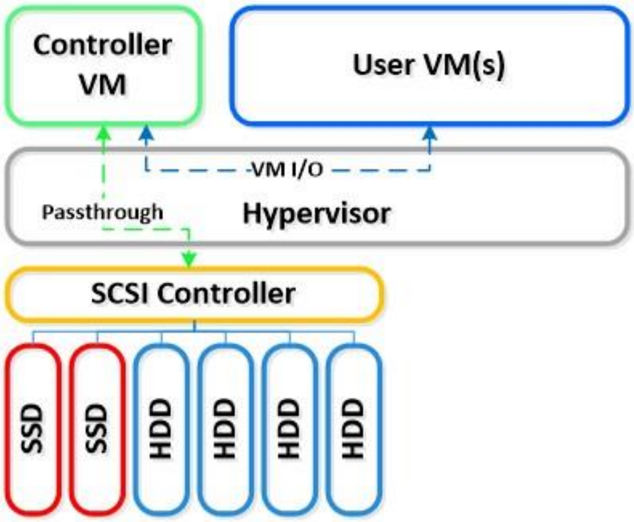

In addition, local storage from all nodes is virtualized into a unified pool by the Nutanix Distributed File System (NDFS). In effect, NDFS acts like an advanced NAS that uses local SSDs and disks from all nodes to store virtual machine data. Virtual machines running on the cluster write data to NDFS as if they were writing to shared storage.

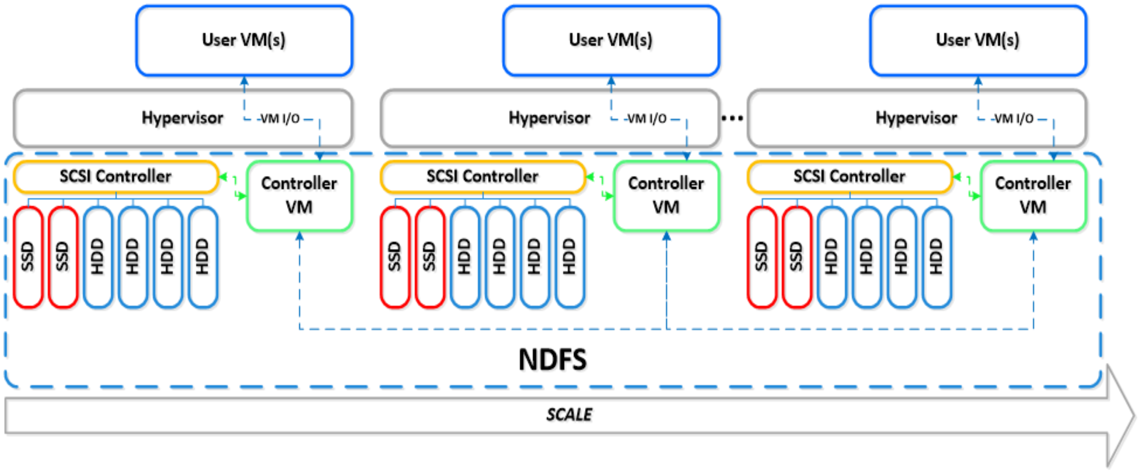

NDFS understands the concept of a virtual machine and provides advanced data management features. It brings data closer to virtual machines by storing the data locally on the system, resulting in higher performance at a lower cost. Nutanix platforms can horizontally scale from as few as three nodes to a large number of nodes, enabling organizations to scale their infrastructure as their needs grow.

The Nutanix Elastic Deduplication Engine is a software-driven, massively scalable and intelligent data reduction technology. It increases the effective capacity in the disk tier, as well as the RAM and flash cache tiers of the system, by eliminating duplicate data. This substantially increases storage efficiency, while also improving performance due to larger effective cache capacity in RAM and flash. Deduplication is performed by each node individually in the cluster, allowing for efficient and uniform deduplication at scale. This technology is increasingly effective with full/persistent clones or P2V migrations.

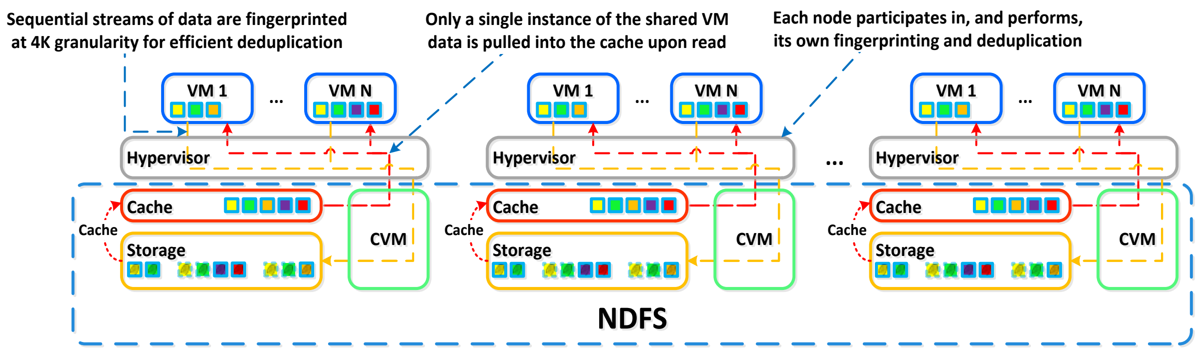

Nutanix Shadow Clones delivers distributed localized caching of virtual disks performance in multi-reader scenarios, such as desktop virtualization using vWorkspace. With Shadow Clones, the CVM actively monitors virtual disk access trends. If there are requests originating from more than two remote CVMs, 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.

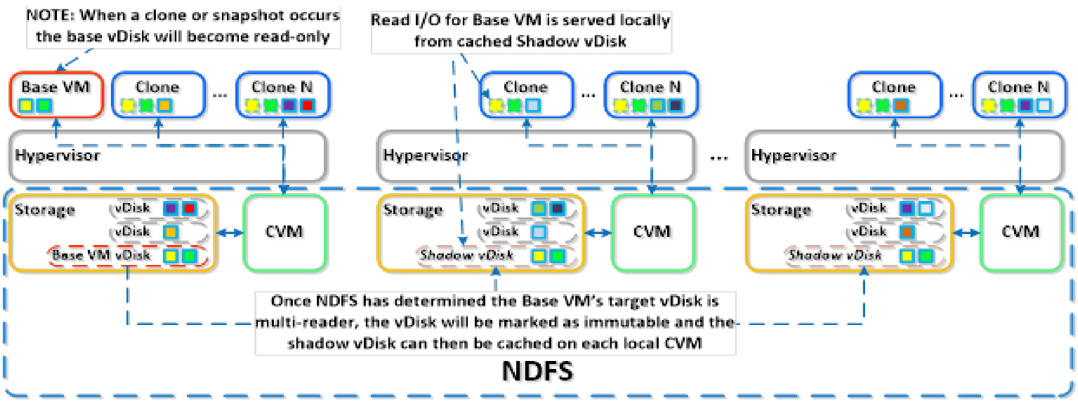

The benefits of the Nutanix Platform are now exposed to scale out vSphere or Hyper-V deployments:

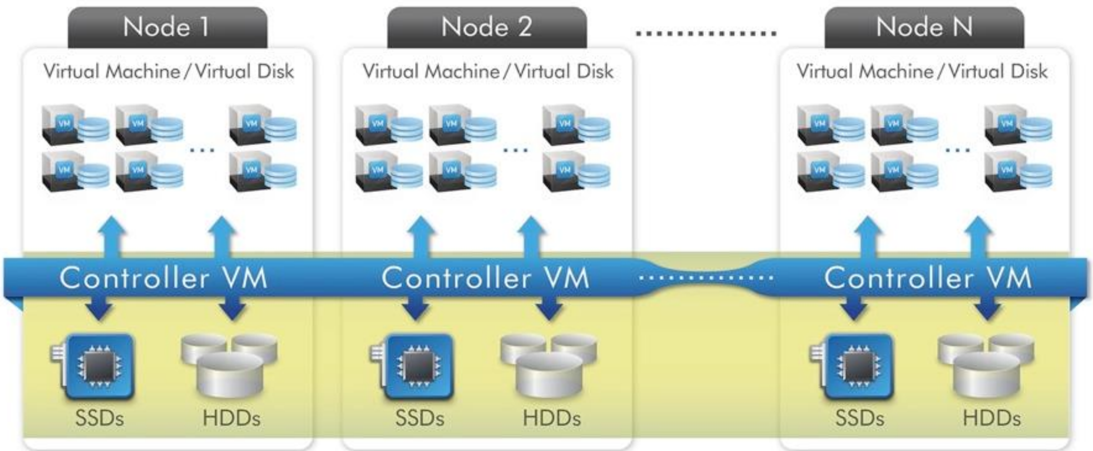

Nutanix Web-scale Converged Infrastructure

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 a Nutanix block, node, storage pool, and container:

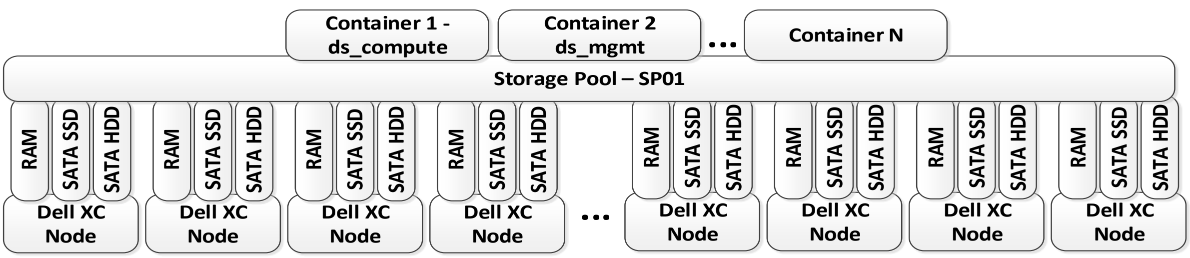

Dell XC Web Scale appliance allows organizations to deliver virtualized or remote desktops and applications through a single platform and support end users with access to all of their desktops and applications in a single place.

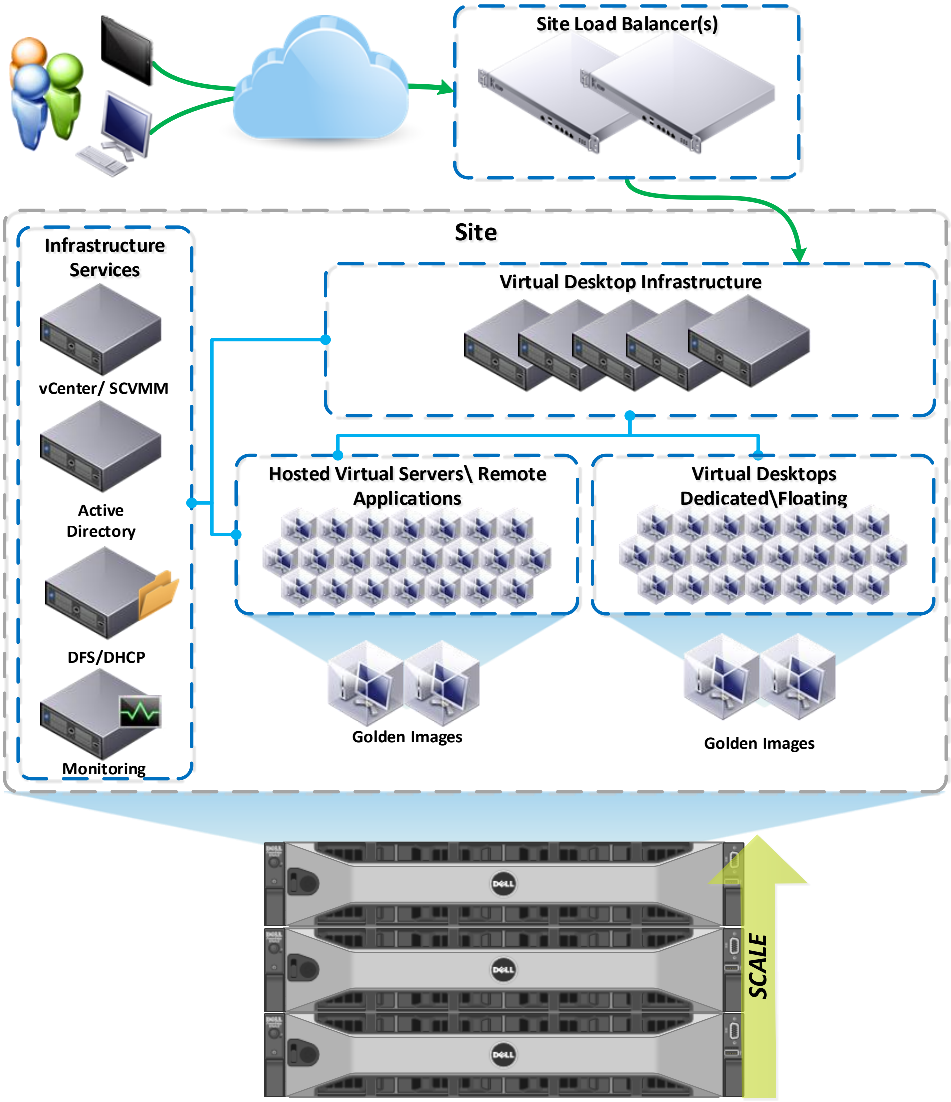

2.3 Dell XC Web Scale – Solution Pods

The networking layer consists of the 10Gb Dell Networking S4810 utilized to build a world-class leaf/spine architecture with robust 1Gb switching in the S60 for iDRAC connectivity.

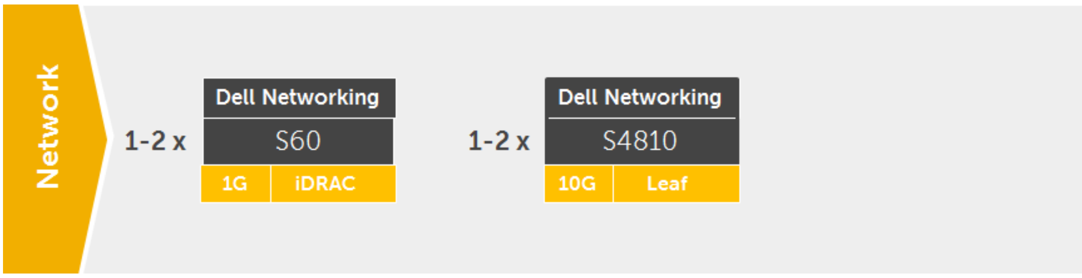

The compute, management and storage layers are 'converged' into a single XC appliance, hosting either VMware vSphere or Microsoft Hyper-V hypervisors. A minimum of three nodes per cluster is required. The recommended boundaries of an individual pod follow two schools of thought: hypervisor-based or VDI solution-based. Limits can be established based on the number of nodes supported within a given hypervisor cluster, 32 nodes for vSphere or 64 nodes for Hyper-V or the limits of a given VDI solution scale. A single Nutanix NDFS cluster can span an unlimited number of appliance nodes with several hypervisor clusters contained within.

Dell recommends that the VDI management infrastructure nodes be separated from the compute resources onto their own appliance with a common storage namespace shared between them based on NFS for vSphere and SMB for Hyper-V. One node for VDI management is required, minimally, and expanded based on size of the pod. The designations ds_rdsh, ds_compute, and ds_mgmt are logical 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 for RDSH. Distinct clusters should be built for management and compute host types for HA, respectively, to plan predictable failover, scale and load across the pod. The NFS or SMB namespace can be shared across the two hypervisor clusters adding disk capacity and performance for each distinct cluster.

