Dell XC Web Scale Appliance Architecture for Citrix

Dell Wyse Technical Marketing

October 2014

A Dell Appliance Architecture

Revisions

Date	Description
October 2014	Initial release
November 2014	Final release

Contents

1 Introduction

1.1 Purpose

This document addresses the architecture design, configuration and implementation considerations for the key components required to deliver full-clone persistent virtual desktops via Citrix® XenDesktop® and XenApp® on Microsoft® Windows Server® Hyper-V® 2012 R2 or VMware® vSphere® 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 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.

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.

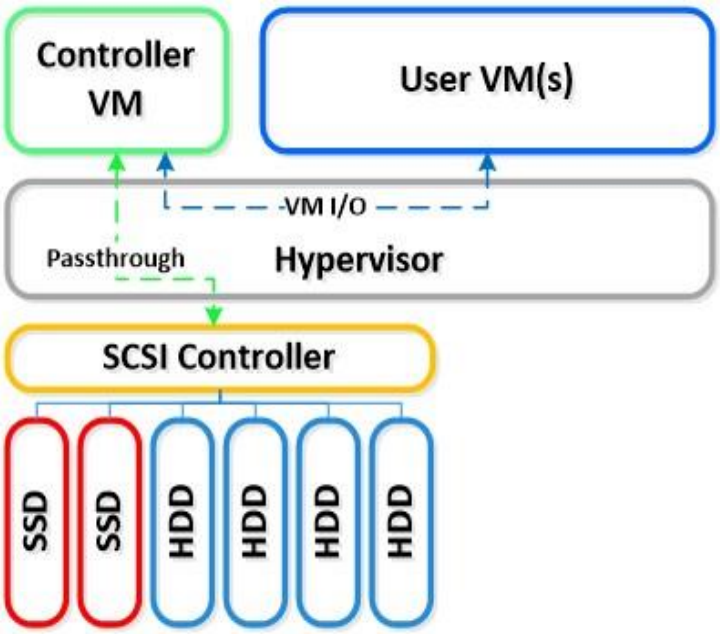

In addition, local storage from all nodes is virtualized into a unified pool by the Nutanix Distributed File System (NDFS). In effect, NDFS acts like an advanced NAS that uses local SSDs and disks from all nodes to store virtual machine data. Virtual machines running on the cluster write data to NDFS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

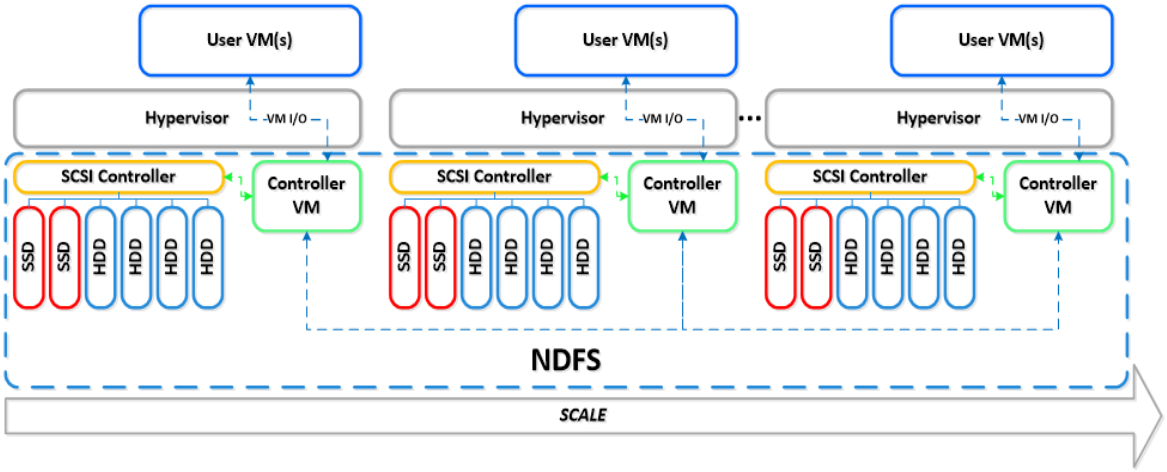

NDFS understand the concept of a virtual machine and provides advanced data management features. It brings data closer to virtual machines by storing the data locally on the system, resulting in higher performance at a lower cost. Nutanix platforms can horizontally scale from as few as three nodes to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

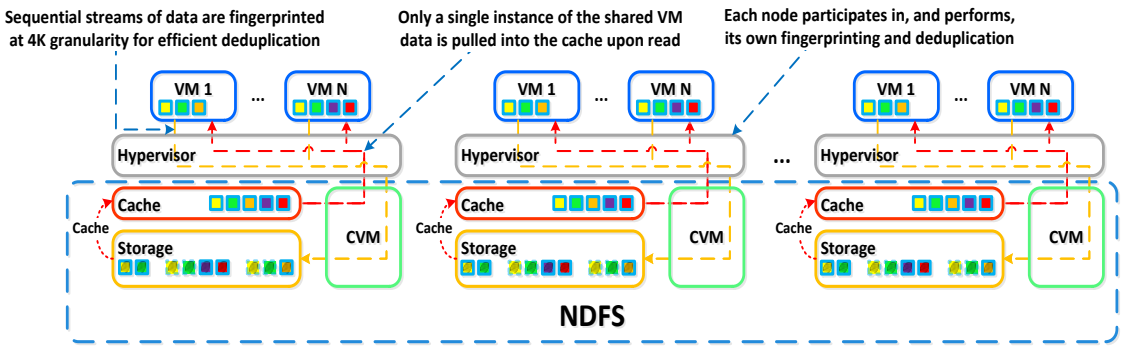

Nutanix Shadow Clones delivers distributed localized caching of virtual disks performance in multi-reader scenarios, such as desktop virtualization using Citrix XenDesktop or XenApp. With Shadow Clones, the CVM actively monitors virtual disk access trends. If there are requests originating from more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

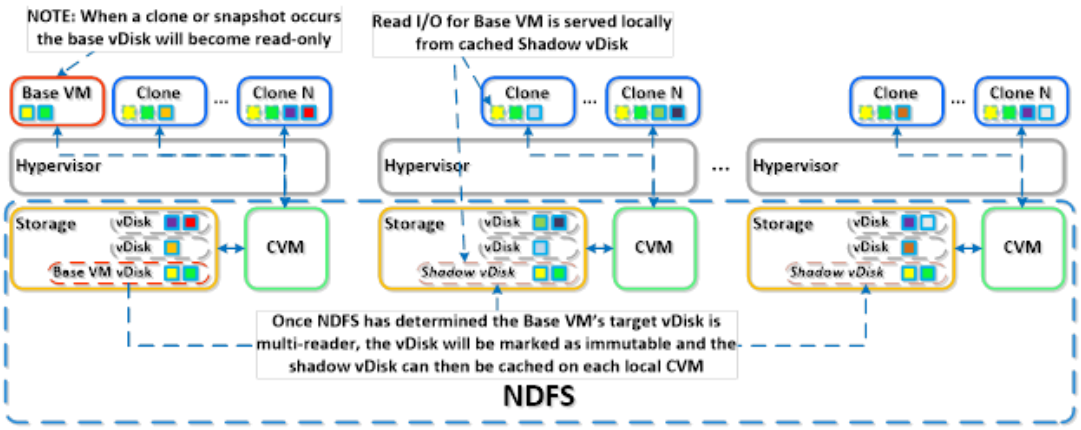

The benefits of the Nutanix Platform are now exposed to scale out vSphere or Hyper-v deployments:

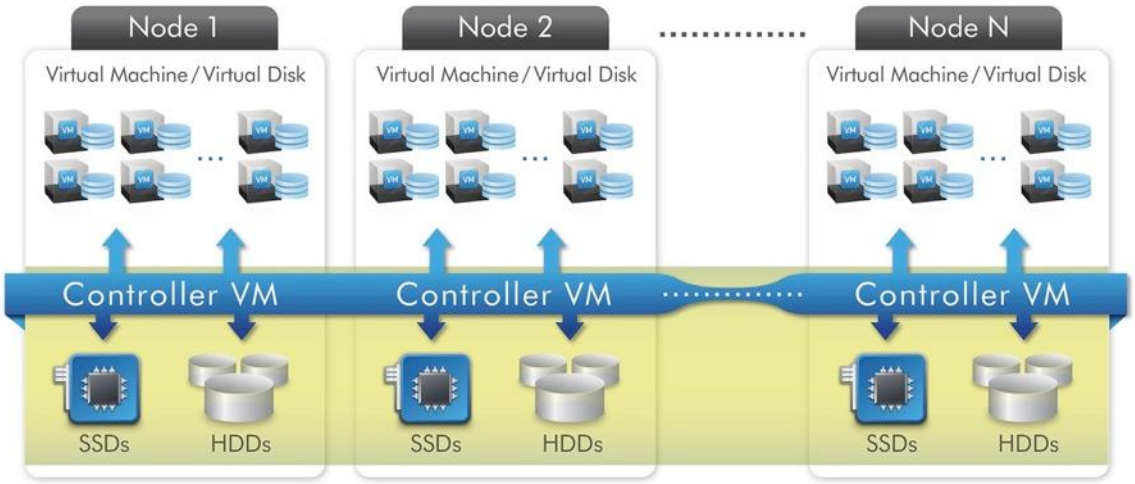

Nutanix Web-scale Converged Infrastructure

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:

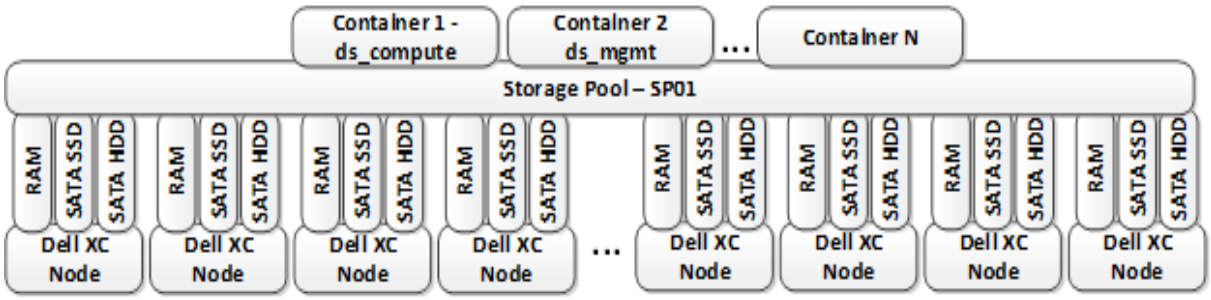

Dell XC Web Scale allows organizations to deliver virtualized or remote desktops and applications through a single platform and support end users with access to all of their desktops and applications in a single place.