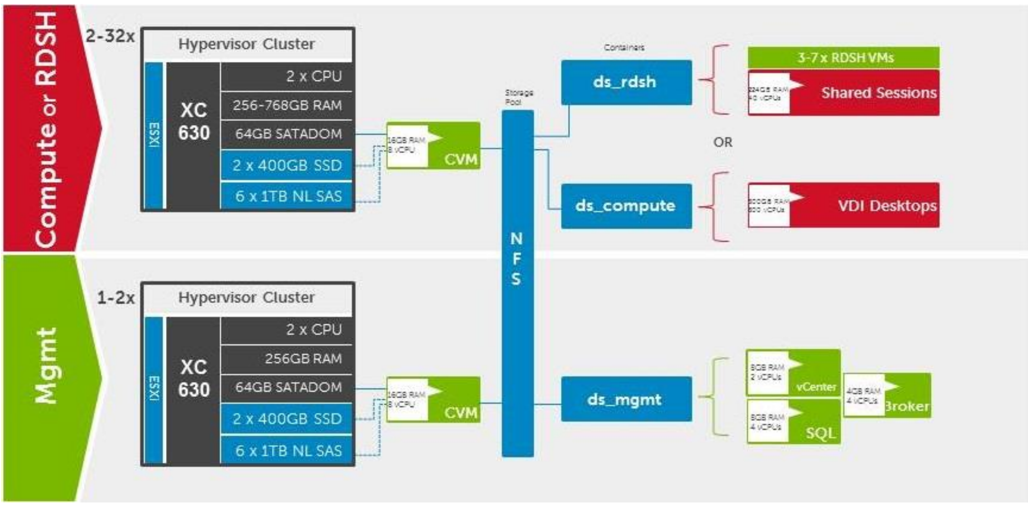

The Hyper-V solution pod scales a bit higher due to the upper node limit of 64 for a Hyper-V cluster and alternatively uses SMB as the file access protocol. The hardware platforms, logical containers and basic architecture are the same. If HA (High Availability) is required, then additional nodes should be included for compute or management, respectively.

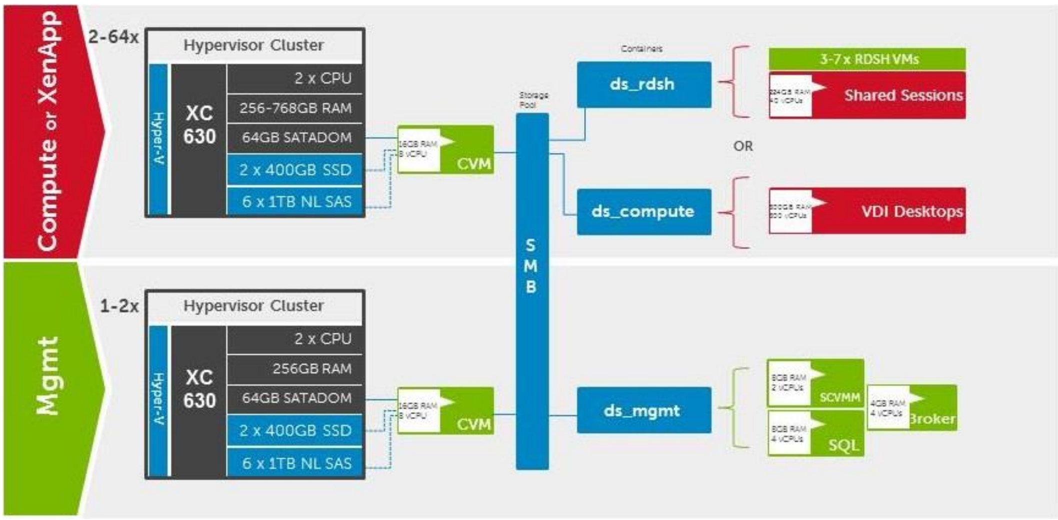

2.3.1 Network Architecture

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. This provides scale from 3 XC nodes to thousands without any impact to available bandwidth:

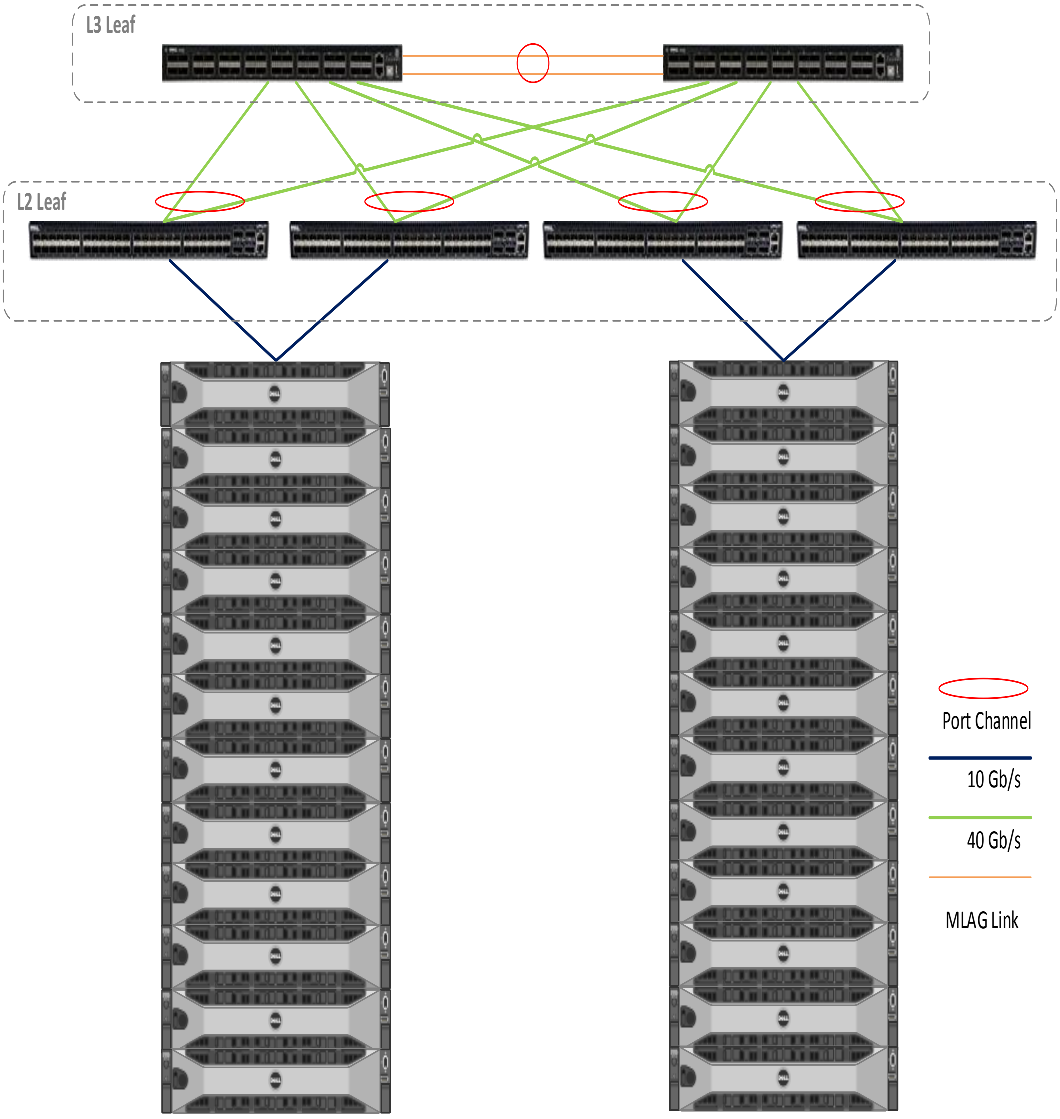
L3 Leaf
L2 Leaf
Port Channel
10 Gb/s
40 Gb/s
MLAG Link

3 Hardware Components

3.1 Network

The following sections contain the core network components for the Dell XC Web-Scale Converged Appliance for vWorkspace. 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 S60 (1Gb ToR Switch)

The Dell Networking S60 is a high-performance 1Gb access switch optimized for lowering operational costs at the network edge and is recommended for iDRAC connectivity. The S60 answers the key challenges related to network congestion in data center ToR (Top-of-Rack) and service provider aggregation deployments. As the use of bursty applications and services continue to increase, huge spikes in network traffic that can cause network congestion and packet loss, also become more common. The S60 is equipped with the industry's largest packet buffer (1.25 GB), enabling it to deliver lower application latency and maintain predictable network performance even when faced with significant spikes in network traffic. Providing 48 line-rate Gb ports and up to four optional 10Gb uplinks in just 1-RU, the S60 conserves valuable rack space. Further, the S60 design delivers unmatched configuration flexibility, high reliability, and power and cooling efficiency to reduce costs.

Model	Features	Options	Uses
Dell Networking S60	44 x BaseT (10/100/1000) + 4 x SFP High performance High Scalability	Redundant PSUs	1Gb connectivity for iDRAC
		4 x 1Gb SFP ports that support copper or fiber	
		12Gb or 24Gb stacking (up to 12 switches)	
		2 x modular slots for 10Gb uplinks or stacking modules	

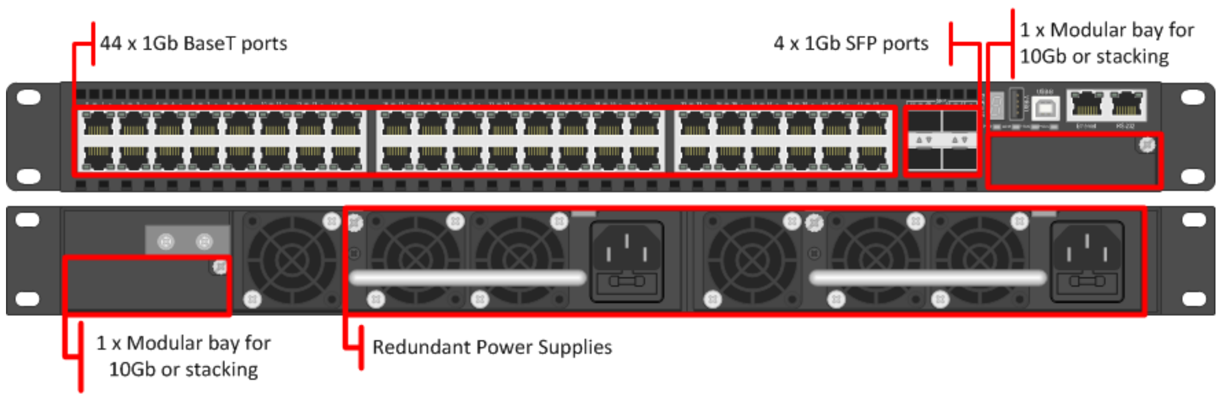

Guidance:

- 10Gb uplinks to a core or distribution switch are the preferred design choice using the rear 10Gb uplink modules. If 10Gb to a core or distribution switch is unavailable the front 4 x 1Gb SFP ports are used.
- The front 4 SFP ports can support copper cabling and are upgraded to optical if a longer run is needed.
- The S60 is appropriate for use in solutions scaling higher than 6000 users.

For more information on the S60 switch and Dell Networking products, please visit: LINK

3.1.1.1 S60 Stacking

The S60 switches are optionally stacked with 2 or more switches, if greater port count or redundancy is desired. Each switch will need a stacking module plugged into a rear bay and connected with a stacking cable. The best practice for switch stacks greater than 2 is to cable in a ring configuration with the last switch in the stack cabled back to the first. Uplinks need to be configured on all switches in the stack back to the core to provide redundancy and failure protection.

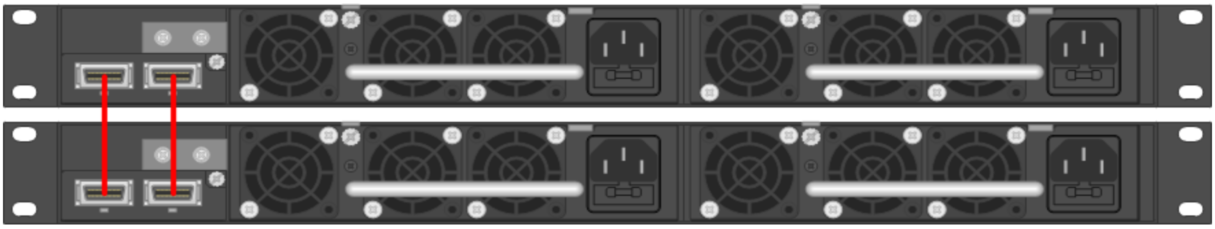

3.1.2 Dell Networking S4810 (10Gb ToR Switch)

The Dell Networking S-Series S4810 is an ultra-low latency 10/40Gb Top-of-Rack (ToR) switch purpose-built for applications in high-performance data center and computing environments. Leveraging a non-blocking, cut-through switching architecture, the S4810 delivers line-rate L2 and L3 forwarding capacity with ultra-low latency to maximize network performance. The compact S4810 design provides industry-leading density of 48 dual-speed 1/10Gb (SFP+) ports as well as four 40Gb QSFP+ uplinks to conserve valuable rack space and simplify the migration to 40Gb in the data center core (Each 40Gb QSFP+ uplink can support four 10Gb ports with a breakout cable). Priority-based Flow Control (PFC), Data Center Bridge Exchange (DCBX), Enhance Transmission Selection (ETS), coupled with ultra-low latency and line rate throughput, make the S4810 ideally suited for converged leaf/spine environments.

Model	Features	Options	Uses
Dell Networking S4810	48 x SFP+ (1Gb/10Gb) + 4 x QSFP+ (40Gb) Redundant Power Supplies	Single-mode/ multimode optics, TwinAx, QSFP+ breakout cables	ToR switch for 10Gb converged connectivity
		Stack up to 6 switches or 2 using VLT, using SFP or QSFP ports	

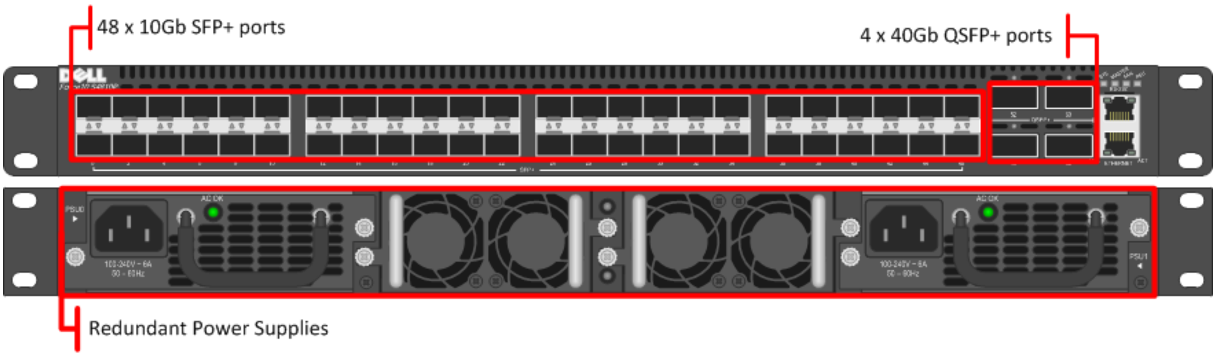

Guidance:

- The 40Gb QSFP+ ports are split into 4 x 10Gb ports using breakout cables for stand-alone units, if necessary. This is not supported in stacked configurations.
- 10Gb or 40Gb uplinks to a core or distribution switch is the preferred design choice.
- The front 4 SFP ports can support copper cabling and are upgraded to optical if a longer run is needed.
- The S60 is appropriate for use in solutions scaling higher than 6000 users.

For more information on the S4810 switch and Dell Networking, please visit: LINK

3.1.2.1 S4810 Stacking

The S4810 switches are optionally stacked up to 6 switches or configured to use Virtual Link Trunking (VLT) up to 2 switches. Stacking is supported on either SFP or QSFP ports as long as that port is configured for stacking. The best practice for switch stacks greater than 2 is to cable in a ring configuration with the last switch in the stack cabled back to the first. Uplinks need to be configured on all switches in the stack back to the core to provide redundancy and failure protection. It is recommended that the S4810 be configured for use in the leaf layer connective up to another switch in the spine layer.

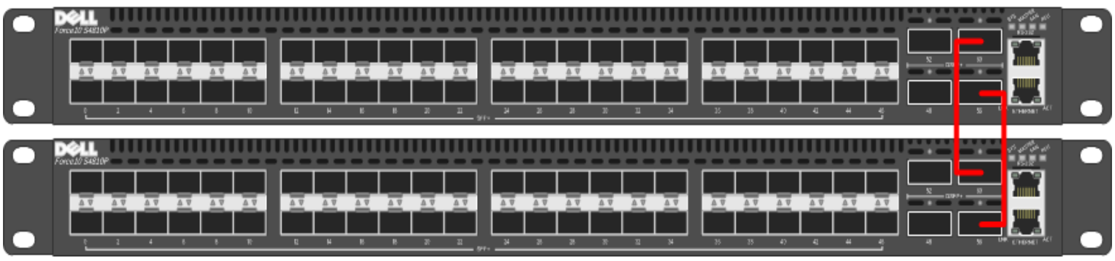

3.2 Appliance

The server platform for this appliance is the Dell XC630 (13G). It offers an optimal balance of storage utilization, performance and cost with an optional in-server hybrid storage configuration that can support tiering and capacity for up to 28 drives in a 2S/2U system, including up to 18 x 1.8" SATA SSDs. For more information please visit: [Link]

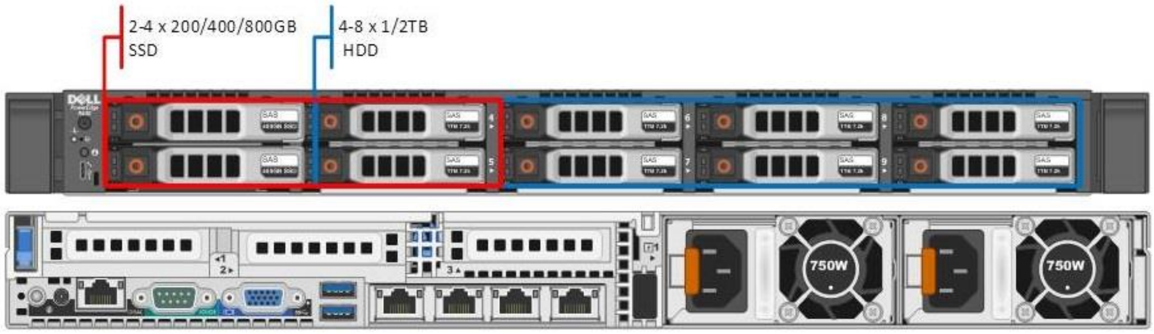

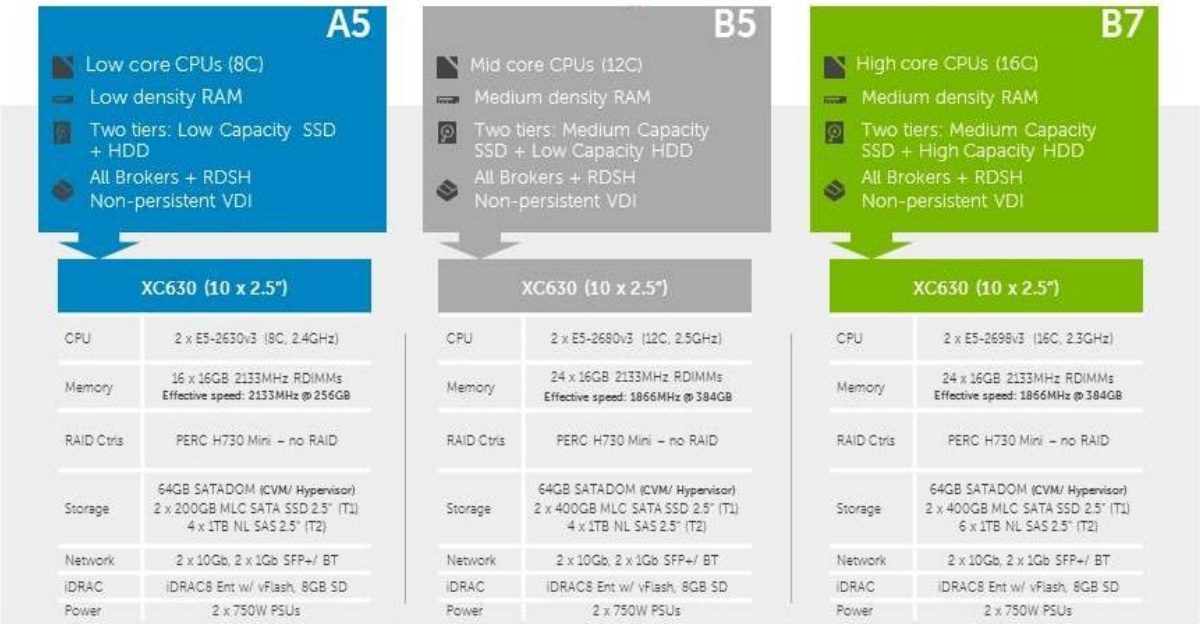

3.2.1 Dell XC A5

The Dell XC A5 platform is perfect for POCs, lighter user workloads, shared sessions or application virtualization. Each appliance comes equipped with dual 8-core CPUs and 256GB of high-performance DDR4 RAM. The 64GB SATADOM module hosts the hypervisor and Nutanix Controller VM while the H730 PERC is configured in pass-through mode to present the SSDs and HDDs to the NDFS cluster. A minimum of six disks come in each host, 2 x 200GB SSD for the performance tier (Tier1) and 4 x 1TB NL SAS disks for the capacity tier (Tier2), each tier can be expanded as required. These six disks are presented to the Nutanix Controller VM running locally on each host. Each platform can be outfitting with SFP+ or BaseT NICs.

A5 – Dell XC630	
CPU	2 x E5-2630v3 (8C, 2.4GHz)
Memory	16 x 16GB 2133MHz RDIMMs **Effective speed: 2133MHz @ 256GB**
RAID Ctrls	PERC H730 Mini – no RAID
Storage	64GB SATADOM **(CVM/ HV)** 2 x 200GB MLC SATA SSD 2.5" (T1) 4 x 1TB NL SAS 2.5" (T2)
Network	2 x 10Gb, 2 x 1Gb SFP+/ BT
iDRAC	iDRAC8 Ent w/ vFlash, 8GB SD
Power	2 x 750W PSUs

3.2.2 Dell XC B5

The Dell XC B5 platform is perfect for larger POCs, medium user workloads, shared sessions or application virtualization. Each appliance comes equipped with dual 12-core CPUs and 384GB of high-performance RAM. The 64GB SATADOM module hosts the hypervisor and Nutanix Controller VM while the H730 PERC is configured in pass-through mode to present the SSDs and HDDs to the NDFS cluster. A minimum of six disks come in each host, 2 x 400GB SSD for the performance tier (Tier1) and 4 x 1TB NL SAS disks for the capacity tier (Tier2), each tier can be expanded as required. These six disks are presented to the Nutanix Controller VM running locally on each host. Each platform can be outfitting with SFP+ or BaseT NICs.

B5 – Dell XC630	
CPU	2 x E5-2680v3 (12C, 2.5GHz)
Memory	24 x 16GB 2133MHz RDIMMs **Effective speed: 1866 MHz @ 384GB**
RAID Ctrls	PERC H730 Mini – no RAID
Storage	64GB SATADOM **(CVM/ HV)** 2 x 400GB MLC SATA SSD 2.5" (T1) 4 x 1TB NL SAS 2.5" (T2)
Network	2 x 10Gb, 2 x 1Gb SFP+/ BT
iDRAC	iDRAC8 Ent w/ vFlash, 8GB SD
Power	2 x 750W PSUs

3.2.3 Dell XC B7

The Dell XC B7 platform is ideal for high performance requirements, heavy user workloads, shared sessions or application virtualization. Each appliance comes equipped with dual 16-core CPUs and 384GB of high-performance RAM. The 64GB SATADOM module hosts the hypervisor and Nutanix Controller VM while the H730 PERC is configured in pass-through mode to present the SSDs and HDDs to the NDFS cluster. A minimum of six disks come in each host, 2 x 400GB SSD for the performance tier (Tier1) and 6 x 1TB NL SAS disks for the capacity tier (Tier2), each tier can be expanded as required. These six disks are presented to the Nutanix Controller VM running locally on each host. Each platform can be outfitting with SFP+ or BaseT NICs.

B7 – Dell XC630	
CPU	2 x E5-2698v3 (16C, 2.3GHz)
Memory	24 x 16GB 2133MHz RDIMMs **Effective speed: 1866 MHz @ 384GB**
RAID Ctrls	PERC H730 Mini – no RAID
Storage	64GB SATADOM **(CVM/ HV)** 2 x 400GB MLC SATA SSD 2.5" (T1) 6 x 1TB NL SAS 2.5" (T2)
Network	2 x 10Gb, 2 x 1Gb SFP+/ BT
iDRAC	iDRAC8 Ent w/ vFlash, 8GB SD
Power	2 x 750W PSUs

3.3 Dell Wyse Cloud Clients

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

3.3.1 Dell Wyse T10D

The T10D handles everyday tasks with ease and also provides multimedia acceleration for task workers who need video. Users will enjoy integrated graphics processing and additional WMV9 & H264 video decoding capabilities from the Marvell ARMADA™ PXA2128 1.2 GHz Dual Core ARM System-on-Chip (SoC) processor. In addition, the T10D is one of the only affordable thin clients to support dual monitors with monitor rotation, enabling increased productivity by providing an

extensive view of task work. Designing smooth playback of high bit-rate HD video and graphics in such a small box hasn't been at the expense of energy consumption and heat emissions either. Using less than 7 watts of electricity, the T10D's small size enables discrete mounting options: under desks, to walls, and behind monitors, creating cool workspaces in every respect.

3.3.2 ThinOS – D10D

Designed for knowledge workers and power users, the new Dell Wyse D10D is a high-performance thin client based on Dell Wyse ThinOS, the virus-immune firmware base designed for optimal thin client security, performance, and ease-of-use. Highly secure, compact and powerful, the D10D combines Dell Wyse ThinOS with a dual-core AMD 1.4 GHz processor and a revolutionary unified graphics engine for an outstanding user experience. The D10D 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 D10D offers a unique combination of performance, efficiency, and affordability. The D10D 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

3.3.3 Dell Wyse D90D8

A strong, reliable thin client, the D90D8 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 D90D8 offers cool, quiet operations, potentially lowering your overall carbon footprint.

WE8S

3.3.1 Dell Wyse Z90D8

The versatile Z90D8 gives people the freedom to mix and match a broad range of legacy and cutting edge peripheral devices. Ports for parallel, serial, and USB 3.0 offer fast, flexible connectivity. Like all Dell Wyse cloud clients, the new Dell Wyse Z90D8 is one cool operator. Its energy efficient processor – which out-performs other more power-hungry alternatives – and silent fan-less design, all contribute to help lower an

WE8S

organization's carbon footprint through power requirements that are a fraction of traditional desktop PCs.

3.3.2 Dell Wyse Cloud Connect

Designed to promote bring-your-own-device (BYOD) environments, Dell Wyse Cloud Connect allows you to securely access and share work and personal files, presentations, applications and other content from your business or your home. Managed through Dell Wyse Cloud Client Manager software-as-a-service (SaaS), IT managers can ensure that each Cloud Connect device is used by the appropriate person with the right permissions and access to the appropriate apps and content based on role, department and location. Slightly larger than a USB memory stick, Cloud Connect is an ultra-compact multimedia-capable device. Simply plug it into any available Mobile High-Definition Link (MHL) / HDMI port on a TV or monitor, attach a Bluetooth keyboard and mouse, and you're off and running. Easy to slip into your pocket or bag, it enables an HD-quality window to the cloud, great for meetings and presentations while on business travel, or for cruising the internet and checking email after a day of work. For more information, please visit: Link

3.3.3 Dell Chromebook 11

The lightweight, easy-to-use Dell Chromebook 11 helps turn education into exploration - without the worries of safety or security. Priced to make 1:1 computing affordable today, Chromebook 11 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 4th Gen Intel® processor provides up to 10 hours of power. Encourage creativity with the Chromebook 11 and its multimedia features that include an 11.6" screen, stereo sound and webcam.

4 Software Components

4.1 Broker Technology

The solution is based on Wyse vWorkspace™ which provides a complete end-to-end solution delivering Microsoft Windows virtual desktops or server-based hosted shared sessions 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.

Wyse vWorkspace 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.

4.1.1 Wyse vWorkspace

Wyse vWorkspace is an enterprise class desktop virtualization management solution which enables blended deployment and support of virtual desktops, shared sessions and virtualized applications. The core components of vWorkspace are:

- **Connection Broker**

  The vWorkspace Connection Broker helps users connect to their virtual desktops, applications, and other hosted resource sessions. The user's [endpoint?] sends a request to the connection broker to access their virtual environment. The connection broker processes the request by searching for available desktops, and then redirects the user to the available managed desktop or application.

- **Management Database**

  The vWorkspace Management Database is required to perform administrative functions. The management database stores all the information relevant to a vWorkspace farm, such as configuration data, administrative tasks and results, and information regarding client connections to virtual desktops and RDSH environments.

- **Management Console**

  The vWorkspace Management Console is an integrated graphical interface that helps you perform various management and administrative functions and can be installed on any workstation or server.

- **Data Collector Service**

  The vWorkspace Data Collector service is a Windows service on RDSH servers, virtual desktops, and Hyper-V hosts in a vWorkspace farm that sends a heartbeat signal and other information to the connection broker.

- **Hyper-V Catalyst Components**

  vWorkspace Hyper-V Catalyst Components increase the scalability and performance of virtual computers on Hyper-V Hosts. Hyper-V catalyst components consist of two components: HyperCache and HyperDeploy. HyperCache provides read IOPS savings and improves virtual desktop performance through selective RAM caching of parent VHDs.

HyperDeploy manages parent VHD deployment to relevant Hyper-V hosts and enables instant cloning of Hyper-V virtual computers.

- **Diagnostics and Monitoring**

  Built on Dell Software's Foglight platform, vWorkspace Diagnostics and Monitoring provides real-time and historical data for user experience, hypervisor performance, RDSH servers/applications, virtual desktops, Connection Broker servers, Web Access servers, Secure Access servers, profile servers, EOP Print servers, and farm databases.

- **User Profile Management**

  vWorkspace User Profile Management uses virtual user profiles as an alternative to roaming profiles in a Microsoft Windows environment including virtual desktops and RD Session Hosts. The virtual user profiles eliminate potential profile corruption and accelerate logon and logoff times by combining the use of a mandatory profile with a custom persistence layer designed to preserve user profile settings between sessions.

- **Web Access**

  vWorkspace Web Access is a web application that acts as a web-based portal to a vWorkspace farm. It helps users to retrieve the list of available applications and desktops by using their web browser. After successful authentication, their published desktops and applications are displayed in the web browser.

- **Secure Access**

  vWorkspace Secure Access is an SSL gateway that simplifies the deployment of applications over the Internet and can provide proxy connections to vWorkspace components such as RDP sessions, the Web Access client, and connection brokers.

- **EOP Print Server**

  vWorkspace EOP Print is a single-driver printing solution that satisfies both client-side and network printing needs in a vWorkspace environment by providing bandwidth usage control, intelligent font embedding, native printer feature support and clientless support for LAN connected print servers and remote site print servers.

vWorkspace 8.5 includes support for Microsoft Windows Server R2, Windows 8.1, Lync 2013, and App-V 5.0 as well as provides several enhancements to Diagnostics and Monitoring, Hyper-V Catalyst Components, Dell EOP and more.

For additional information about the enhancements in Wyse vWorkspace 8.5, please visit: LINK

4.1.2 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 ensures that a CPU can access its own directly-connected RAM banks faster than those banks of the other processor which are accessed via the Quick Path Interconnect (QPI). Ensuring that your virtual RDSH servers do not span physical CPUs will ensure the greatest possible performance benefit.

The general guidance for RDSH on the Dell XC appliance is as follows:

4.1.2.1 A5

Eight physical cores per CPU in the A5 platform, assuming 2x oversubscription per core, gives us a total of 32 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.