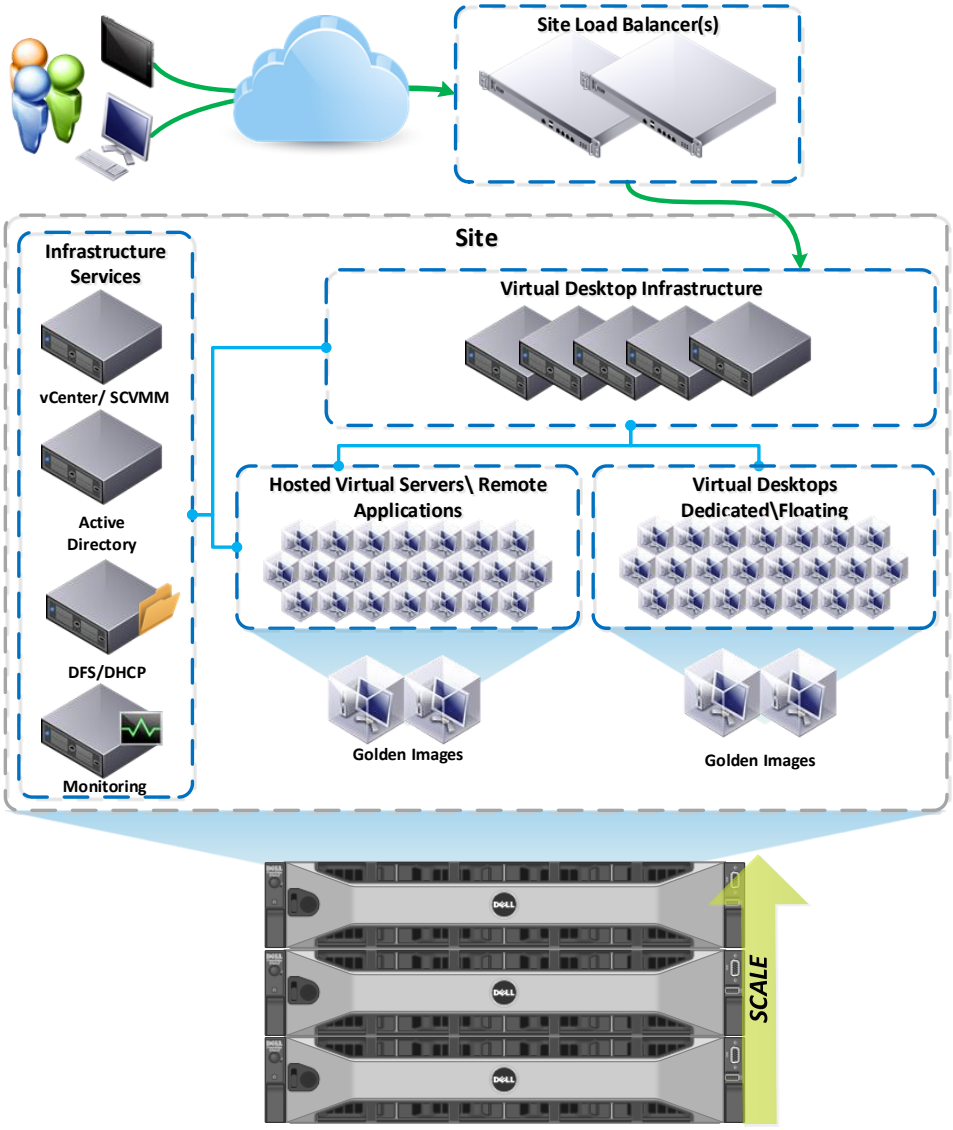

2.3 Dell XC Web Scale – Solution Pods

The networking layer consists of the 10Gb Dell Force10 S4810 utilized to build a world-class leaf/spine architecture with robust 1Gb switching in the S60 for iDRAC connectivity.

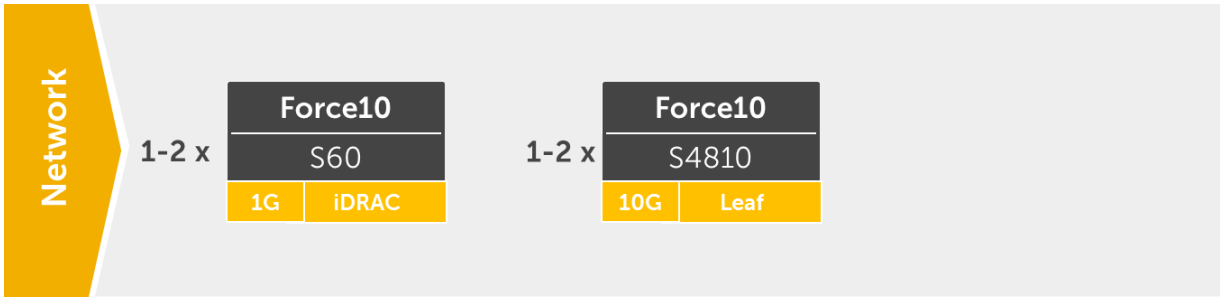

The compute, management and storage layers are 'converged' into a single server XC Series appliance, hosting either VMware vSphere or Microsoft Hyper-V hypervisors. The recommended boundaries of an individual pod are based on number of nodes supported within a given hypervisor cluster, 32 nodes for vSphere or 64 nodes for Hyper-V.

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 interchangeably for XenApp or 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 clusters adding disk capacity and performance for each distinct cluster.

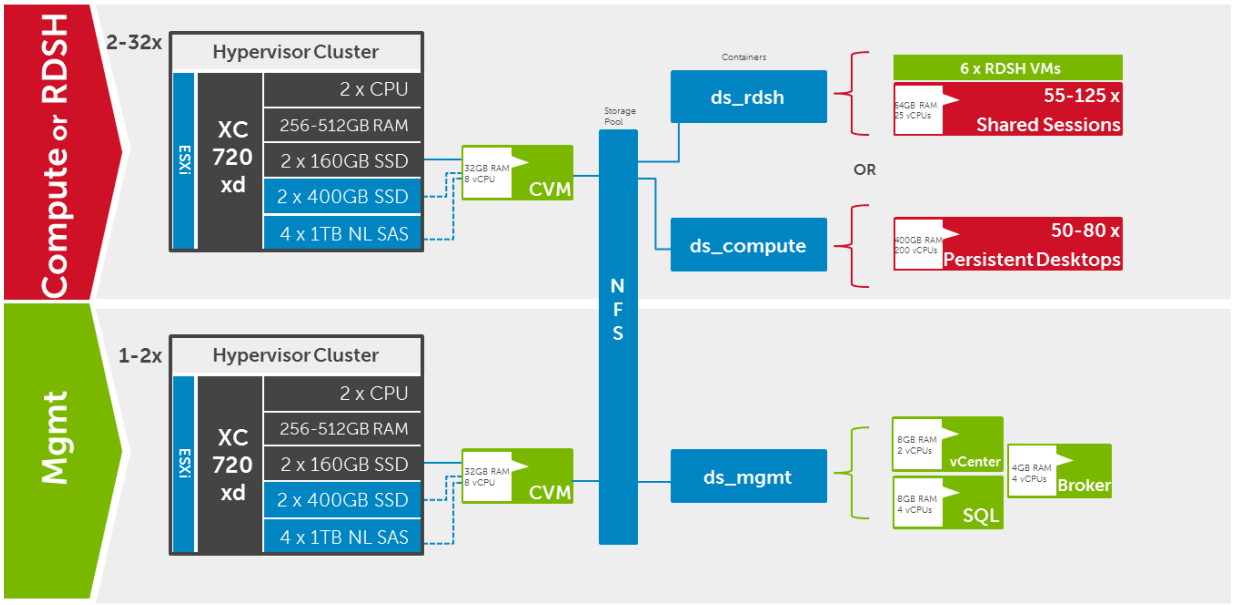

The Hyper-V solution pod scales a bit higher due to the upper node limit of 64 for a Hyper-V cluster and alternatively uses SMB as the file access protocol. The hardware platforms, logical containers and basic architecture are the same.

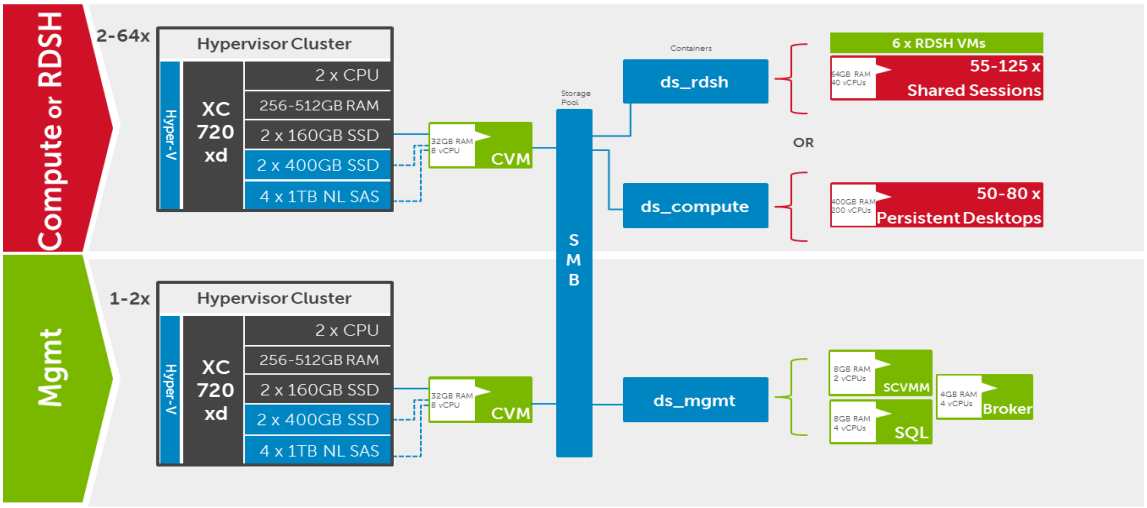

2.3.1 Network Architecture

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 providing scale from 3 XC nodes to thousands without any impact to available bandwidth:

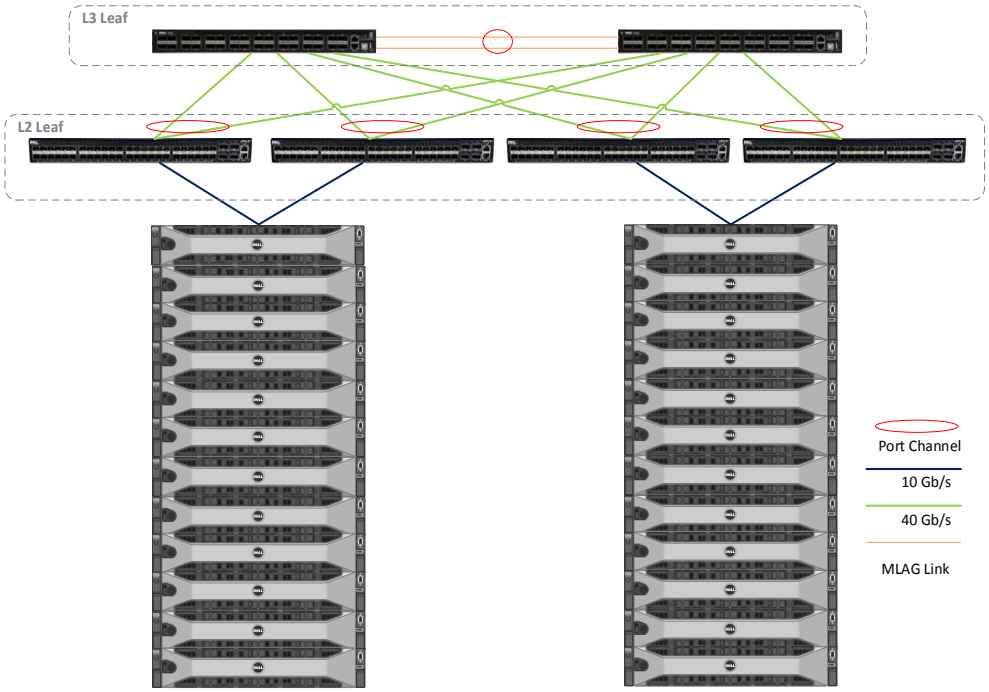

3 Hardware Components

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 Force10 S60 (1Gb ToR Switch)

The Dell Force10 S-Series S60 is a high-performance 1/10Gb 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
Force10 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 the support copper or fiber	
		12Gb or 24Gb stacking (up to 12 switches)	
		2 x modular slots for 10Gb uplinks or stacking modules	

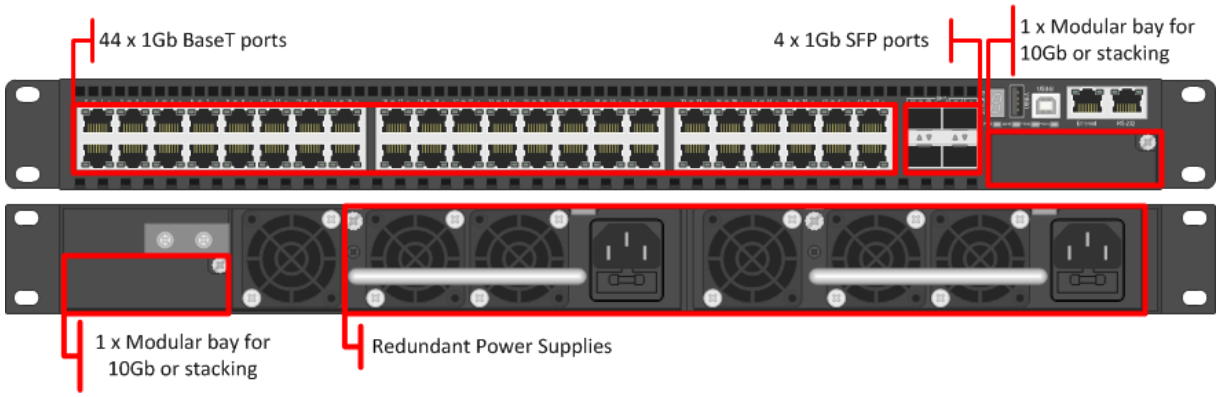

Guidance:

- 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 Force10 networking, 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.