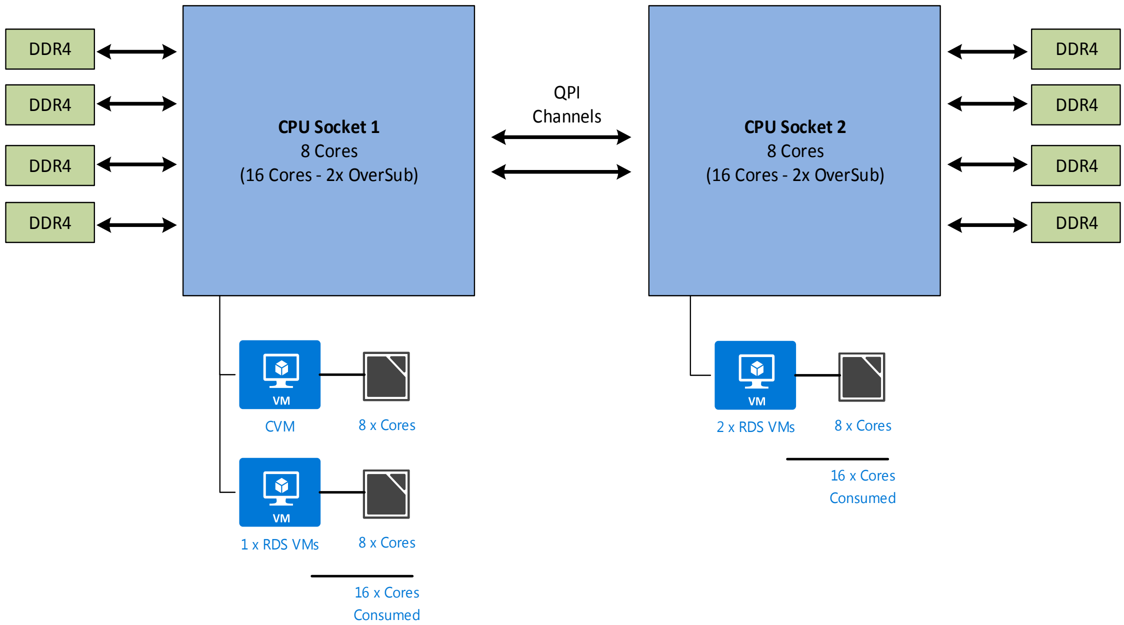

4.1.2.2 B5

12 physical cores per CPU in the B5 platform, assuming 2x oversubscription per core, nets a total of 48 consumable cores per appliance. The Nutanix CVM receives its vCPU allotment from the first physical CPU so configuring the RDSH VMs as shown below will ensure that no NUMA spanning occurs which can lower performance.

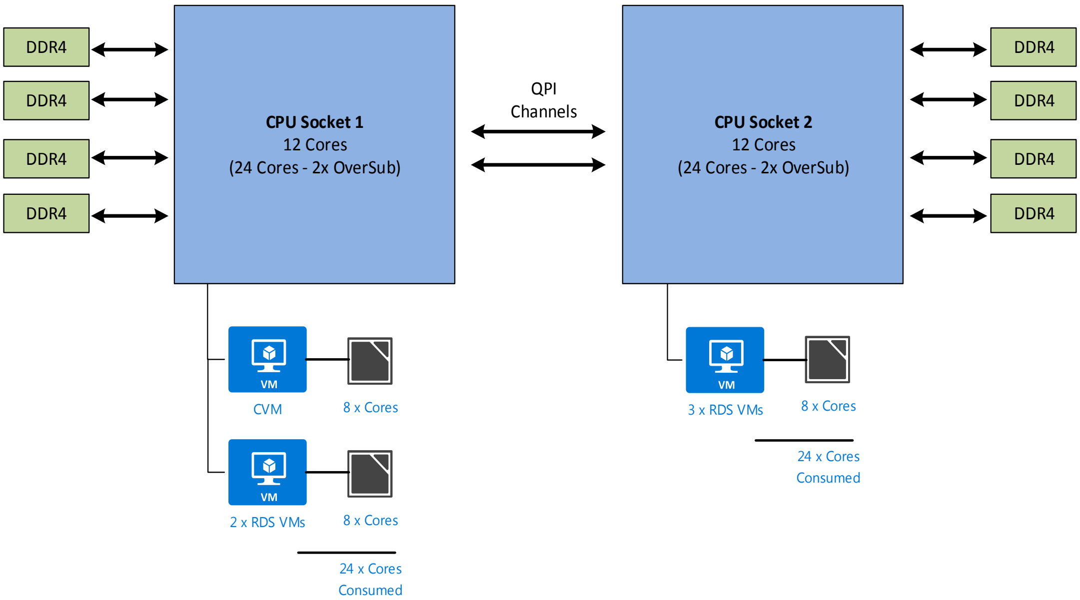

4.1.2.3 B7

16 physical cores per CPU in the B7 platform, assuming 2x oversubscription per core, nets a total of 64 consumable cores per appliance. The Nutanix CVM receives its vCPU allotment from the first physical CPU so configuring the RDSH VMs as shown below will ensure that no NUMA spanning occurs which can lower performance.

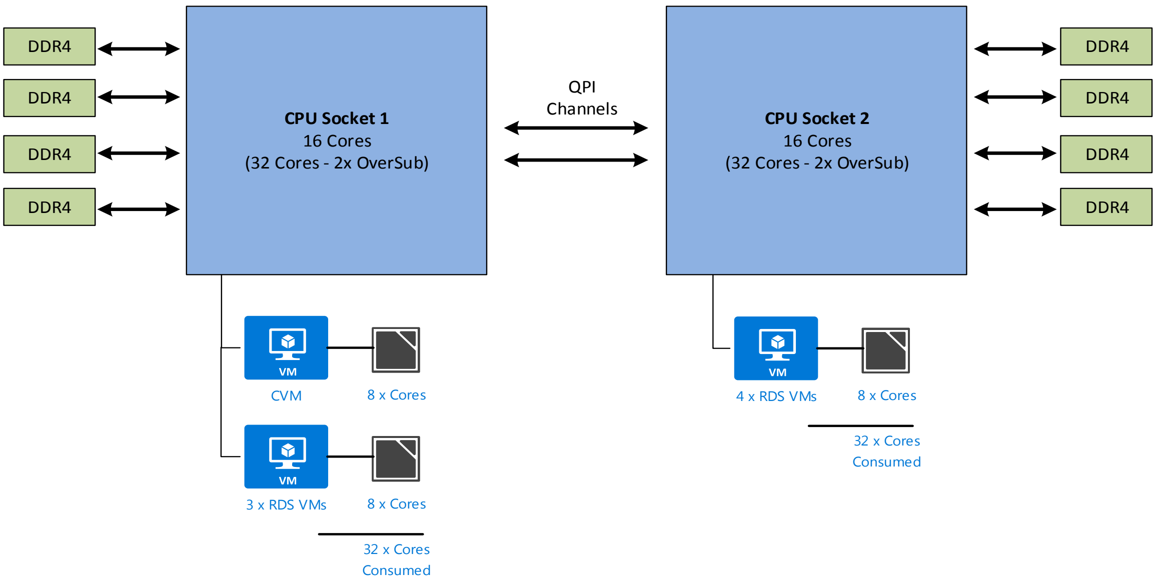

4.2 Hypervisor Platform

4.2.1 Microsoft Hyper-V

Windows Server 2012 R2 Hyper-V ™ is a powerful virtualization technology that enables businesses to leverage the benefits of virtualization. Hyper-V reduces costs, increases hardware utilization, optimizes business infrastructure, and improves server availability. Hyper-V works with virtualization-aware hardware to tightly control the resources available to each virtual machine. The latest generation of Dell servers includes virtualization-aware processors and network adapters.

From a network management standpoint, virtual machines are much easier to manage than physical computers. To this end, Hyper-V includes many management features designed to make managing virtual machines simple and familiar, while enabling easy access to powerful VM-specific management functions. The primary management platform within a Hyper-V based vWorkspace virtualization environment is Microsoft Systems Center Virtual Machine Manager 2012 R2 (SCVMM).

SCVMM provides centralized and powerful management, monitoring, and self-service provisioning for virtual machines. SCVMM host groups are a way to apply policies and to check for problems across several VMs at once. Groups are organized by owner, operating system, or by custom names such as "Development" or "Production". The interface also incorporates Remote Desktop Protocol (RDP); double-click a VM to bring up the console for that VM—live and accessible from the management console. For additional information about the enhancements to Hyper-V in Microsoft Windows Server 2012 R2, please visit: LINK

4.2.2 VMware vSphere 5.5

VMware vSphere 5 is a virtualization platform used for building VDI and cloud infrastructures. vSphere 5 represents a migration from the ESX architecture to the ESXi architecture. VMware vSphere 5 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 client and the vSphere web client.

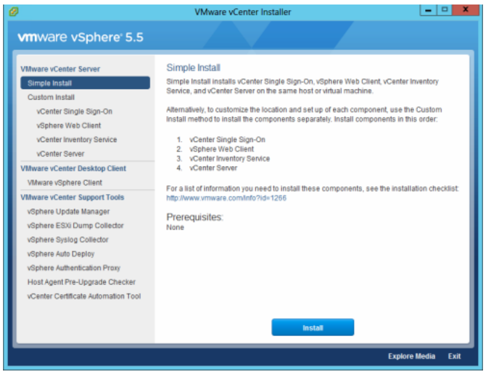

Throughout the Dell Wyse Datacenter solution, all VMware and Microsoft best practices and prerequisites for core services are adhered to (NTP, DNS, Active Directory, etc.). The vCenter 5 VM used in the solution is a single Windows Server 2008 R2 VM or vCenter 5 virtual appliance, residing on a host in the management Tier. SQL server is a core component of the Windows version of vCenter and is hosted on another VM also residing in the management Tier. It is recommended that all additional vWorkspace components be installed in a distributed architecture, one role per server VM.

4.3 Operating Systems

4.3.1 Microsoft Windows Server 2012 R2

Microsoft Windows Server 2012 R2 is the latest iteration of the Windows Server operating system environment. This release introduces a host of new features and enhancements, including virtualization, storage, networking, management, security and applications. With this release also come the introduction of Microsoft Cloud OS and an update of products and services to further enable customers' shift to cloud enablement.

For additional information about the enhancements in Microsoft Windows Server 2012 R2, please visit: LINK

4.3.2 Microsoft Windows 8.1

Microsoft Windows 8.1 is an update to the latest Windows desktop operating system, providing several user centric features. With updates to the user interface, applications, online services, security and more, Windows 8.1 helps keeps a consistent user experience across virtual and physical instances.

For additional information about the enhancements in Microsoft Windows 8.1, please visit: LINK

4.4 Application Virtualization

Microsoft Application Virtualization (App-V) provides multiple methods to deliver virtualized applications to RDS environments, virtual desktops, physical desktops, connected as well as disconnected clients. App-V can help reduce the costs and time associated with managing gold master VM and PC images with integrated applications. App-V also removes potential conflicts, such as legacy application compatibility, since virtualized applications are never installed on an end point. Once an application has been packaged using the Microsoft Application Virtualization Sequencer, it can be saved to removable media, streamed to desktop clients or presented to session-based users on a RDSH host. App-V provides a scalable framework that can be managed by System Center Configuration Manager for a complete management solution.

To learn more about application virtualization and how it integrates into a RDS environment please visit: LINK

For more information about vWorkspace and App-V integration, reviews the administration guide: LINK

4.4.1 RDSH Integration into Dell Wyse Datacenter Architecture

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

- One or more Server OS instances running the PNTools added to the vWorkspace site

The total number of required virtual RDSH servers is dependent on application type, quantity and user load and appliance capability ensuring that proper NUMA architecture boundaries are adhered to. Deploying RDSH virtually and in a multi-server farm configuration increases overall farm performance, application load balancing as well as farm redundancy and resiliency.

5 Solution Architecture for vWorkspace

5.1 Management Infrastructure

5.1.1 Hyper-V

The Management role requirements for Hyper-V based solutions are summarized below:

Role	vCPU	Startup RAM (GB)	Dynamic Memory			NIC	OS vDisk	
			Min|Max	Buffer	Weight		Size (GB)	Location
Nutanix CVM	8*	32	Dynamic Memory Disabled		High	2	-	C:\ (SATADOM)
Broker + Licensing	4	4	512MB | 8GB	20%	Med	1	40	SDS: ds_mgmt
Web Access + SGW	2	4	512MB | 6GB	20%	Med	1	40	SDS: ds_mgmt
vFoglight	2	4	512MB | 6GB	20%	Med	1	60	SDS: ds_mgmt
vFoglight AM	2	4	512MB | 6GB	20%	Med	1	60	SDS: ds_mgmt
Primary SQL	4	8	512MB | 10GB	20%	Med	1	40 + 200	SDS: ds_mgmt
Total	**24 vCPUs**	**60GB**	**3.5GB | 44GB**	-	-	**8vNICs**	**500GB**	-

5.1.2 ESXi

The Management role requirements for ESXi based solutions are summarized below:

Role	vCPU	Startup RAM (GB)	NIC	OS vDisk	
				Size (GB)	Location
Nutanix CVM	8*	32	2	-	C:\ (SATADOM)
Broker + Licensing	4	8	1	40	SDS: ds_mgmt
Web Access + SGW	2	6	1	40	SDS: ds_mgmt
vFoglight	2	6	1	60	SDS: ds_mgmt
vFoglight AM	2	6	1	60	SDS: ds_mgmt
Primary SQL	4	10	1	40 + 200	SDS: ds_mgmt
vCenter	2	8	1	60	SDS: ds_mgmt
Total	**24**	**60**	**8**	**500**	**-**

5.1.3 SQL Databases

The vWorkspace databases will be hosted by a single dedicated SQL 2012 Server VM in the Management layer. This architecture provides configuration guidance using a dedicated SQL Server VM to serve the environment. Use caution during database setup to ensure that SQL data, logs, and TempDB are properly separated onto their respective volumes and auto-growth is enabled for each database. Initial placement of all databases into a single SQL instance is fine unless performance becomes an issue, in which case the database needs to be separated into separate named instances. Best practices defined by Dell and Microsoft are to be adhered to, to ensure optimal database performance.

5.1.4 DNS

DNS plays a crucial role in the environment not only as the basis for Active Directory but will be used to control access to the various Dell 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 during the initial deployment. CNAMEs and the round robin DNS mechanism should be employed 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 would be 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 XC Web Scale appliances come with two tiers of storage by default, SSD for performance and HDD for capacity. A single common software defined 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 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 CPU demand on the Controller VMs.

Container	Purpose	Replication Factor	Perf Tier Deduplication	Capacity Tier Deduplication	Compression
Ds_compute	Desktop VMs	2	Enabled	Disabled	Disabled
Ds_mgmt	Mgmt Infra VMs	2	Enabled	Disabled	Disabled
Ds_rdsh	RDSH Server VMs	2	Enabled	Disabled	Disabled

5.3 Networking

5.3.1 Hyper-V

The Hyper-V configuration, while identical in core requirements and hardware, is executed differently due to how Hyper-V and Windows Server 2012 R2 implement networking and virtual switches. As shown in the diagram below, native Windows Server 2012 R2 NIC Teaming is utilized to load balance and provide resiliency for network connections. For the compute host in this scenario, a single LBFO NIC team is configured to connect to a Hyper-V switch for external traffic and one internal Hyper-V switch is used for the Nutanix CVM. All vNICs associated with the Management OS connect directly to the external Hyper-V switch.

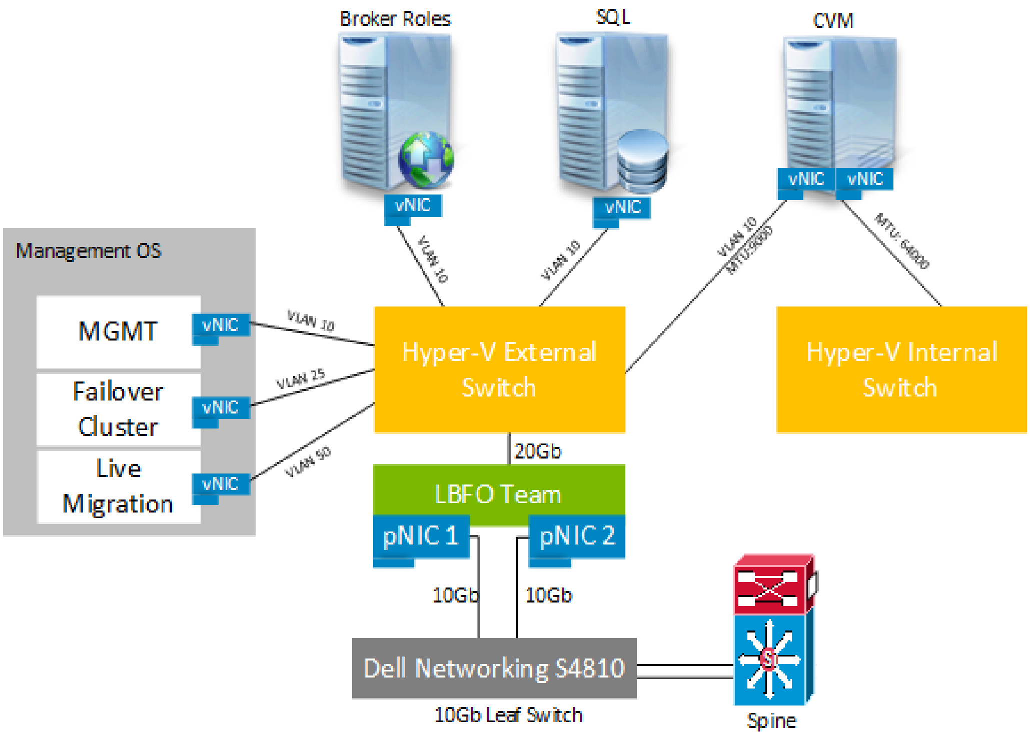

The NIC team for the Hyper-V switch is configured as switch independent, dynamic for the load balancing mode with all adapters set to active. This team is used exclusively by Hyper-V.

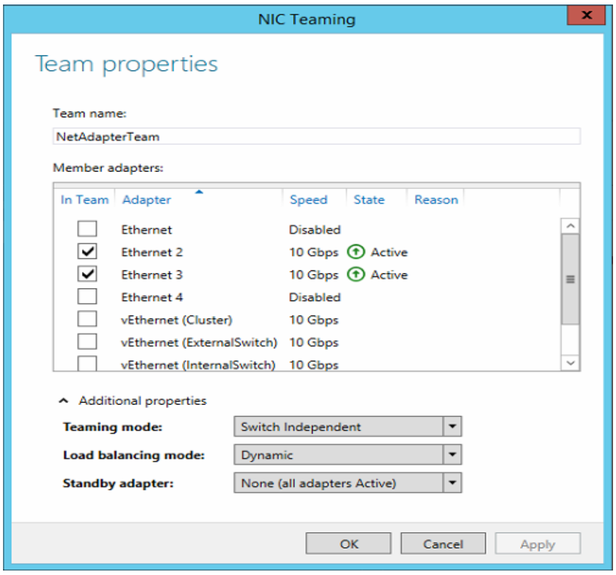

The dedicated compute host configuration is shown in the diagram below and configured very similarly to the management host configuration.

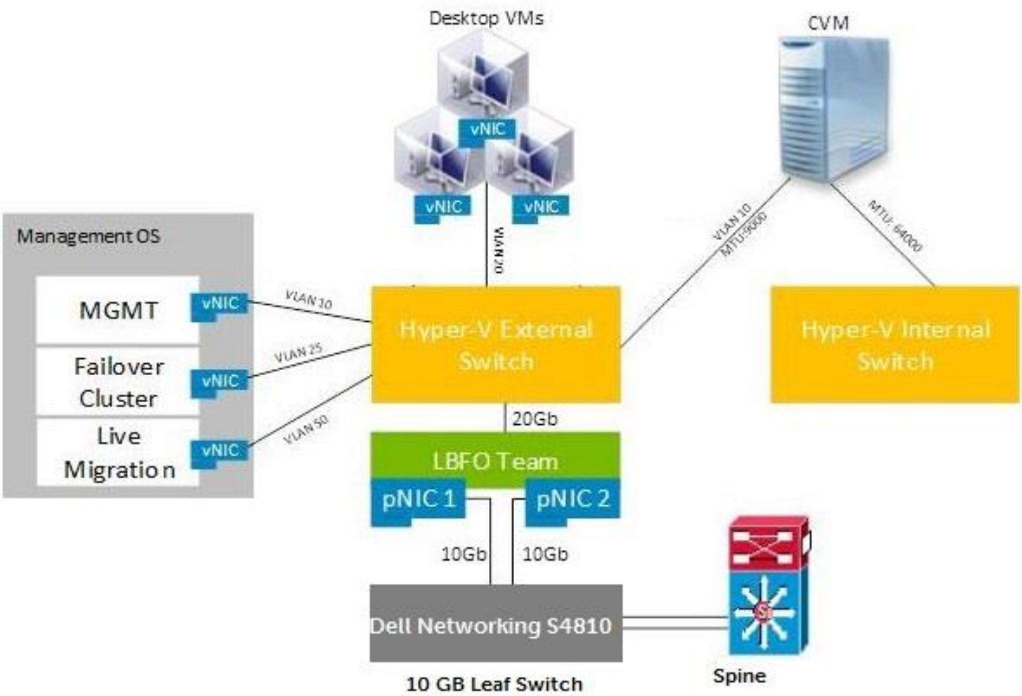

5.3.2 vSphere

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.

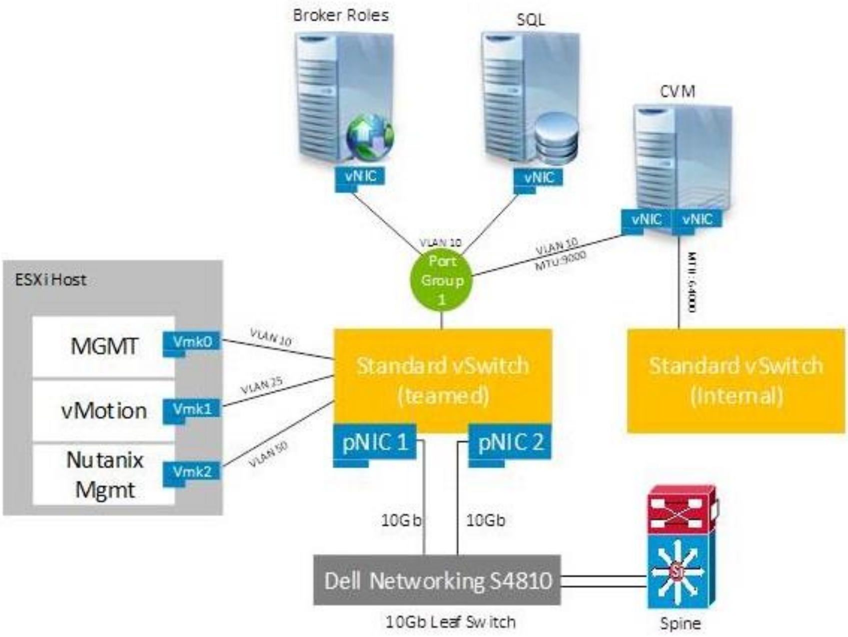

The Compute hosts are configured in the same basic manner with the desktop VMs connecting to the primary port group on the external vSwitch.

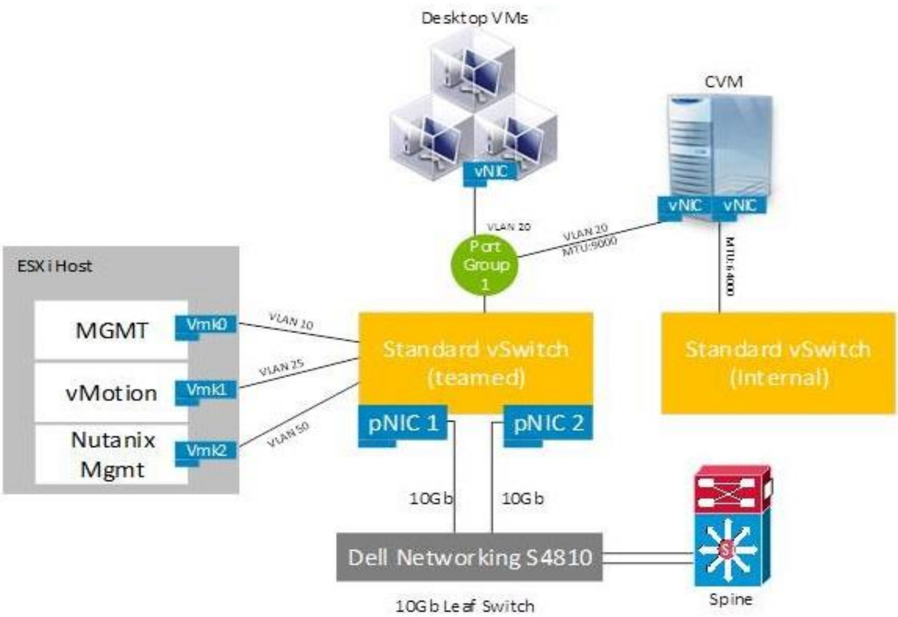

5.4 Scaling Guidance

Each component of the solution architecture scales independently according to the desired number of supported users. Brokers scale differently from other management VMs, as does the hypervisor in use.

- 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
Broker Servers	Desktops per instance (dependent on SQL performance as well)	Additional servers added to the farm	Additional virtual machine resources (RAM and CPU)
RDSH Servers	Desktops per instance	Additional virtual RDSH servers added to the farm	Additional physical servers to host virtual RDSH servers
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
Monitoring Services	Managed agents/units (dependent on SQL performance as well)	Add additional monitoring servers and migrate databases to a dedicated SQL server	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, Hyper-V clustering is introduced in both the management and compute layers, SQL is mirrored or clustered and NetScaler is leveraged for load balancing. Storage protocol switch stacks and NAS selection will vary based on chosen solution architecture.

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

- An additional switch is added at the network Tier which is configured with the original as a stack and thereby equally spreading each host's network connections across both switches.
- At the compute Tier an additional Hyper-V host is added to provide N+1 protection.
- A number of enhancements occur at the Management Tier, the first of which is the addition of another host. The Management hosts are configured in an HA cluster. All applicable vWorkspace infrastructure server roles can then be duplicated on the new host where connections to each are load balanced via the addition of a load balancer appliance. SQL Server databases also receive greater protection through the addition and configuration of a 3-way SQL mirror with a witness or AlwaysOn configuration.

5.6 Wyse vWorkspace Communication Flow

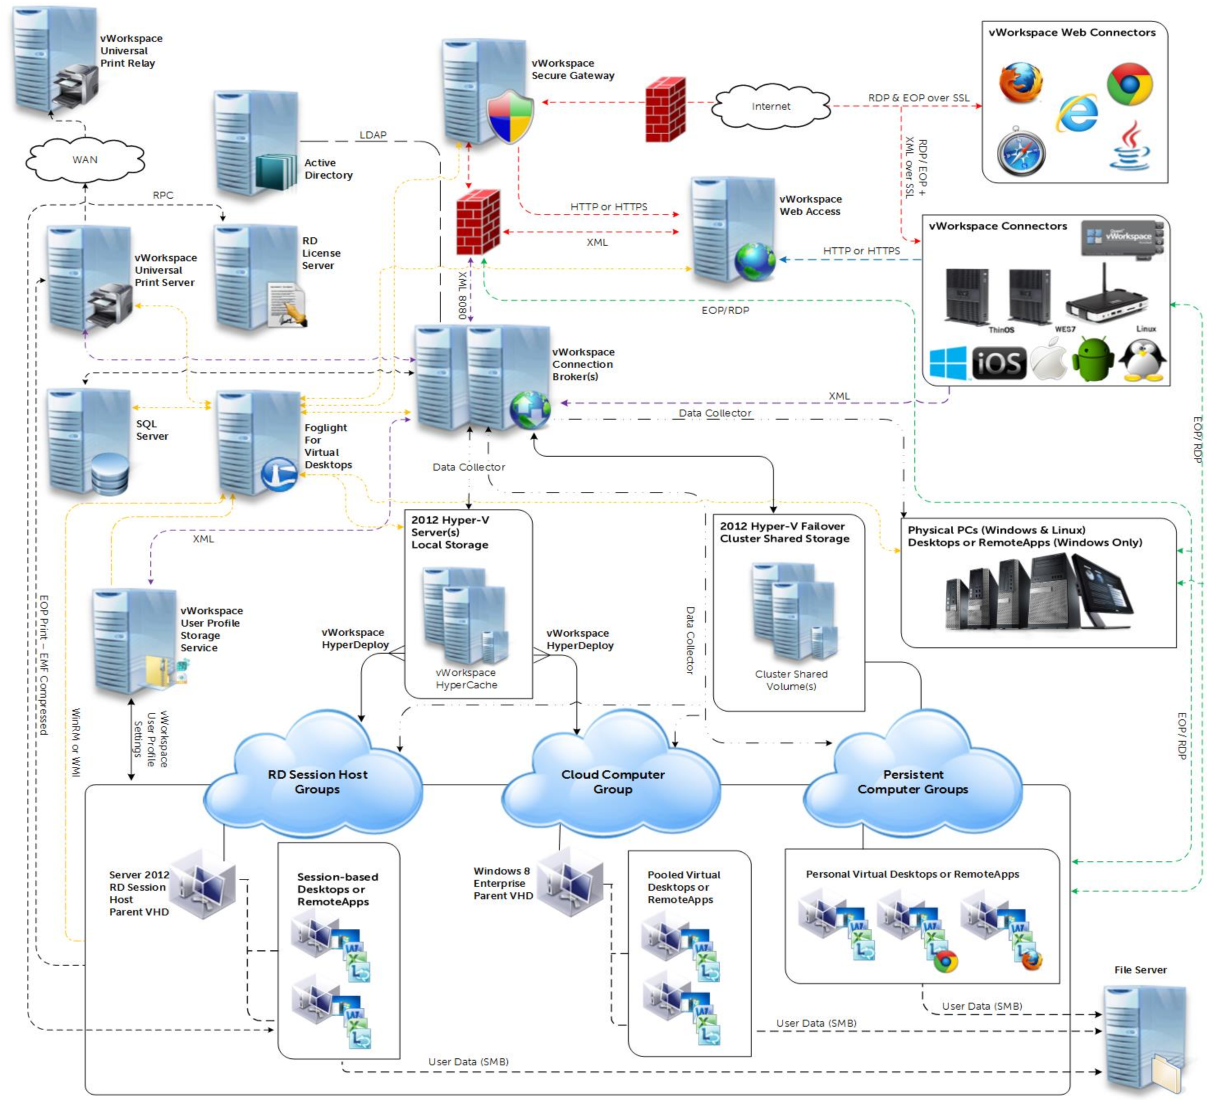

6 Solution Performance and Testing

6.1 Load Generation and Monitoring

6.1.1 Login VSI 4 – Login Consultants

Login VSI is the de-facto industry standard tool for testing VDI environments and server-based computing / terminal services environments. It uses launcher systems to connect a specified number of users to available desktops within the environment and then executes a standard collection of desktop application software (e.g. Microsoft Office, Adobe Acrobat Reader etc.) on each VDI desktop/session. Once the user is connected, a logon script configures the user environment and then starts the test workload. Each launcher system can launch connections to multiple 'target' machines (i.e. VDI desktops). A centralized management console is used to configure and manage the Login VSI environment and launcher systems.

6.1.2 Microsoft Perfmon

Microsoft Perfmon was utilized to collect performance data for tests performed on the Hyper-V platform.

6.2 Performance Analysis Methodology

In order to ensure the optimal combination of end-user experience (EUE) and cost-per-user, performance analysis and characterization (PAAC) on Dell Wyse Datacenter solutions is carried out using a carefully designed, holistic methodology that monitors both hardware resource utilization parameters and EUE during load-testing. This methodology is based on the three pillars shown below. Login VSI is currently the load-testing tool used during PAAC of Dell Wyse Datacenter solutions; Login VSI is the de-facto industry standard for VDI and server-based computing (SBC) environments and is discussed in more detail below.