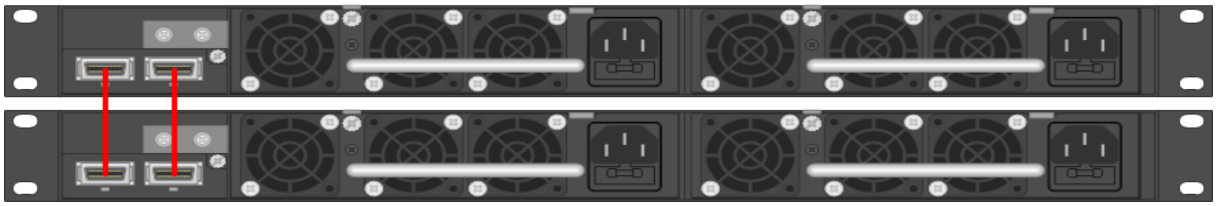

3.1.2 Force10 S4810 (10Gb ToR Switch)

The Dell Force10 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
Force10 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	

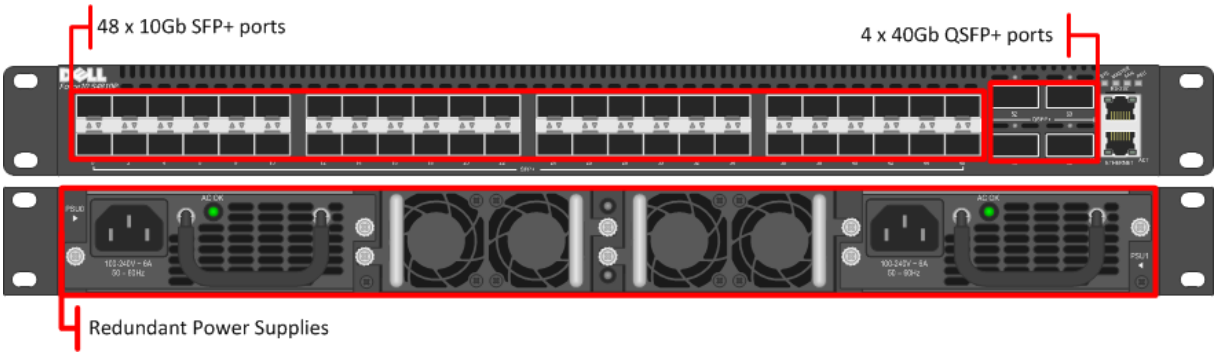

Guidance:

- 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 Force10 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.

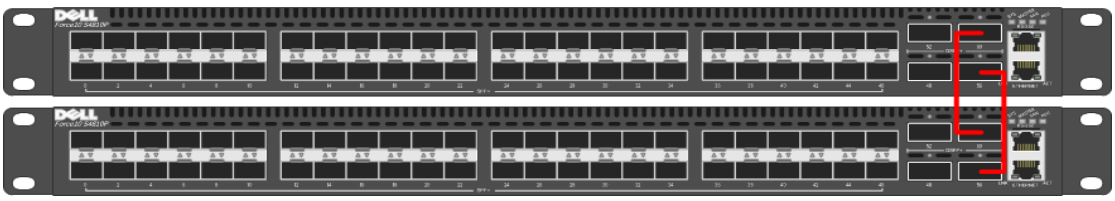

3.2 PowerEdge XC720xd

The rack server platform for the Dell Wyse Datacenter solution is the best-in-class Dell PowerEdge XCR720xd (12G). This dual socket CPU platform runs the fastest Intel Xeon E5-2600 family of processors, can host up to 768GB RAM, and supports up to 12 x 3.5" SAS disks. The Dell PowerEdge XCR720xd offers uncompromising performance and scalability in a 2U form factor. For more information please visit: Link

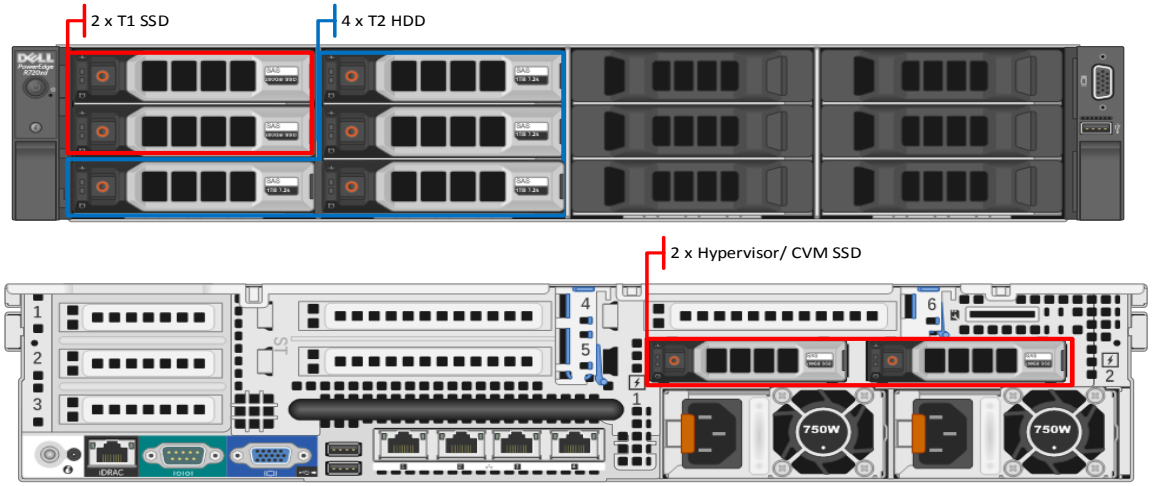

3.2.1 Dell XC Entry Level

The Dell CX entry level platform is perfect for POCs, lighter user workloads, shared sessions or application virtualization. Each appliance comes equipped with dual 6-core CPUs and 256GB of high-performance RAM. The H310 PERC connects to the 2 x 160GB SSDs in the rear flexbay housing the hypervisor and Nutanix Contoller VM. Six additional disks come in each host, 2 x 200GB SSD for the hot tier (Tier1) and 4 x 1TB NL SAS disks for the cold tier (Tier2). 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.

Entry Level – Dell XC720xd	
CPU	2 x E5-2620v2 (6C, 2.1GHz)
Memory	16 x 16GB 1866MHz RDIMMs **Effective speed: 1866MHz @ 256GB**
Disk Ctrls	H310 – RAID1 – Rear Flexbay LSI9207i – no RAID – front drives
Storage	2 x 160GB SSD 2.5" Rear **(CVM/ HV)** 2 x 200GB MLC SATA SSD 3.5" (T1) 4 x 1TB NL SAS 3.5" (T2)
Network	2 x 10Gb, 2 x 1Gb SFP+/ BT
iDRAC	iDRAC7 Ent w/ vFlash, 8GB SD
Power	2 x 750W PSUs

3.2.2 Dell XC Mid-Level

The Dell CX mid-level platform is perfect for larger POCs, medium user workloads, shared sessions or application virtualization. Each appliance comes equipped with dual 8-core CPUs and 384GB of high-performance RAM. The H310 PERC connects to the 2 x 160GB SSDs in the rear flexbay housing the hypervisor and Nutanix Contoller VM. Six additional disks come in each host, 2 x 400GB SSD for the hot tier (Tier1) and 4 x 1TB NL SAS disks for the cold tier (Tier2). 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.

Mid-Level – Dell XC720xd	
CPU	2 x E5-2650v2 (8C, 2.6GHz)
Memory	24 x 16GB 1866MHz RDIMMs **Effective speed: 1333MHz @ 384GB**
Disk Ctrls	H310 – RAID1 – Rear Flexbay LSI9207i – no RAID – front drives
Storage	2 x 160GB SSD 2.5" Rear **(CVM/ HV)** 2 x 400GB MLC SATA SSD 3.5" (T1) 4 x 1TB NL SAS 3.5" (T2)
Network	2 x 10Gb, 2 x 1Gb SFP+/ BT
iDRAC	iDRAC7 Ent w/ vFlash, 8GB SD
Power	2 x 750W PSUs

3.2.3 Dell XC High Level

The Dell CX high level platform is ideal for high performance requirements, heavy user workloads, shared sessions or application virtualization. Each appliance comes equipped with dual 8-core CPUs and 512GB of high-performance RAM. The H310 PERC connects to the 2 x 160GB SSDs in the rear flexbay housing the hypervisor and Nutanix Contoller VM. Six additional disks come in each host, 2 x 400GB SSD for the hot tier (Tier1) and 4 x 1TB NL SAS disks for the cold tier (Tier2). 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.

High Level – Dell XC720xd	
CPU	2 x E5-2680v2 (10C, 2.8GHz)
Memory	16 x 32GB 1866MHz LRDIMMs **Effective speed: 1600MHz @ 512GB**
Disk Ctrls	H310 – RAID1 – Rear Flexbay LSI9207i – no RAID – front drives
Storage	2 x 160GB SSD 2.5" Rear **(CVM/ HV)** 2 x 400GB MLC SATA SSD 3.5" (T1) 4 x 1TB NL SAS 3.5" (T2)
Network	2 x 10Gb, 2 x 1Gb SFP+/ BT
iDRAC	iDRAC7 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 Citrix user experience and are the recommended choices for this solution.

3.3.1 ThinOS – T10D

T10D sets the standard for thin clients. Providing an exceptional user experience, the T10D features the incredibly fast Dell Wyse ThinOS, for environments in which security is critical—there's no attack surface to put your data at risk. Boot up in just seconds and log in securely to almost any network. The T10D delivers a superior Citrix VDI user experience, along with usability and management features found in premium thin clients. The T10D delivers outstanding performance based on its dual core system-on-a-chip (SoC) design, and a built-in media processor delivers smooth multimedia, bi-directional audio and Flash playback. Flexible mounting options let you position the T10D vertically or horizontally on your desk, on the wall or behind your display. Using about 7-watts of power in full operation, the T10D creates very little heat for a greener, more comfortable working environment. Link

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. For more information, please visit: Link

3.3.3 Windows Embedded 7 – Z90Q7

The Dell Wyse Z90Q7 is a super high-performance Windows Embedded Standard 7 thin client for virtual desktop environments. Featuring a quad-core AMD processor, and an integrated graphics engine that significantly boost performance; the Z90Q7 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 the Z90Q7 does the rest automatically—no need to reboot. Scale to tens of thousands of endpoints with Dell Wyse WDM software or leverage your existing Microsoft System Center Configuration Manager platform. The Z90Q7 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

3.3.4 Windows Embedded 8 – Z90Q8

Dell Wyse Z90Q8 is a super high-performance Windows Embedded 8 Standard thin client for virtual desktop environments. Featuring a quad-core AMD processor, the Z90Q8 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. And you can scale to tens of thousands of Z90Q8 endpoints with Dell Wyse Device Manager (WDM) software, or leverage your existing Microsoft System Center Configuration Manager platform. With single-touch or multi-touch capable displays, the Z90Q8 adds the ease of an intuitive touch user experience. The Z90Q8 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

3.3.5 Suse Linux – Z50D

Designed for power users, the new Dell Wyse Z50D is the highest performing thin client on the market. Highly secure and ultra-powerful, the

Z50D combines Dell Wyse-enhanced SUSE Linux Enterprise with a dual-core AMD 1.65 GHz processor and a revolutionary unified engine for an unprecedented user experience. The Z50D 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

3.3.6 Dell Wyse Zero – Xenith 2

Establishing a new price/performance standard for zero clients for Citrix, the new Dell Wyse Xenith 2 provides an exceptional user experience at a highly affordable price for Citrix XenDesktop and XenApp environments. With a zero attack surface, the ultra-secure Xenith 2 offers network-borne viruses and malware no target for attacks. Xenith 2 boots up in just seconds and delivers exceptional performance for Citrix XenDesktop and XenApp users while offering usability and management features found in premium Dell Wyse cloud client devices. Xenith 2 delivers outstanding performance based on its system-on-chip (SoC) design optimized with its Dell Wyse zero architecture and a built-in media processor delivers smooth multimedia, bi-directional audio and Flash playback. Flexible mounting options let you position Xenith 2 vertically or horizontally. Using about 7 Watts of power in full operation, the Xenith 2 creates very little heat for a greener working environment. For more information, please visit: Link

3.3.7 Dell Wyse Zero – Xenith Pro 2

Dell Wyse Xenith Pro 2 is the next-generation zero client for Citrix HDX and Citrix XenDesktop, delivering ultimate performance, security and simplicity. With a powerful dual core AMD G-series CPU; Xenith Pro 2 is faster than competing devices. This additional computing horsepower allows dazzling HD multimedia delivery without overtaxing your server or network. Scalable enterprise-wide management provides simple deployment, patching and updates—your Citrix XenDesktop server configures it out-of-the-box to your preferences for plug-and-play speed and ease of use. Virus and malware immune, the Xenith Pro 2 draws under 9 watts of power in full operation—that's less than any PC on the planet. For more information please visit: Link

3.3.8 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

3.3.9 Dell Venue 11 Pro

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

3.3.10 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 Citrix XenDesktop

The solution is based on Citrix XenDesktop 7.1 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.

Citrix XenDesktop 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.

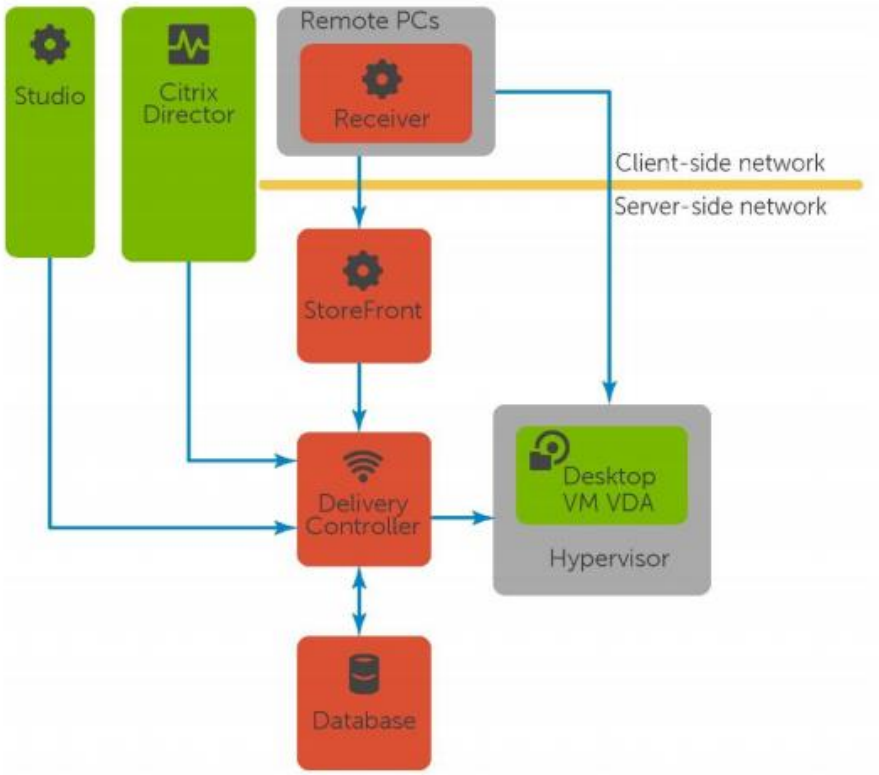

The core XenDesktop components include:

- **Studio**

  Studio is the management console that enables you to configure and manage your deployment, eliminating the need for separate management consoles for managing delivery of applications and desktops. Studio provides various wizards to guide you through the process of setting up your environment, creating your workloads to host applications and desktops, and assigning applications and desktops to users.

- **Director**

  Director is a web-based tool that enables IT support teams to monitor an environment, troubleshoot issues before they become system-critical, and perform support tasks for end users. You can also view and interact with a user's sessions using Microsoft Remote Assistance.

- **Receiver**

  Installed on user devices, Citrix Receiver provides users with quick, secure, self-service access to documents, applications, and desktops from any of the user's devices including smartphones, tablets, and PCs. Receiver provides on-demand access to Windows, Web, and Software as a Service (SaaS) applications.

- **Delivery Controller** (DC)

  Installed on servers in the data center, the controller authenticates users, manages the assembly of users' virtual desktop environments, and brokers connections between users and their virtual desktops.

- **StoreFront**

  StoreFront authenticates users to sites hosting resources and manages stores of desktops and applications that user's access.

- **License Server**

  The Citrix License Server is an essential component at any Citrix-based solution. Every Citrix product environment must have at least one shared or dedicated license server. License servers are computers that are either partly or completely dedicated to storing and managing licenses. Citrix products request licenses from a license server when users attempt to connect.

- **Machine Creation Services** (MCS)

  A collection of services that work together to create virtual servers and desktops from a master image on demand; optimizing storage utilization and providing a pristine virtual machine to users every time they log on. Machine Creation Services is fully integrated and administrated in Citrix Studio.

- **Provisioning Services** (PVS)

  The Provisioning Services infrastructure is based on software-streaming technology. This technology allows computers to be provisioned and re-provisioned in real-time from a single shared-disk image.

- **Virtual Delivery Agent** (VDA)

  The Virtual Desktop Agent is a transparent plugin that is installed on every virtual desktop or XenApp host (RDSH) and enables the direct connection between the virtual desktop and users' endpoint devices.

4.1.1 Machine Creation Services (MCS)

Citrix Machine Creation Services is the native provisioning mechanism within Citrix XenDesktop for virtual desktop image creation and management. Machine Creation Services uses the hypervisor APIs to create, start, stop, and delete virtual desktop images. Desktop images are organized in a

Machine Catalog and within that catalog there are a number of options available to create and deploy virtual desktops:

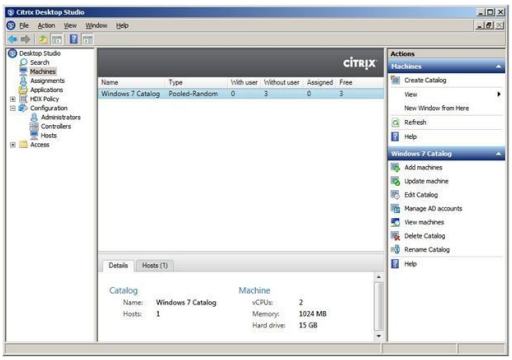

- **Random**: Virtual desktops are assigned randomly as users connect. When they logoff, the desktop is reset to its original state and made free for another user to login and use. Any changes made by the user are discarded at log off.
- **Static**: Virtual desktops are assigned to the same user every time with three options for how to handle changes made to the desktop: Store on local vDisk, Personal vDisk, or discarded on user log off.

All the desktops in a random or static catalog are based off a master desktop template which is selected during the catalog creation process. MCS then takes snapshots of the master template and layers two additional virtual disks on top: an Identity vDisk and a Difference vDisk. The Identity vDisk includes all the specific desktop identity information such as host names and passwords. The Difference vDisk is where all the writes and changes to the desktop are stored. These Identity and Difference vDisks for each desktop are stored on the same data store as their related clone.

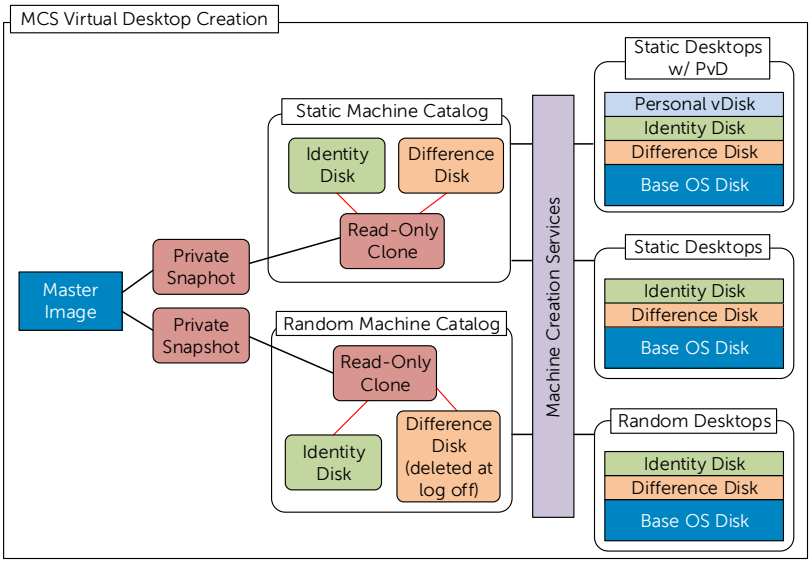

While traditionally used for small to medium sized XenDesktop deployments, MCS can bring along with it some substantial Tier 1 storage cost savings because of the snapshot/identity/difference disk methodology. The Tier 1 disk space requirements of the identity and difference disks when layered on top of a master image snapshot, is far less than that of a dedicated desktop architecture.

4.1.2 Citrix Profile Manager

Citrix Profile Management is a component of the XenDesktop suite which is used to manage user profiles and minimize many of the issues associated with traditional Windows roaming profiles in an environment where users may have their user profile open on multiple devices at the same time. The profile management toolset has two components: the profile management agent, installed on any device where the user profiles is managed, and a Group Policy Administrative Template, which is imported to a group policy.

In order to further optimize, the profile management folders within the user profile is redirected the users' home drive. The folder redirection is managed via group policy objects within Active Directory. The following folders are redirected:

- Contacts
- Downloads
- Favorites
- Links
- My Documents
- Searches
- Start Menu
- Windows
- My Music
- My Pictures
- My Videos
- Desktop

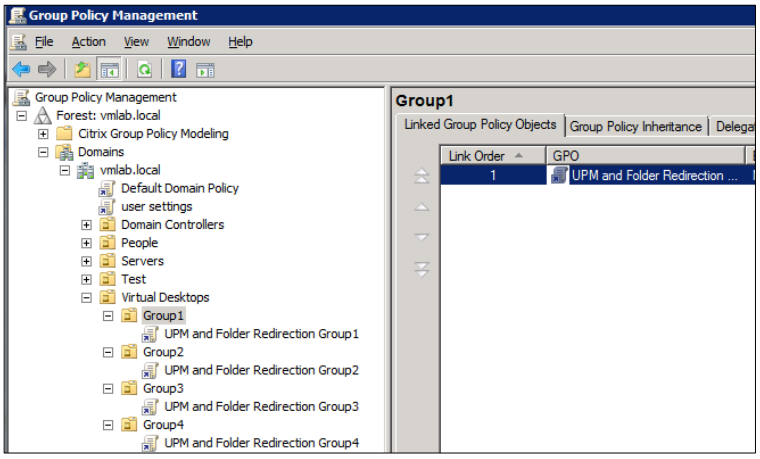

4.1.3 XenApp

Hosted desktops and applications, delivered via Microsoft XenApp, provide a sensible and low-cost entry into VDI by consolidating many basic task user sessions within a single Windows Server instance (formerly Terminal Services with XenApp). This technology can also be used to augment a XenDesktop deployment by delivering virtualized applications to user desktops via XenApp or App-V. XenDesktop 7 brings greater simplicity to the server hosted model simply by leveraging the Citrix Virtual Delivery Agent (VDA) installed within the Windows Server OS. The management infrastructure between desktop OS (desktop VMs) and server OS (XenApp) deployments is identical.