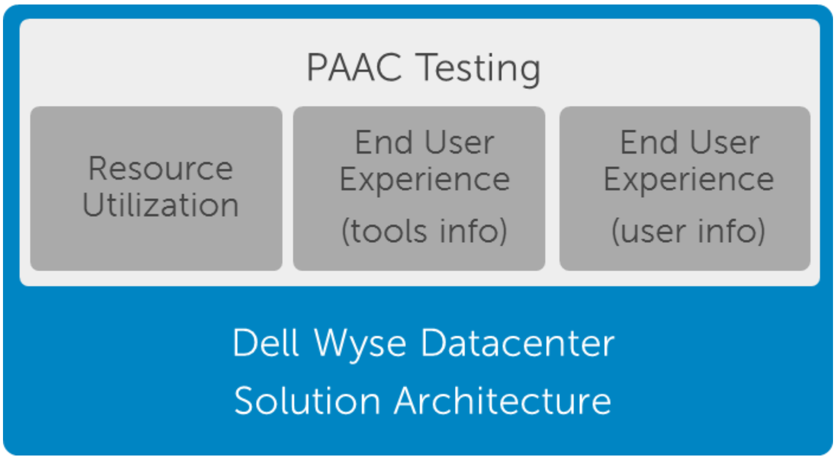

Each of the pillars shown in the above diagram is discussed in more detail below.

6.2.1 Resource Utilization

Poor end-user experience is one of the main risk factors when implementing desktop virtualization but the 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 has not happened (and that it is not close to happening), PAAC on Dell Wyse Datacenter solutions monitors the relevant resource utilization parameters and applies relatively conservative thresholds as shown in the table below. As discussed above, these 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. These thresholds are used to decide the number of virtual 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 Wyse Datacenter RA.

Resource Utilization Thresholds	
Parameter	**Pass / Fail Threshold**
Physical Host CPU Utilization	85%
Physical Host Memory Utilization	85%
Network Throughput	85%
Storage IO Latency	20ms

6.2.2 Dell Wyse Datacenter Workloads and Profiles

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, we use five workload / profile levels, each of which is bound by specific metrics and capabilities. In addition, we use workloads and profiles that are targeted at graphics-intensive use cases. We present more detailed information in relation to these workloads and profiles below but first it is useful to define the terms "workload" and "profile" 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 Wyse Datacenter solutions e.g. Microsoft Office applications, PDF Reader, Internet Explorer etc.

6.2.3 Dell Wyse Datacenter Profiles

The table shown below presents the persistent user profiles used during PAAC on this solution. These profiles have been carefully selected to provide the optimal level of resources for common use cases.

User Profile	vCPUs	Hyper-V Start up Memory	Hyper-V Minimum Memory	Hyper-V Max Memory	ESXi Memory Reservation	ESXi Memory Configured	OS
Standard	1	2GB	1GB	2GB	2GB	2GB	x86
Enhanced	2	3GB	1GB	3GB	3GB	3GB	x86
Professional	2	4GB	1GB	4GB	4GB	4GB	x64

Standard user profile equates to a task worker, Enhanced to a knowledge worker and Professional to a power user.

6.2.4 Dell Wyse Datacenter Workloads

Load-testing on each of the profiles described in the above table is carried out using an appropriate workload that is representative of the relevant use case. In the case of the non-graphics workloads, these workloads are Login VSI workloads and in the case of graphics workloads, these are specially designed workloads that stress the VDI environment to a level that is appropriate for the relevant use case. This information is summarized in the table below:

Profile Name	Workload	OS Image
Standard	Login VSI Light	Full Clone Copies
Enhanced	Login VSI Medium	Full Clone Copies
Professional	Login VSI Heavy	Full Clone Copies

As noted in the table above, further information for each of the workloads is given below. It is noted that for Login VSI testing, the following login and boot paradigm is used:

- Users are logged in every 30 seconds.
- All desktops are pre-booted in advance of logins commencing.

For all testing, all virtual desktops run an industry-standard anti-virus solution (currently McAfee VirusScan Enterprise) in order to fully replicate a customer environment.

6.2.4.1 Login VSI 3 versus Login VSI 4

PAAC on Dell Wyse Datacenter solutions is currently carried out using Login VSI version 4. However, some previous Dell Wyse Datacenter solutions used Login VSI version 3 for this PAAC work. Login VSI version 3 used a slightly different set of workloads to those used by Login VSI 4 and in order to allow comparison of results obtained using these different Login VSI versions, it is useful to be aware of the information presented in the figure below. This information (for Login VSI medium workload) illustrates higher CPU utilization and lower disk IO for Login VSI 4 (green / upper set of graphs) when compared to Login VSI3 (red / lower set of graphs). The exact variation between these Login VSI versions will vary between environments.

6.2.4.2 Login VSI Light Workload

Compared to the Login VSI medium workload described below, the light workload runs fewer applications (mainly Excel and Internet Explorer with some minimal Word activity) and starts/stops the applications less frequently. This results in lower CPU, memory and disk IO usage.

6.2.4.3 Login VSI Medium Workload

The Login VSI medium workload is designed to run on 2vCPU's per desktop VM. This workload emulates a medium knowledge worker using Office, IE, PDF and Java/FreeMind. The Login VSI medium workload has the following characteristics

- Once a session has been started the workload will repeat (loop) every 48 minutes.
- The loop is divided in four segments; each consecutive Login VSI user logon will start a different segment. This ensures that all elements in the workload are equally used throughout the test.
- The medium workload opens up to five applications simultaneously.
- The keyboard type rate is 160 ms for each character.
- Approximately two minutes of idle time is included to simulate real--world users.

Each loop opens and uses:

- 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.

6.2.4.4 Login VSI Heavy Workload

The heavy workload is based on the medium workload except that the heavy workload:

- Begins by opening four instances of Internet Explorer. These instances stay open throughout the workload loop.
- Begins by opening two instances of Adobe Reader. These instances stay open throughout the workload loop.
- There are more PDF printer actions in the workload.
- Instead of 480p videos a 720p and a 1080p video are watched.
- Increased the time the workload plays a flash game.
- The idle time is reduced to two minutes.

6.3 Testing and Validation

6.3.1 Testing Process

The purpose of the single server testing is to validate the architectural assumptions made around the server stack. Each user load is tested against four runs. First, a pilot run to validate that the infrastructure is functioning and valid data can be captured, and then, three subsequent runs allowing correlation of data.

At different stages of the 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.

For all workloads, the performance analysis scenario will be to launch a user session every 30 seconds. Once all users have logged in, all will run workload activities at steady-state for 60 minutes and then logoffs will commence.

6.4 vWorkspace Test Results

6.4.1 Configuration

The following results represent 300 Premium workload users on a 3 node Nutanix cluster. One host was dedicated to Management while the remaining two hosts each delivered 150 Virtual Desktops.

Below is the configuration of the compute hosts used for this testing.

Platform	Config	CPU	Memory	RAID Ctlr	HD Config	Network	Login VSI Workload
XC630	B7	E5-2698v3 (16Core, 2.3GHz)	384GB	H730P	1 X 64 GB SATADOM (CVM/ HV) 2 x 400GB, Intel S3700, SATA SSD's 2.5" (T1) 8 x 1TB NL SAS 2.5" (T2)	2 x 10Gb, 2 x 1Gb SFP+/ BT	Heavy

1GB networking was used for the deployment of the Bensen XC appliances while all 10GB networking was used for the PAAC testing.

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.

- Node3 – XC630– Dedicated Management (vWorkspace Broker, SQL Server & File Services)
- Node2 – XC630 – Dedicated Compute
- Node1 – XC630 – Dedicated Compute

Please refer to Section 5.1 for the configuration of each management virtual machine.

The virtual machines were non-persistent vWorkspace desktops each configured on Windows 8.1 aligning with the Login VSI 4.X virtual machine configuration, Office2010 was used with each Virtual Machine sized at 25GB. User Workload configuration of the load generation virtual machines is shown in the table below:

Role	Config	vCPU	Startup RAM	Dynamic Memory			NIC	OS vDisk	
				Min|Max	Buffer	Weight		Size(GB)	Location
VM	B7	2	1GB	1024MB | 4GB	20%	Med	1	25	SDS: ds_compute

Test Results

Validation was performed using CCC standard testing methodology using LoginVSI 4 load generation tool for VDI benchmarking that simulates production user workloads. Dedupe on the performance Tier was turned on for this testing.

Each test run adhered to PAAC best practices with a 30 second session logon interval, 16 minutes of steady state after which sessions would begin logging off.

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

Hyper-visor	Provisioning	Workload	Density Per Host	Peak CPU %	Avg Memory Consumed	Avg IOPS/ User	Avg Net Kbps/User
Hyper-V	vWorkspace	Professional	155	75%	249 GB	7.45	666

CPU Utilization* CPU % for Hyper-V Hosts was adjusted to account for the fact that on Intel E5-2698v3 series processors the ESX host CPU metrics will exceed the rated 100% for the host if Turbo Boost is enabled (by default). The Adjusted CPU % Usage is based on 100% usage and but is not reflected in the charts. The figure shown in the table is the Compute host steady state peak CPU Usage. One Nutanix CVM Controller virtual machine is located on each node and is allocated 10000MHz of CPU reservation.

Memory utilization. The figure shown in the table above is the average memory **consumed** per Compute host over the recorded test period.

The IOPS results are calculated from the average Nutanix **Cluster Disk IOPS** figure over the test period divided by the number of users.

Network Utilization The figure shown in the table is the average **Bytes/ps/User per Compute host** over the recorded test period.

Results summary

B7 Configuration Hyper-V / vWorkspace User Test, Heavy Workload

The architecture for this solution includes a dedicated management host and two compute hosts. Each of the compute hosts was populated with 70 full virtual machines and one Nutanix CVM per host.

The Nutanix CVM's took up approximately 12% of the compute hosts CPU Usage at the start of the test run.

This chart does not include the additional 10% of CPU available from the Turbo boost feature. With Turbo Boost included an additional 10% CPU is available for us.

The below graph shows the performance data for 150 user sessions per host on a pair of Compute hosts. The CPU reaches a steady state peak of 75% during the test cycle when approximately 150 users are logged on to each compute host.

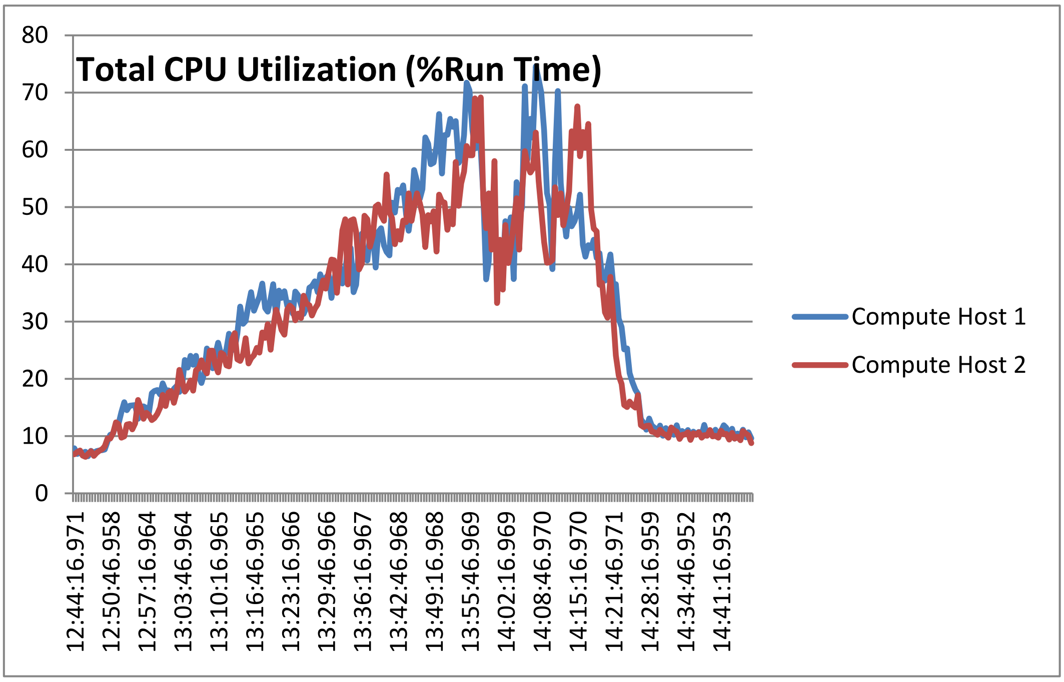

The Management host in the cluster runs the Cluster Shared Volume and vWorkspace management virtual machines and a Nutanix CVM virtual machine. Its CPU utilization is significantly lower than the compute hosts in the cluster. The CPU utilization for the management host does not exceed 10% at any point in the test cycle.

Memory consumption for the cluster, out of a total of 384 GB available memory per node there were no constraints for any of the hosts. The Compute Hosts reached a max memory consumption of 249 GB. This graph illustrates the level of available memory throughout the test cycle.

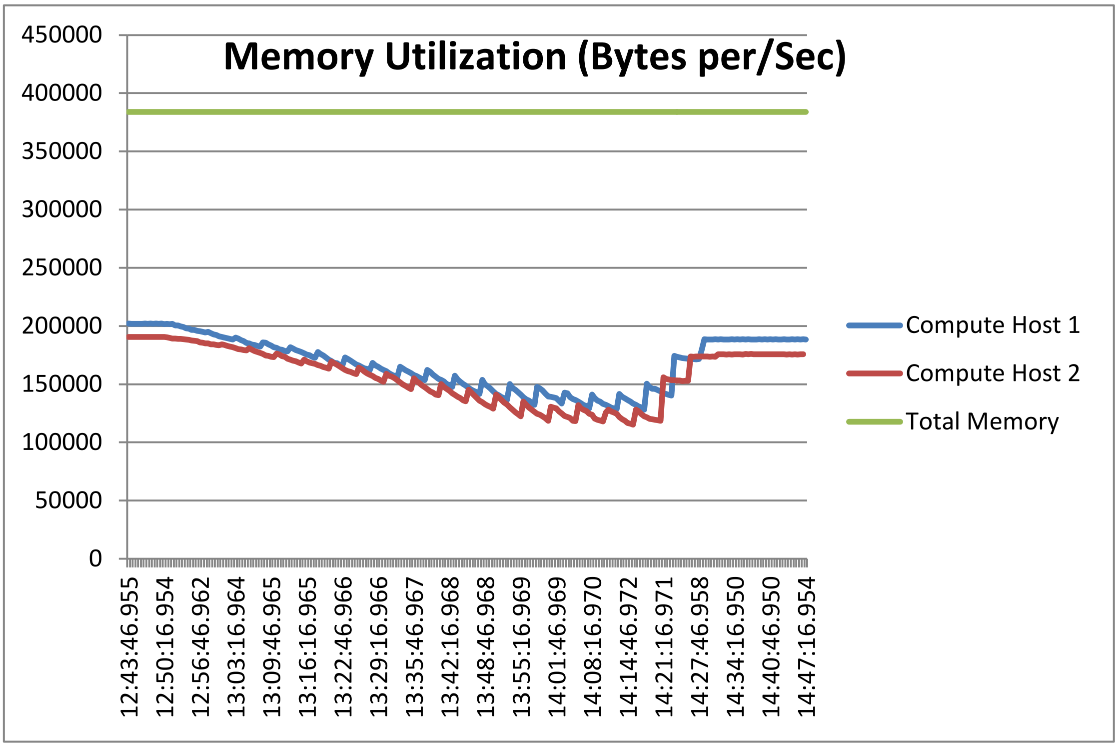

Dual port 10Gb Networking was used for the hosts. The graph below shows the utilization for each compute host in the cluster.