Hosted applications give users a customized application set with a "locally-installed" experience even though the applications are centrally installed and managed on XenApp servers. This can dramatically simplify the XenDesktop environment by leveraging a widely shared virtual desktop image, while at the same time extending the scalability of XenDesktop by alleviating resource load from the desktop compute servers. This two-Tiered approach to desktop and application delivery brings management simplification, a much quicker return on investment and the absolute best end-user experience.

Benefits of hosted desktop sessions and applications:

- Management of applications (single instance)
- Management of simple desktop images (no applications installed)
- PVS to stream XenApp servers as well as user desktops
- Scalability of XenDesktop compute hosts: CPU and IOPS reduction via application offload
- Shared storage scalability: less IOPS = more room to grow

Citrix XenDesktop with XenApp integration can effectively deliver a desktop/application hybrid solution as well. Specifically where a single or small number of shared VDI desktop images are deployed via XenDesktop, each with common shared applications installed within the golden image. A user-specific application set is then deployed and made accessible via the hosted application compute infrastructure, accessible from within the virtual desktop.

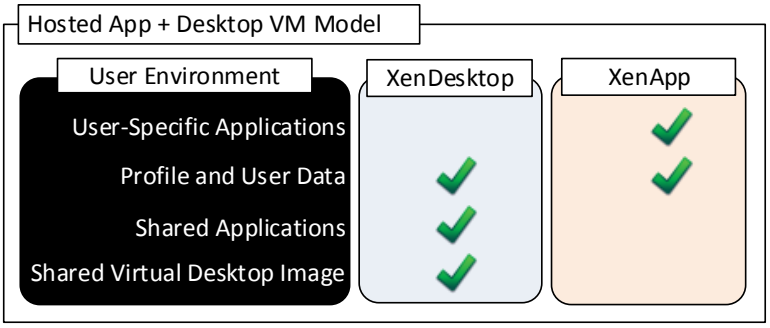

Alternatively, XenApp provides a platform for delivering Windows server-based sessions to users who may not need a full desktop VM. Hosted desktops increase infrastructure resource utilization while reducing complexity as all applications and sessions are centrally managed.

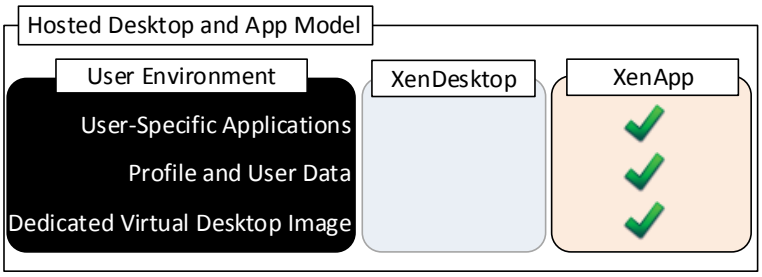

4.1.3.1 XenApp Integration into Dell Wyse Datacenter Architecture

The XenApp servers can exist as physical or virtualized instances of Windows Server 2012 R2. A minimum of one (1), up to a maximum of eight (8) virtual servers are installed per physical compute host. Since XenApp instances are easily added to an existing XenDesktop stack, the only additional components required are:

- One or more Server OS instances running the Citrix VDA added to the XenDesktop site

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

4.2 VDI Hypervisor Platforms

4.2.1 VMware vSphere 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.

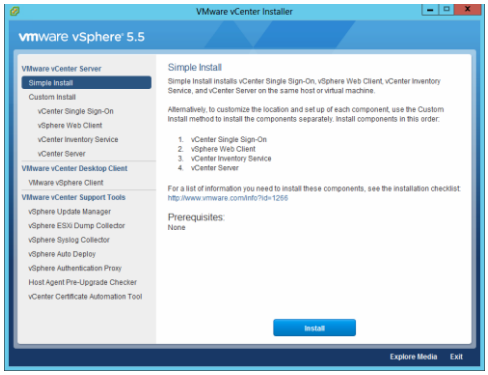

VMware vSphere 5 includes three major layers: Virtualization, Management and Interface. The Virtualization layer includes infrastructure and application services. The Management layer is central for configuring, provisioning and managing virtualized environments. The Interface layer includes the vSphere 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 XenDesktop components be installed in a distributed architecture, one role per server VM.

4.2.2 Microsoft Windows Server 2012 R2 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 XenDesktop virtualization environment is Microsoft Systems Center Virtual Machine Manager SP1 (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.

4.3 Citrix NetScaler

Citrix NetScaler is an all-in-one web application delivery controller that makes applications run five times better, reduces web application ownership costs, optimizes the user experience, and makes sure that applications are always available by using:

- Proven application acceleration such as HTTP compression and caching
- High application availability through advanced L4-7 load balancer
- Application security with an integrated AppFirewall
- Server offloading to significantly reduce costs and consolidate servers

Where Does a Citrix NetScaler Fit in the Network?

A NetScaler appliance resides between the clients and the servers, so that client requests and server responses pass through it. In a typical installation, virtual servers (vservers) configured on the NetScaler provide connection points that clients use to access the applications behind the NetScaler. In this case, the NetScaler owns public IP addresses that are associated with its vservers, while the real servers are isolated in a private network. It is also possible to operate the NetScaler in a transparent mode as an L2 bridge or L3 router, or even to combine aspects of these and other modes. NetScaler can also be used to host the StoreFront function eliminating complexity from the environment.

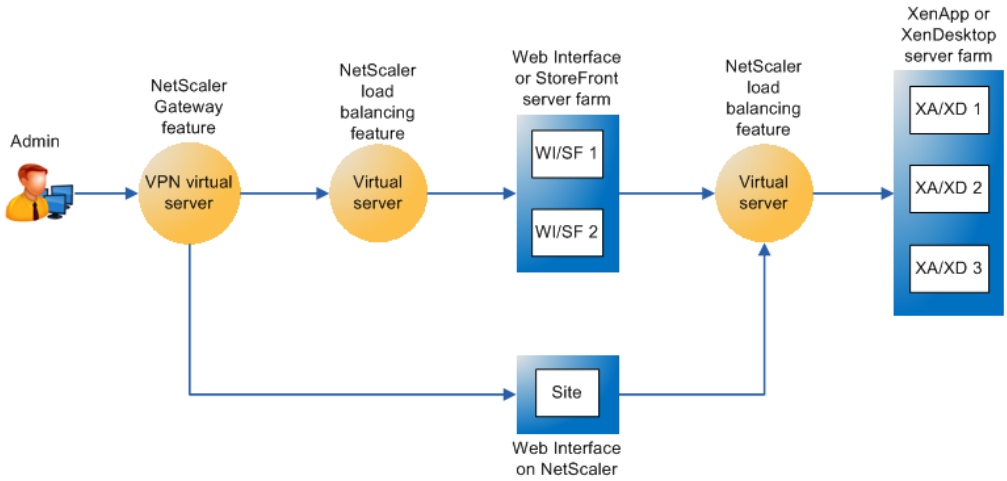

Global Server Load Balancing

GSLB is an industry standard function. It is in widespread use to provide automatic distribution of user requests to an instance of an application hosted in the appropriate data center where multiple processing facilities exist. The intent is to seamlessly redistribute load on an as required basis, transparent to the user community. These distributions are used on a localized or worldwide basis. Many companies use GSLB in its simplest form. They use the technology to automatically redirect traffic to Disaster Recovery (DR) sites on an exception basis. That is, GSLB is configured to simply route user load to the DR site on a temporary basis only in the event of a catastrophic failure or only during extended planned data center maintenance. GSLB is also used to distribute load across data centers on a continuous load balancing basis as part of normal processing.

XenDesktop HA with NetScaler White Paper: Link

Several of the management components of the XenDesktop stack are made highly-available using NetScaler to load balance traffic. The following management components require the use of a load balancer to function in a high availability mode:

- StoreFront Servers
- Licensing Server
- XenDesktop XML Service
- XenDesktop Desktop Director
- Provisioning Services TFTP Service

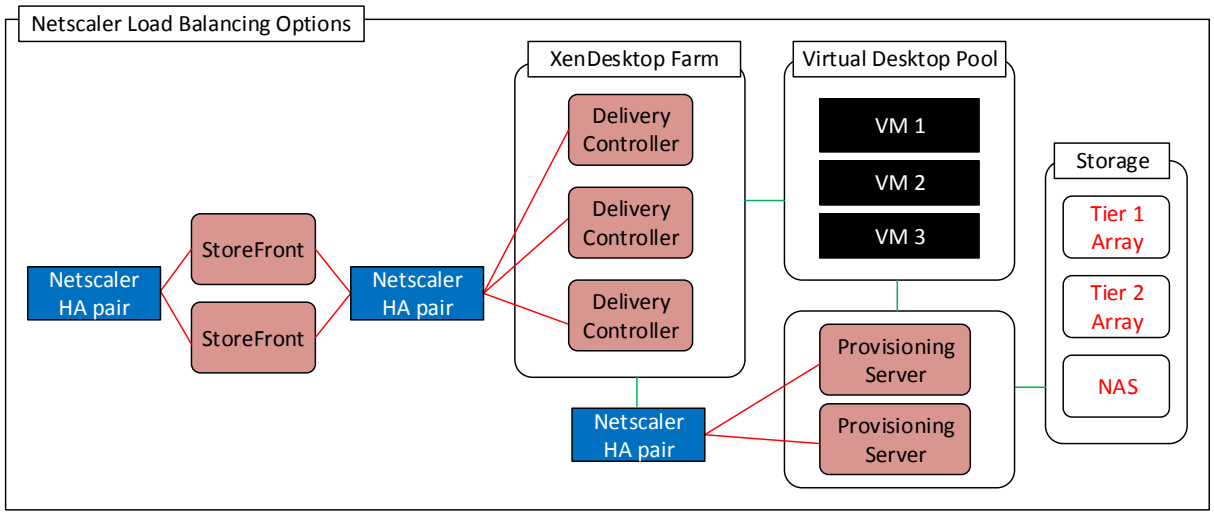

Citrix NetScaler is added to the Dell Wyse Datacenter mgmt. stack at any time and runs on the existing server infrastructure.

4.4 Citrix CloudBridge

Citrix CloudBridge provides a unified platform that connects and accelerates applications, and optimizes bandwidth utilization across public cloud and private networks. The only WAN optimization solution with integrated, secure, transparent cloud connectivity, CloudBridge allows enterprises to augment their data center with the infinite capacity and elastic efficiency provided by public cloud providers. CloudBridge delivers superior application performance and end-user experiences through a broad base of features, including:

- Market-leading enhancements for the Citrix XenDesktop user experience including HDX WAN optimization
- Secure, optimized networking between clouds
- Compression, de-duplication and protocol acceleration
- Acceleration of traditional enterprise applications
- Sophisticated traffic management controls and reporting
- Faster storage replication times and reduced bandwidth demands
- Integrated video delivery optimization to support increasing video delivery to branch offices
- Deliver a faster experience for all users

CloudBridge is ICA aware and enables IT organizations to accelerate, control and optimize all services – desktops, applications, multimedia and more – for corporate office, branch office and mobile users while dramatically reducing costs. With CloudBridge, branch office users experience a better desktop experience with faster printing, file downloads, video streaming and application start-up times.

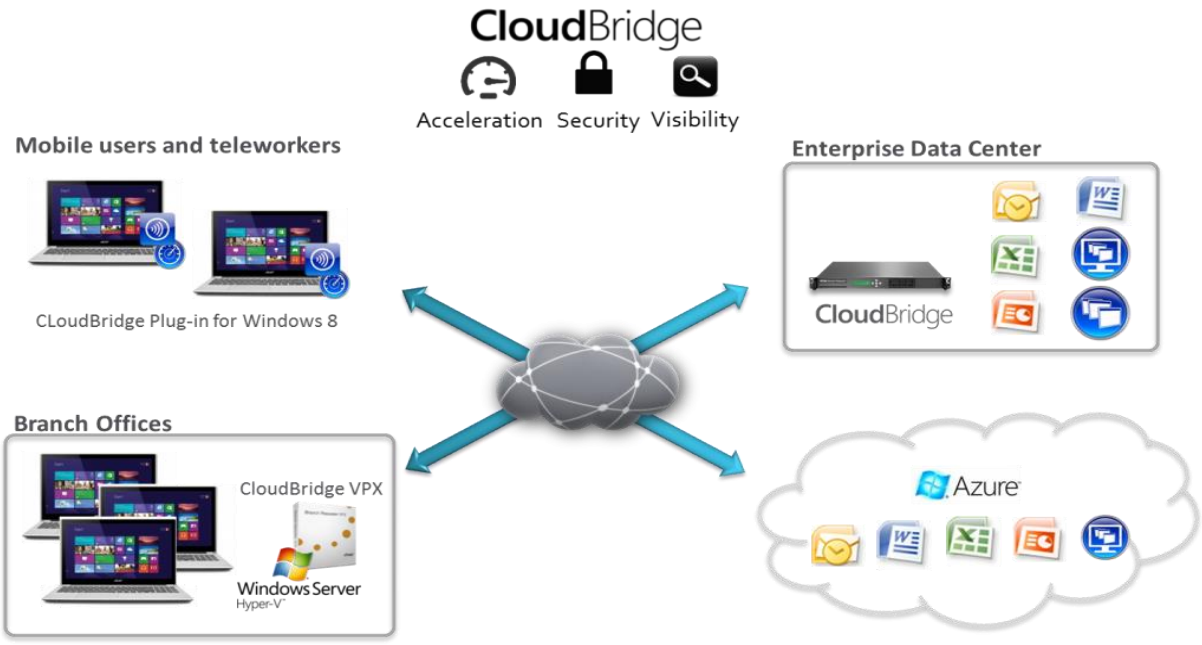

For more information please visit: Link

5 Solution Architecture for XenDesktop 7

5.1 Management Infrastructure

5.1.1 vSphere Management Role Requirements

Role	vCPU	vRAM (GB)	NIC	OS vDisk	
				Size (GB)	Location
Nutanix CVM	8*	32	2	-	C:\ (rear flexbay)
DDC + Lic	4	8	1	40	SDS: ds_mgmt
Storefront	2	4	1	40	SDS: ds_mgmt
Primary SQL	4	8	1	40 + 200	SDS: ds_mgmt
vCenter Appliance	2	8	1	125	SDS: ds_mgmt
Total	**20 vCPUs**	**60GB**	**6 vNICs**	**445GB**	**-**

5.1.2 Hyper-V Management Role Requirements

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:\ (rear flexbay)
DDC + Lic	4	8	512MB | 10GB	20%	Med	1	40	SDS: ds_mgmt
Storefront	2	4	512MB | 6GB	20%	Med	1	40	SDS: ds_mgmt
Primary SQL	4	8	512MB | 10GB	20%	Med	1	40 + 200	SDS: ds_mgmt
SCVMM	2	8	512MB | 10GB	20%	Med	1	125	SDS: ds_mgmt
Total	**20 vCPUs**	**60GB**	**2GB | 36GB**	**-**	**-**	**6 vNICs**	**445GB**	**-**

5.1.3 XenApp on vSphere

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

Role	Platform	VMs per host	vCPU	RAM (GB)	NIC	OS vDisk	
						Size (GB)	Location
XenApp VM	Entry	6	3	32	1	80	SDS: ds_rdsh
XenApp VM	Mid	6	4	32	1	80	SDS: ds_rdsh
XenApp VM	High	6	5	32	1	80	SDS: ds_rdsh

5.1.4 XenApp on Hyper-V

The recommended number of XenApp VMs and their configurations on Hyper-V are summarized below based on applicable hardware platform.

Role	Platform	VMs per host	vCPU	Startup RAM (GB)	Dynamic Memory			NIC	OS vDisk	
					Min|Max	Buffer	Weight		Size (GB)	Location
XenApp VM	Economy	6	3	16	512MB | 32GB	20%	Med	1	80	SDS: ds_rdsh
XenApp VM	Mid	6	4	16	512MB | 32GB	20%	Med	1	80	SDS: ds_rdsh
XenApp VM	High	6	5	16	512MB | 32GB	20%	Med	1	80	SDS: ds_rdsh

5.1.5 SQL Databases

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

- Citrix XenDesktop
- vCenter or SCVMM

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 Citrix, Microsoft and VMware are to be adhered to, to ensure optimal database performance.

The EqualLogic PS series arrays utilize a default RAID stripe size of 64K. To provide optimal performance, configure disk partitions to begin from a sector boundary divisible by 64K.

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).

5.1.6 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 Citrix 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, Citrix 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.

5.1.6.1 DNS for SQL

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

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

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

5.2 Storage Architecture Overview

All Dell 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 resource demands of the Controller VMs. Capacity tier deduplication requires that CVMs be configured with 32GB RAM.

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

5.3 Networking

5.3.1 vSphere

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

5.3.1.1 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.

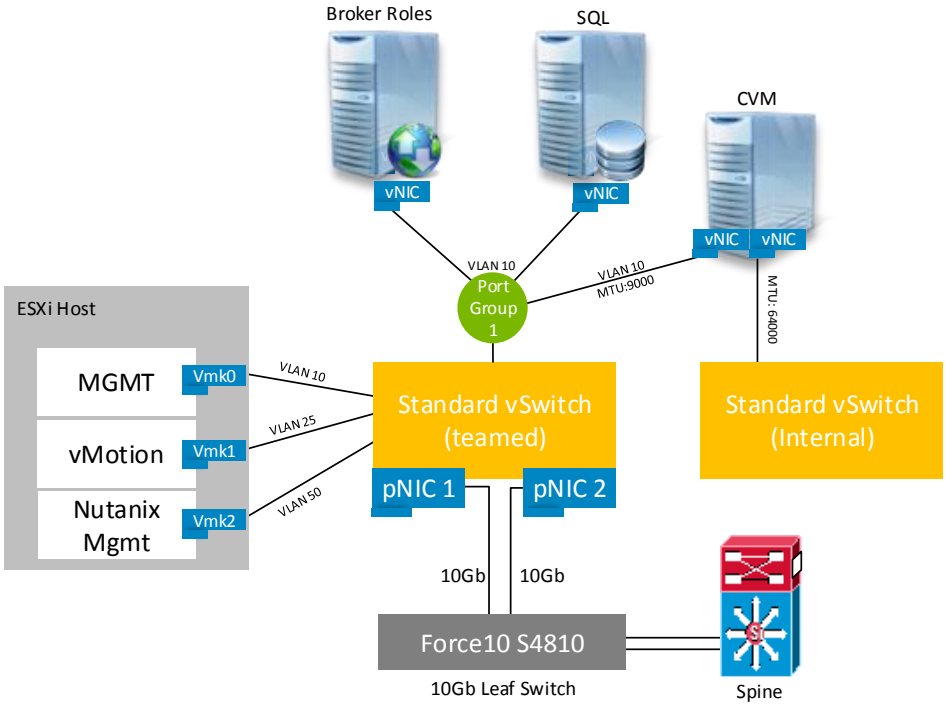

The Compute hosts are configured in the same basic manner with the desktop VMs connecting to the primary port group on the external vSwitch.

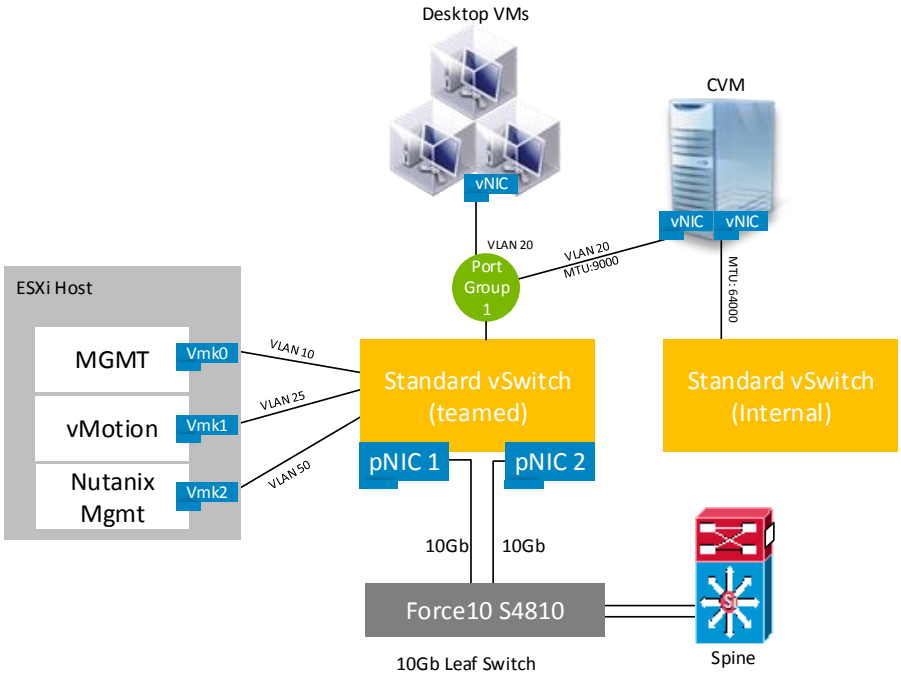

5.3.2 Hyper-V

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.

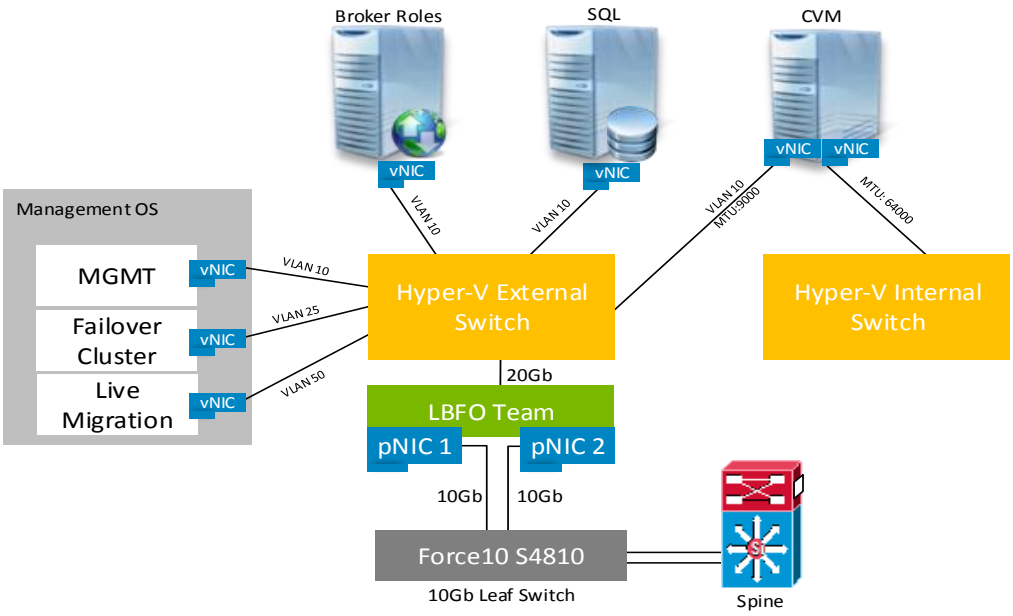

The NIC team for the Hyper-V switch is configured as switch independent, Dynamic for the load balancing mode with all adapters set to active. This team is used exclusively by Hyper-V.