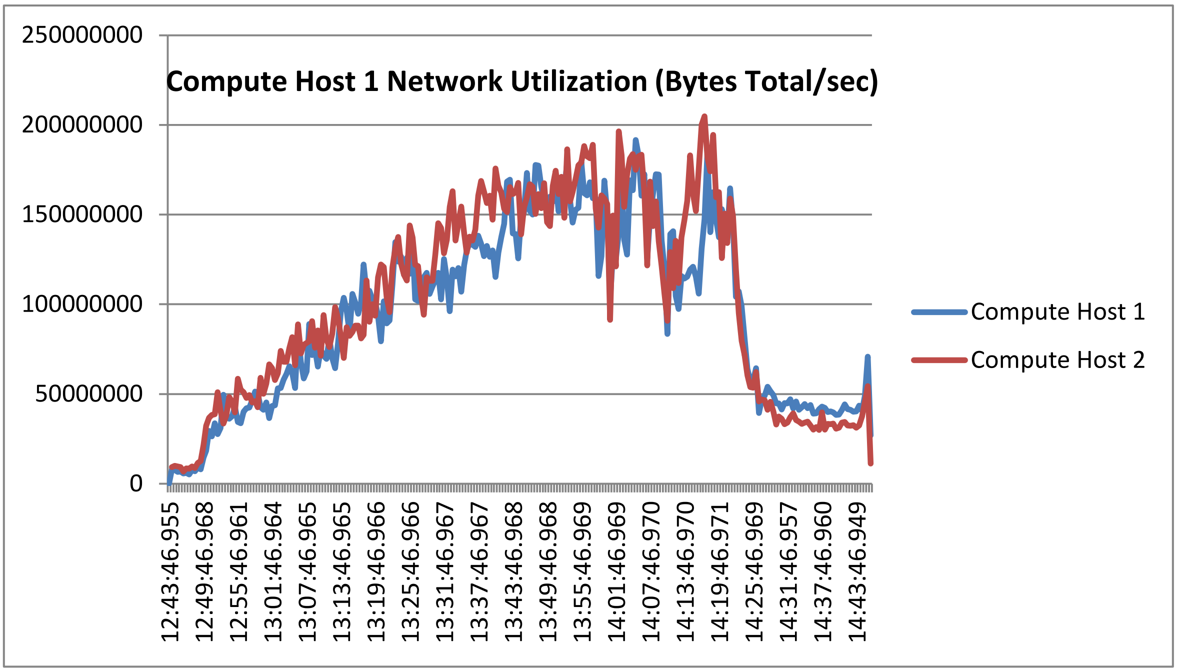

The following graphs show the storage utilization for the Nutanix Cluster. These graphs are taken from Prism. The below graph shows the Nutanix Controller CPU and Memory utilization. Neither exceeded 50% of the allocated resources.

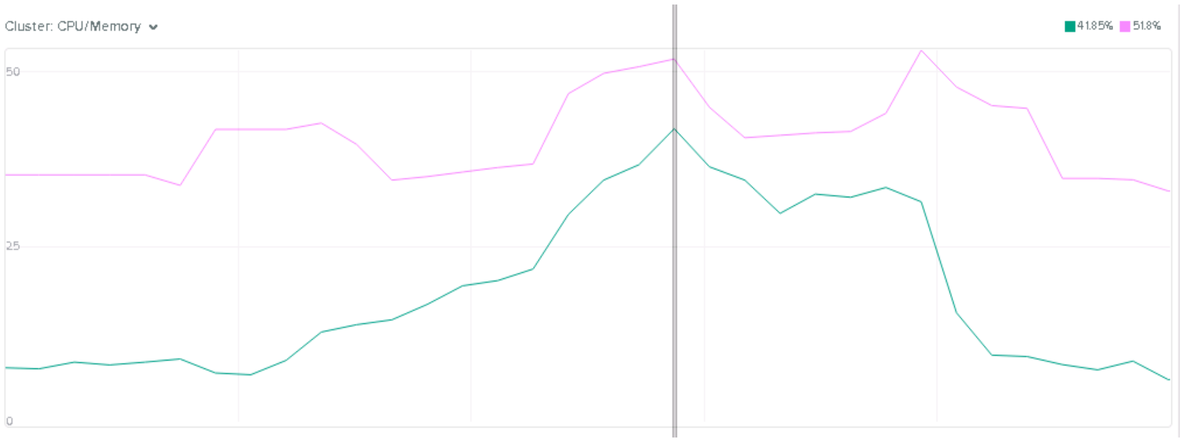

The following graph shows the IOPS recorded by the Nutanix file system at peak of the test. This figure is 2237IOPS.

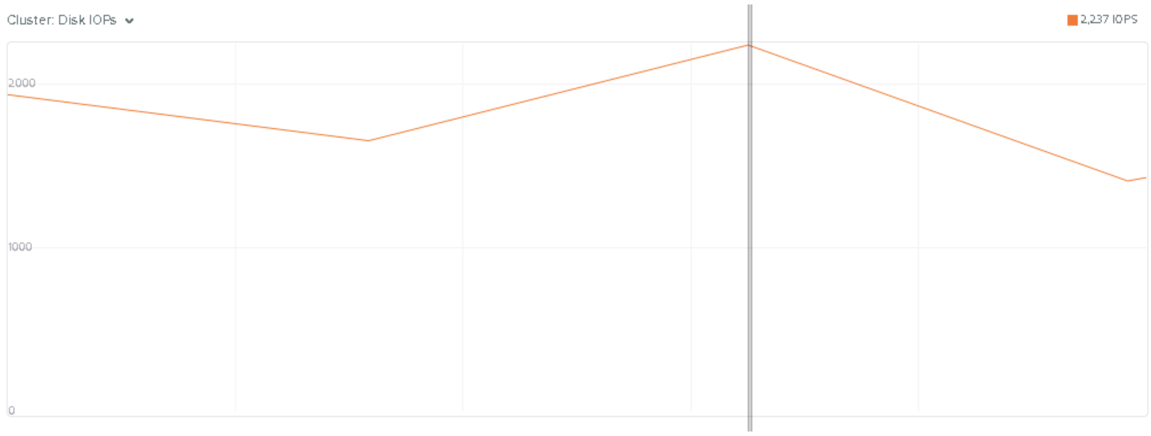

Next we look at the Network utilization at peak of the test. The utilization did not exceed 100Mb/ps.

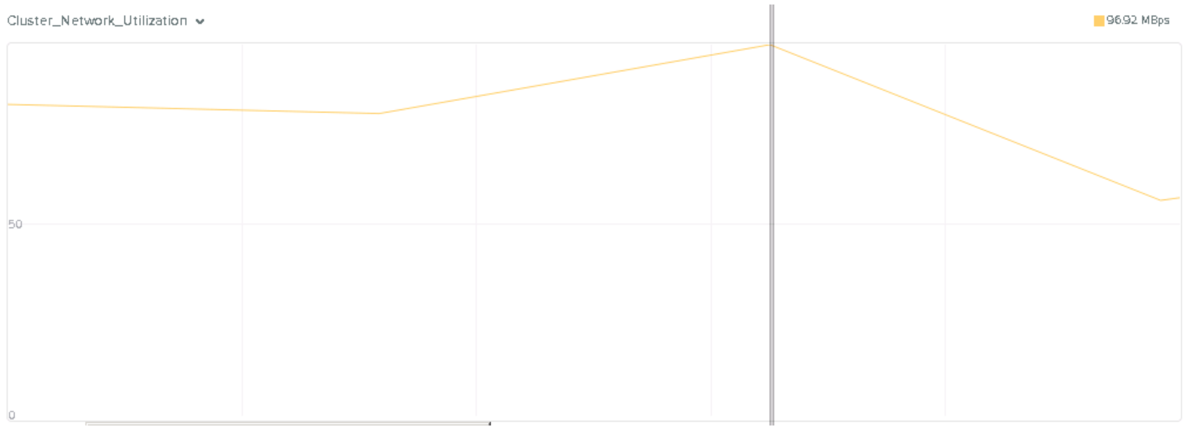

In Summary, the performance of the solution under test was very good. CPU utilization at 75% indicates that additional virtual machines could be added to each compute node. Memory utilization peaked at 249 GB running the premium workload on compute nodes; on the Nutanix cluster the system memory utilization was low at 50% of the available memory. Subjective user test was excellent with no delay or grab experienced. Accessing media content was quick with no delay in the video stream.

6.4.2 RDSH Configuration

The following results represent 500 Standard workload users on a 3 node Nutanix cluster. One host was dedicated to Management while the remaining two hosts each delivered 250 sessions each.

Below is the configuration of the compute hosts used for this testing.

Platform	Config	CPU	Memory	RAID Ctlr	HD Config	Network	Login VSI Workload
XC630	B7	E5-2698v3 (16Core, 2.3GHz)	384GB	H730P	1 X 64 GB SATADOM (CVM/ HV) 2 x 400GB, Intel S3700, SATA SSD's 2.5" (T1) 8 x 1TB NL SAS 2.5" (T2)	2 x 10Gb, 2 x 1Gb SFP+/ BT	Heavy

1GB networking was used for the deployment of the Bensen XC appliances while all 10GB networking was used for the PAAC testing.

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.

- Node3 – XC630– Dedicated Management (vWorkspace Broker, SQL Server & File Services)
- Node2 – XC630 – Dedicated Compute
- Node1 – XC630 – Dedicated Compute

Please refer to Section 5.1 for the configuration of each management virtual machine.

The VM's used were 14 RDSH Session hosts brokered by vWorkspace each configured on Windows Server 2012 R2 aligning with the Login VSI 4.X virtual machine configuration, Office2010 was used. User Workload configuration of the load generation virtual machines is shown in the table below:

Role	Config	vCPU	Startup RAM (GB)	Dynamic Memory			NIC	OS vDisk	
				Min|Max	Buffer	Weight		Size (GB)	Location
RDSH Session Host	B7	8	16GB	16GB | 32GB	20%	Med	1	80	SDS: ds_compute

Test Results

Validation was performed using CCC standard testing methodology using LoginVSI 4 load generation tool for VDI benchmarking that simulates production user workloads. Dedupe on the performance Tier was turned on for this testing.

Each test run adhered to PAAC best practices with a 30 second session logon interval, 16 minutes of steady state after which sessions would begin logging off.

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

Hyper-Visor	Provisioning	Workload	Density Per Host	Peak CPU %	Avg Memory Consumed	Avg IOPS/ User	Avg Net Kbps/User
Hyper-V	vWorkspace Session Hosts	Standard	250	80%	180 GB	2.2	601

CPU Utilization* CPU % for Hyper-V Hosts should be adjusted to account for the fact that on Intel E5-2698v3 series processors the ESX host CPU metrics will exceed the rated 100% for the host if Turbo Boost is enabled (by default). The Adjusted CPU % Usage is based on 100% usage and but is not reflected in the charts. The figure shown in the table is the Compute host steady state peak CPU Usage. One Nutanix CVM Controller virtual machine is located on each node and is allocated 10000MHz of CPU reservation.

Memory utilization. The figure shown in the table above is the average memory **consumed** per Compute host over the recorded test period.

The IOPS results are calculated from the average Nutanix **Cluster Disk IOPS** figure over the test period divided by the number of users.

Network Utilization The figure shown in the table is the average **Bytes/ps/User per Compute host** over the recorded test period.

Results summary

B7 Configuration Hyper-V / vWorkspace User Test, Standard Workload

The architecture for this solution includes a dedicated management host and two compute hosts. Each of the compute hosts was populated with 7 Server 2012 R2 Session hosts and one Nutanix CVM per host.

The Nutanix CVM's took up less than 10% of the compute hosts CPU Usage at the start of the test run.

This chart does not include the additional 10% of CPU available from the Turbo boost feature. With Turbo Boost included an additional 10% CPU is available for us.

The below graph shows the performance data for 250 user sessions per host on a pair of Compute hosts. The CPU reaches a steady state peak of 80% during the test cycle when approximately 250 users are logged on to each compute host.

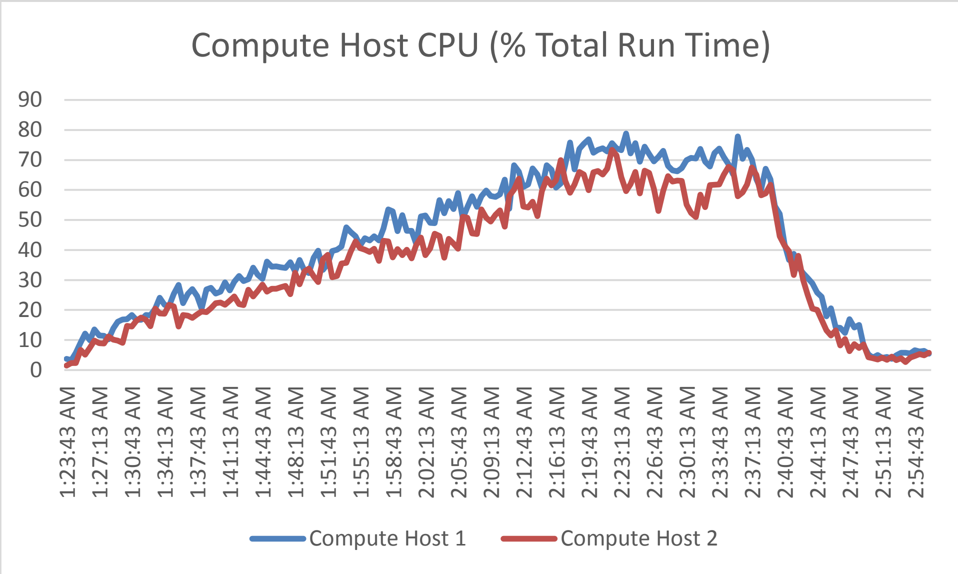

A total of 520 sessions were launched to log in 500 users. 486 sessions went active during the test and 482 sessions logged out.

The Management host in the cluster runs the Cluster Shared Volume and vWorkspace management virtual machines and a Nutanix CVM virtual machine. Its CPU utilization is significantly lower than the compute hosts in the cluster. The CPU utilization for the management host does not exceed 10% at any point in the test cycle.

Memory consumption for the cluster, out of a total of 384 GB available memory per node there were no constraints for any of the hosts. Each host consumed circa 136GB of physical memory on average. Peak consumption during the log off period was 180GB.

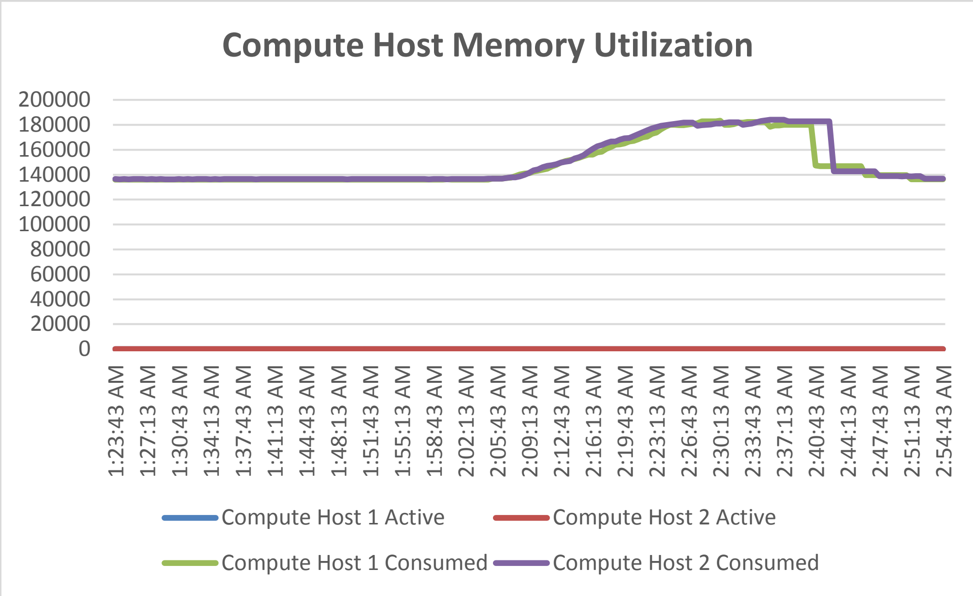

Dual port 10Gb Networking was used for the hosts. The below graph shows the utilization for each compute host in the cluster.