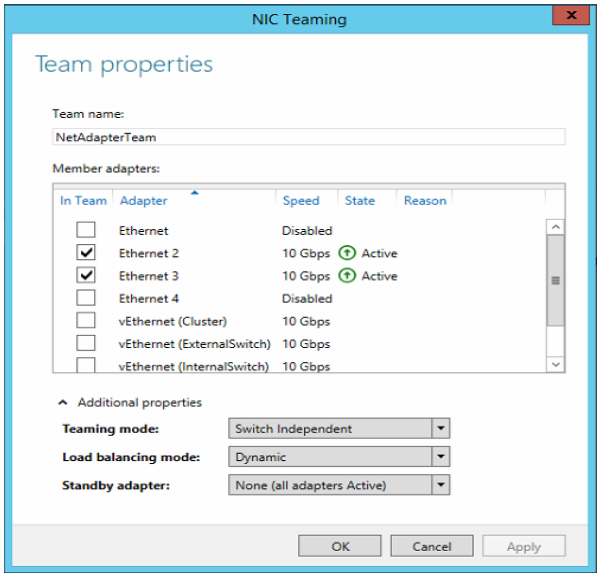

The dedicated compute host configuration is shown in the diagram below and configured very similarly to the management host configuration.

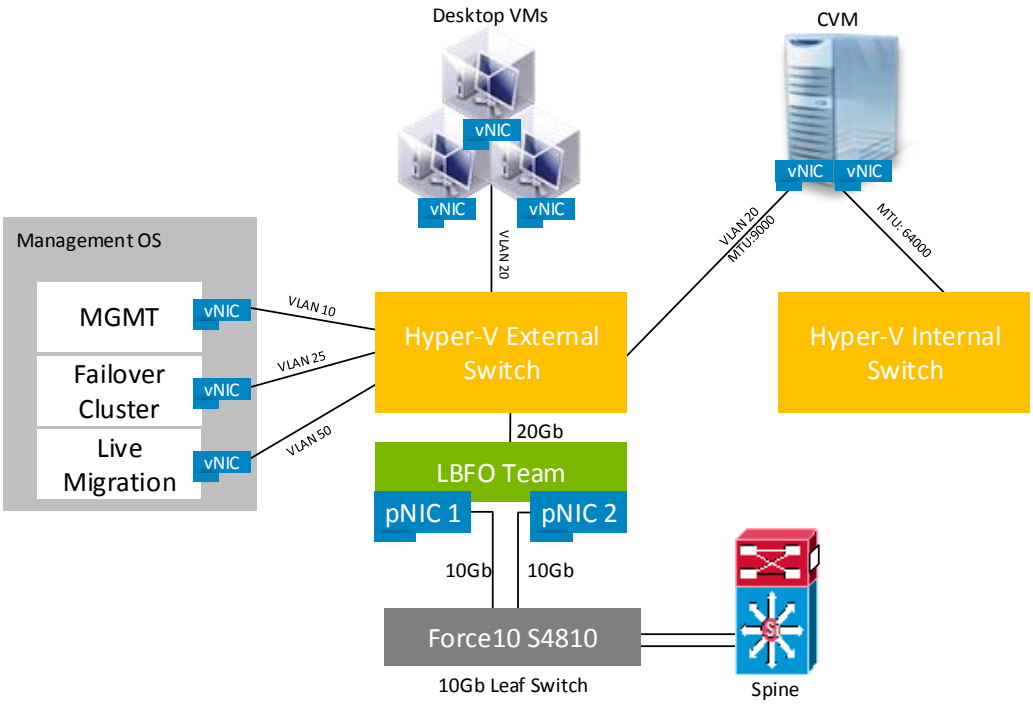

5.4 Scaling Guidance

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 (32 nodes for vSphere, 64 nodes for Hyper-V).

- 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
Provisioning Servers	Desktops per instance	Additional servers added to the Provisioning Server farm	Additional network and I/O capacity added to the servers
Desktop Delivery Servers	Desktops per instance (dependent on SQL performance as well)	Additional servers added to the XenDesktop Site	Additional virtual machine resources (RAM and CPU)
XenApp Servers	Desktops per instance	Additional virtual servers added to the XenDesktop Site	Additional physical servers to host virtual XenApp servers.
Storefront Servers	Logons/ minute	Additional servers added to the Storefront environment	Additional virtual machine resources (RAM and CPU)
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 or 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 ESXi or 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 Citrix infrastructure server roles can then be duplicated on the new host where connections to each are load balanced via the addition of a virtual NetScaler 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 Dell Wyse Datacenter for XenDesktop Communication Flow

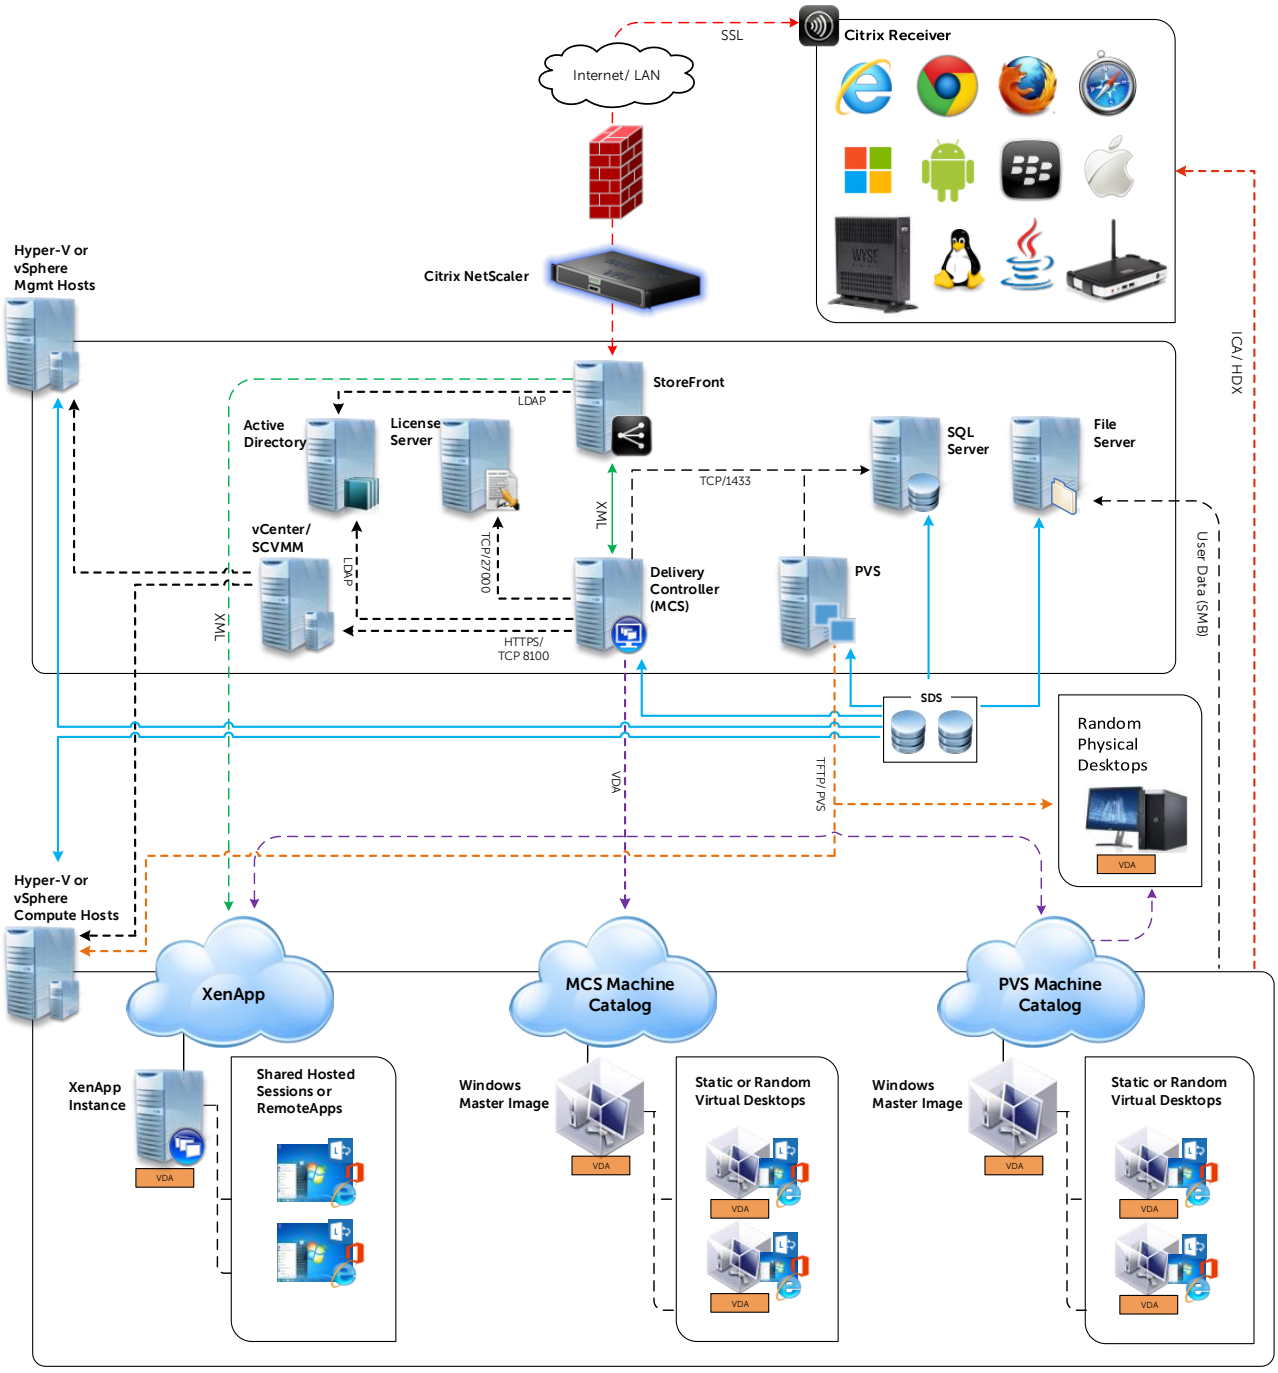

6 Solution Performance and Testing

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 installs a standard collection of desktop application software (e.g. Microsoft Office, Adobe Acrobat Reader) on each VDI desktop; it then uses launcher systems to connect a specified number of users to available desktops within the environment. Once the user is connected the workload is started via a logon script which starts the test script once the user environment is configured by the login script. Each launcher system can launch connections to a number of 'target' machines (i.e. VDI desktops), with the launchers being managed by a centralized management console, which is used to configure and manage the Login VSI environment.