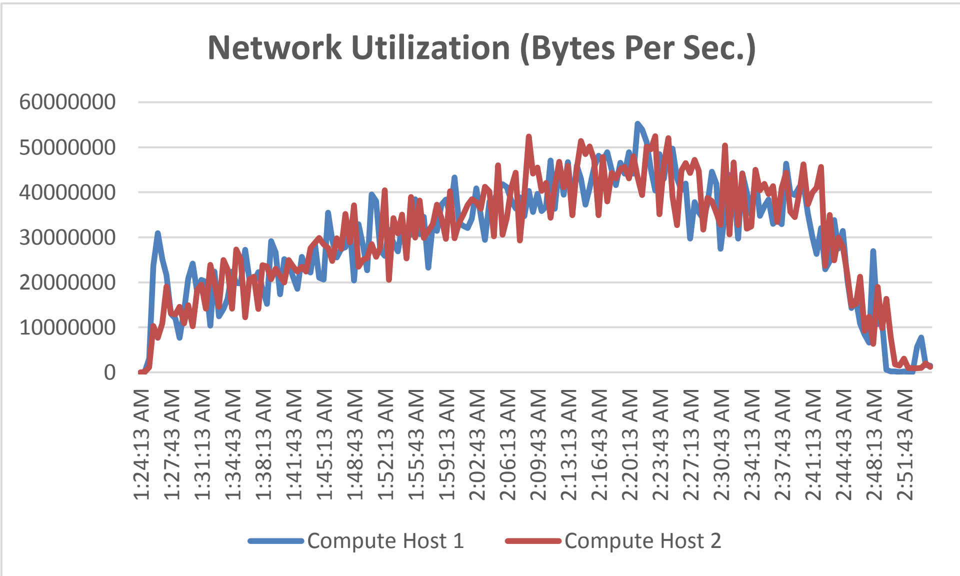

In addition to the compute host metrics, Memory, CPU & IOPS information was collected for the 3 node Nutanix cluster. The following chart illustrates the CPU and Memory utilization during the test period. Neither Memory nor CPU for the Nutanix Cluster exceeded 24& 27% respectively.

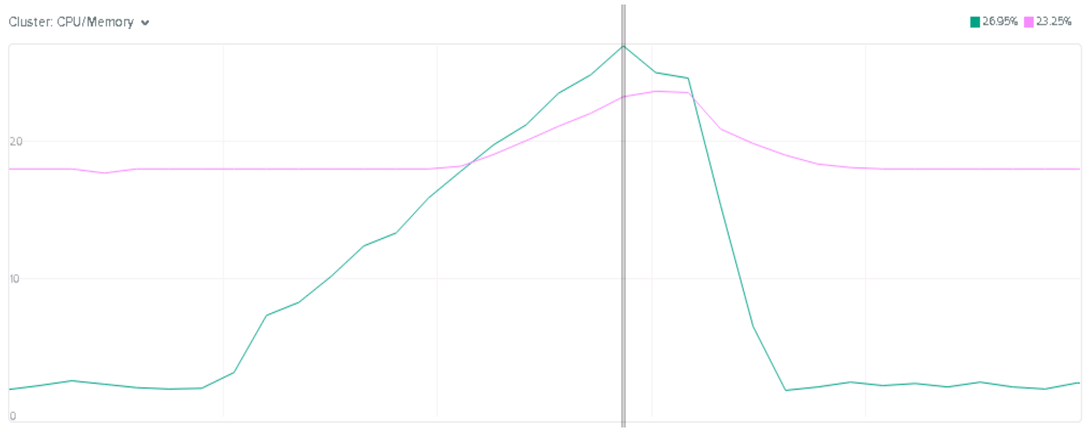

The following graph shows the IOPS recorded by the Nutanix file system at peak of the test. This figure is 1081.

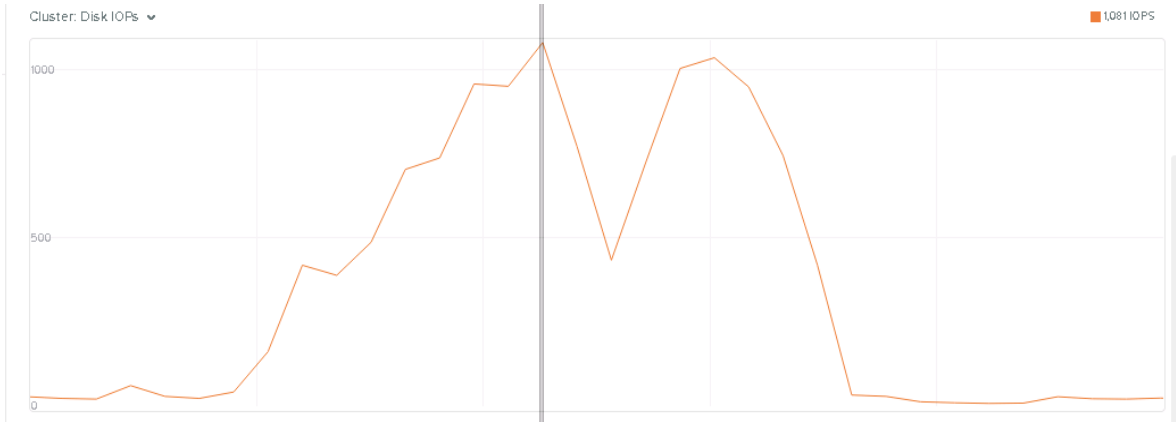

Next we look at the Network utilization which accounted for less than 50% of the available bandwidth. This was not a bottleneck at any stage of the testing period.

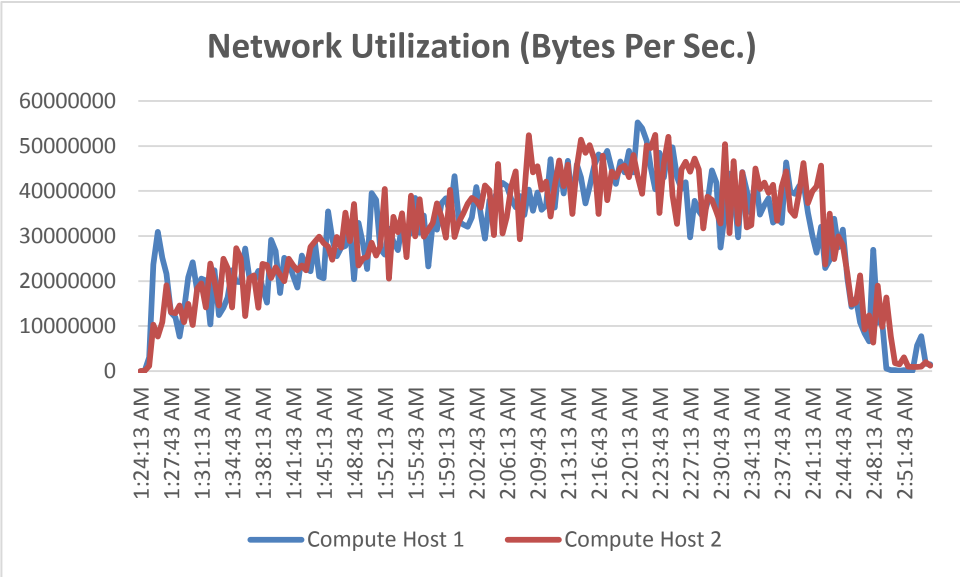

In summary both Memory and CPU stats showed that additional sessions could be hosted per compute node. Please see the single node investigation that follows running 300 RDSH sessions on a single node.

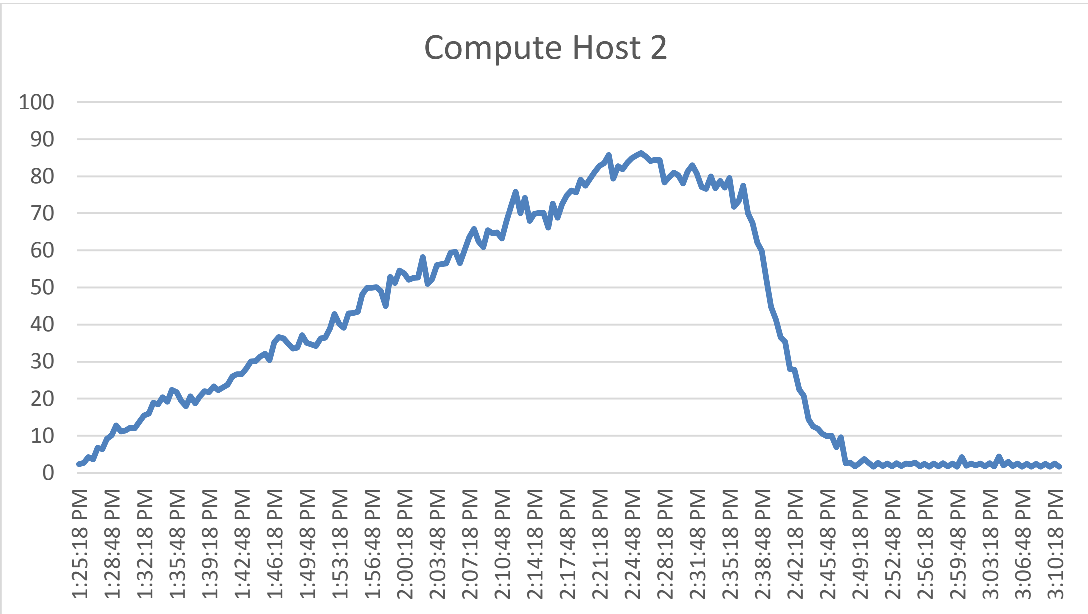

CPU reached 90% with 287 out of 300 sessions logging in. Based on the CPU utilization a maximum of 300 sessions should be run to allow some system over head.

Below is the memory consumption for this host also. A peak of 226GB was consumed during steady state of this test.

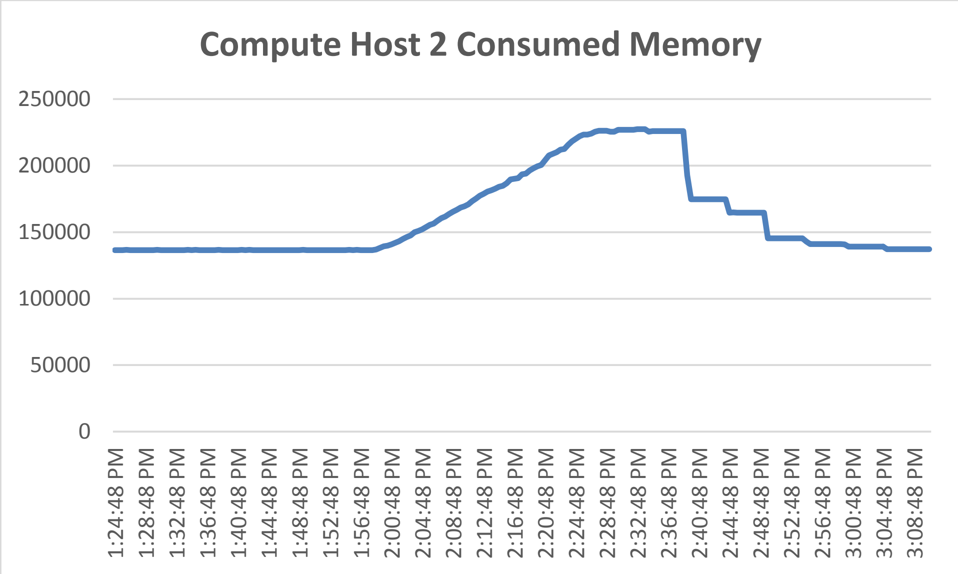

Finally, we look at Network Utilization which showed no strain on the network at any time.

Analysis of the Nutanix File system showed 50% network utilization while both CPU and Memory consumption were low as expected.

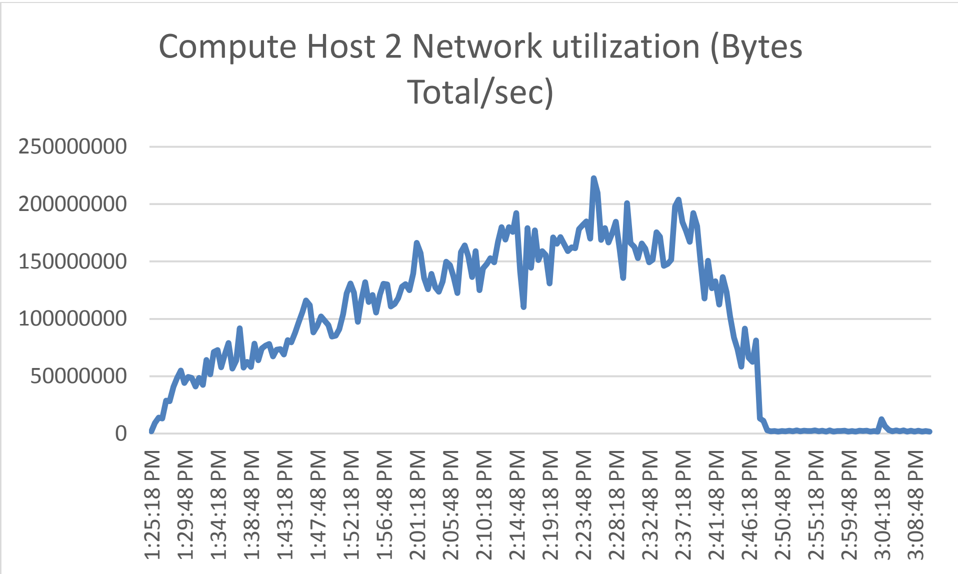

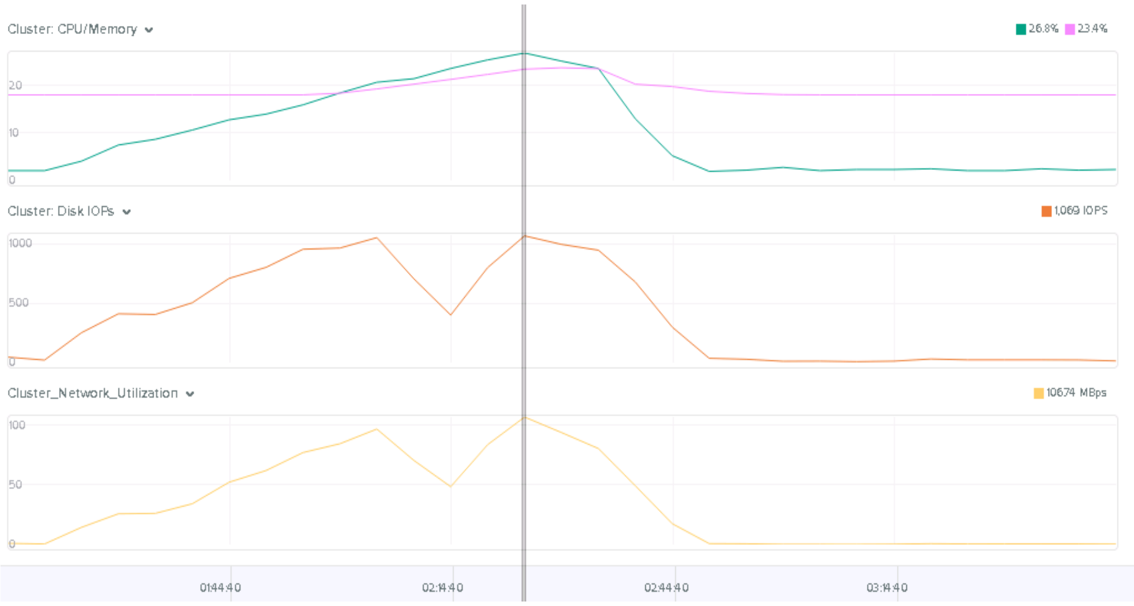

In summary, the single node running 300 sessions showed that Network utilization was high on the Nutanix cluster but did not impact the user performance. User experience was good for each session and a user subjective test showed good performance.

Acknowledgements

Thanks to Dwayne Lessner and Steven Poitras of the Nutanix Engineering team for the technical detail presented in section 2.2 of this document.

About the Authors

Peter Fine is the Sr. Principal Solutions Architect for enterprise VDI solutions at Dell. Peter has extensive experience and expertise on the broader Microsoft, Citrix and VMware solutions software stacks as well as in enterprise virtualization, storage, networking and enterprise data center design.

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.

Paul Wynne is a Principal Engineer within the Cloud Client Computing team at Dell. Paul has over 10 years' experience architecting and supporting VMware, Microsoft & Citrix technologies on a range of Enterprise Storage, Switching and Server technologies.

Andrew Breedy is a staff engineer working in the Dell-Wyse Datacenter engineering group. Andrew has extensive experience in server hardware and on Microsoft and VMWare virtualization solutions as well as a background in engineering test processes and performance analysis and characterization.

Manish Chacko is the Sr. Technical Marketing Advisor for Appliance and Citrix-based solutions at Dell. Before writing about technology, Manish spent time designing, implementing and supporting technology- in IT, Systems Engineering & Network Performance & Monitoring. Manish was a long-time Dell customer & Advocate before becoming a Dell employee.