6.1.2 Liquidware Labs Stratusphere UX

Stratusphere UX was used during each test run to gather data relating to User Experience and desktop performance. Data was gathered at the Host and Virtual Machine layers and reported back to a central server (Stratusphere Hub). The hub was then used to create a series of "Comma Separated Values" (.csv) reports which have then been used to generate graphs and summary tables of key information. In addition the Stratusphere Hub generates a magic quadrate style scatter plot showing the Machine and IO experience of the sessions. The Stratusphere hub was deployed onto the core network therefore its monitoring did not impact the servers being tested. This core network represents an existing customer environment and also includes the following services:

- Active Directory
- DNS
- DHCP
- Anti-Virus

Stratusphere UX calculates the User Experience by monitoring key metrics within the Virtual Desktop environment, the metrics and their thresholds are shown in the following screen shot:

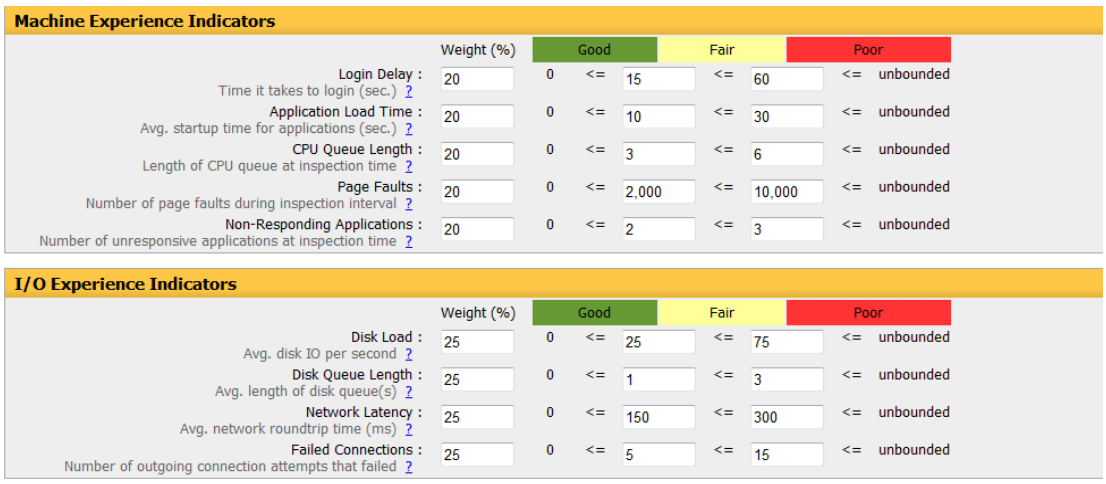

Machine Experience Indicators

	Weight (%)	Good		Fair		Poor	
Login Delay : Time it takes to login (sec.) ?	20	0 <=	15	<=	60	<=	unbounded
Application Load Time : Avg. startup time for applications (sec.) ?	20	0 <=	10	<=	30	<=	unbounded
CPU Queue Length : Length of CPU queue at inspection time ?	20	0 <=	3	<=	6	<=	unbounded
Page Faults : Number of page faults during inspection interval ?	20	0 <=	2,000	<=	10,000	<=	unbounded
Non-Responding Applications : Number of unresponsive applications at inspection time ?	20	0 <=	2	<=	3	<=	unbounded

I/O Experience Indicators

	Weight (%)	Good		Fair		Poor	
Disk Load : Avg. disk IO per second ?	25	0 <=	25	<=	75	<=	unbounded
Disk Queue Length : Avg. length of disk queue(s) ?	25	0 <=	1	<=	3	<=	unbounded
Network Latency : Avg. network roundtrip time (ms) ?	25	0 <=	150	<=	300	<=	unbounded
Failed Connections : Number of outgoing connection attempts that failed ?	25	0 <=	5	<=	15	<=	unbounded

6.1.3 VMware vCenter

VMware vCenter was used for VMware vSphere-based solutions to gather key data (CPU, Memory and Network usage) from each of the desktop hosts during each test run. This data was exported to .csv files for each host and then consolidated to show data from all hosts. While the report does not include specific performance metrics for the Management host servers, these servers were monitored during testing and were seen to be performing at an expected performance level.

6.1.4 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.

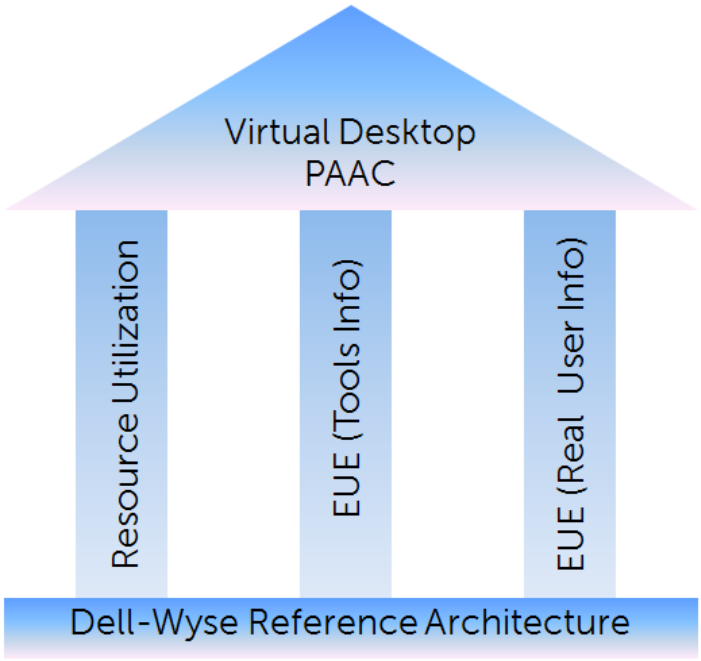

Each of the pillars shown in the above diagram is discussed in more detail below.

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 EUE (Tools Info)

Good EUE is one of the primary factors in determining the success of a VDI implementation. As a result of this fact, a number of vendors have developed toolsets that monitor the environmental parameters that are relevant to EUE. PAAC on Dell Wyse Datacenter solutions uses the Liquidware Labs Stratusphere UX tool to ensure that good EUE is delivered for the density numbers defined in our RAs. More specifically, our PAAC analysis uses a scatter plot provided by Stratusphere UX which presents end-user experience for all load-tested users. Stratusphere UX does this by algorithmically combining relevant parameters in relation to virtual desktop machine experience (e.g. login duration) and virtual desktop IO experience (e.g. disk queue length) to provide a plot that shows end-user experience as good, fair or poor using a magic-quadrant type approach.

6.2.3 EUE (Real User Info)

In order to complement the tools based end-user experience information gathered using Stratusphere UX (as described above) and to provide further certainty around the performance of Dell Wyse Datacenter solutions, PAAC on our solutions also involves a user logging into one of the solutions when they are fully loaded (based on the density specified in the relevant RA) and executing user activities that are representative of the user type being tested (e.g. task, knowledge or power user). An example is a knowledge worker executing a number of appropriate activities in Excel. The purpose of this activity is to verify that the end-user experience is as good as the user would expect on a physical laptop or desktop.

6.2.4 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.

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

6.2.5 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.

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

6.2.6 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	Shared
Enhanced	Login VSI Medium	Shared
Professional	Login VSI Heavy	Shared + Profile Virtualization

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:

- For single-server / single-host testing (typically carried out to determine the virtual desktop capacity of a specific physical server), users are logged in every 30 seconds.
- For multi-host / full solution testing, users are logged in over a period of 1-hour, to replicate the normal login storm in an enterprise environment.
- 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.

6.2.6.1 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.

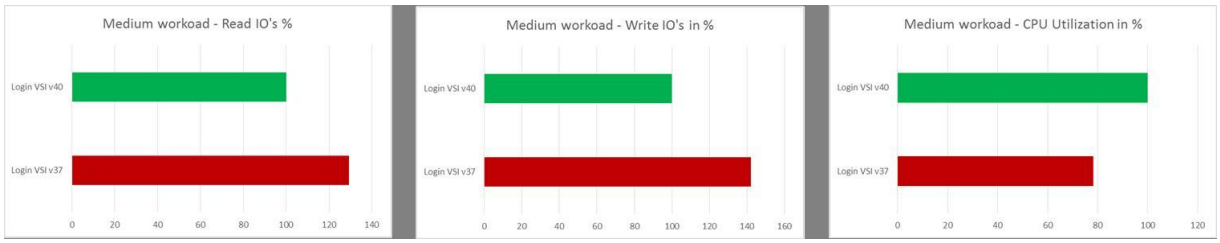

6.2.6.2 Login VSI Light Workload

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.

6.2.6.3 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.

6.2.6.4 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.

6.2.7 Workloads Running on Shared Graphics Profile

Graphics hardware vendors (e.g. Nvidia) typically market a number of graphics cards that are targeted at different markets. Consequently, it is necessary to provide two shared graphics workloads – one for mid-range cards and the other for high-end cards. These workloads are described in more detail below. It is noted that technologies such as the Citrix / Nvidia vGPU technology (where the Nvidia drivers reside on the virtual desktop, giving shared-level density with near pass-through functionality) mean that in some cases, the higher-end workloads, traditionally used for pass-through GPU PAAC, may be more appropriate for shared GPU PAAC. Such scenarios will explicitly state the workload used.

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. A pilot run is conducted to validate that the infrastructure is functioning and valid data is captured. Subsequently three more runs are conducted allowing for correlation of data. Summary of the test results is listed out in the below mentioned tabular format.

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 is 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.

Login VSI has two modes for launching user's sessions:

- Parallel

  Sessions are launched from multiple launcher hosts in a round robin fashion; this mode is recommended by Login Consultants when running tests against multiple host servers. In parallel mode the VSI console is configured to launch a number of sessions over a specified time period (specified in seconds)

- Sequential

  Sessions are launched from each launcher host in sequence, sessions are only started from a second host once all sessions have been launched on the first host and this is repeated for each launcher host. Sequential launching is recommended by Login Consultants when testing a single desktop host server. The VSI console is configure to launch a specified number of session at a specified interval specified in seconds

All test runs which involved the six desktop hosts were conducted using the Login VSI "Parallel Launch" mode and all sessions were launched over an hour to try and represent the typical 9am logon storm. Once the last user session has connected, the sessions are left to run for 15 minutes prior to the sessions being instructed to logout at the end of the current task sequence, this allows every user to complete a minimum of two task sequences within the run before logging out. The single server test runs were configured to launch user sessions every 60 seconds, as with the full bundle test runs sessions were left to run for 15 minutes after the last user connected prior to the sessions being instructed to log out.

6.4 XenDesktop Test Results

This validation was performed for XenDesktop 7.5

Validation was performed using CCC standard testing methodology using LoginVSI 4 load generation tool for VDI benchmarking that simulates production user workloads.

Stratusphere UX was not used in this validation due to compatibility issues/

Each test run adhered to PAAC best practices with a 60minute logon period, 15 minutes of steady state and a 15minute logoff window.

As outlined in section 1 multiple configurations were tested as part of this project.

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

Hyper-visor	Provisioning	Workload	Density Per Host	Avg CPU %	Avg Memory Usage GB	Avg IOPS/User	Avg Net Kbps/User
ESXi	MCS	Standard	60*	N/A	N/A	N/A	N/A
		Enhanced	70	99%	53GB	8.6	1656.5
		Professional	80	98%	65GB	4.5	1503.2
ESXi	XenApp	Standard	70*	N/A	N/A	N/A	N/A
		Enhanced	90	92%	32GB	5.76	872.6
		Professional	125	97%	56GB	3.6	742.3

***Extrapolated**

CPU Utilization CPU % for ESX Hosts was adjusted to account for the fact that on Intel E5-2600v2 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

Memory utilization

Network utilization was higher for the equivalent Hyper-V solution than the ESXi solution for MCS. As expected the PVS solution produced significantly higher network utilization than MCS

The IOPS results above were surprising in that PVS typically would offload IO to the network, relieving the local storage from a significant number of IOPS. In the chart above however it shows that PVS test cases consistently produced more IOPS per user than equivalent MCS test cases. As expected PVS with write cache in RAM showed practically no tier 1 IOPs.

Network Utilization was consistent between the solutions. It was expected that the PVS solution would show higher IOPs than MCS or physical solutions. However, this was not depicted in the results gathered. This is most likely due to caching.

6.5 XenDesktop Test Results

6.5.1 Mid Configuration ESXi/XenDesktop 140 Users

The architecture for this solution includes a dedicated Management host & two compute hosts. Each of the compute hosts we populated with 70 Login VSI Medium users.

The below graph shows the performance data for 70 Users per host on a pair of Compute hosts. CPU reaches an AVG of 99% at peak of the test cycle. This does not account for the additional 10% of available CPU capacity provided by Turbo Boost. Each host will show a total MHz availability of 41584 MHz however with Turbo boost the available MHz value is 45742MHz. You can remove 10000MHz instantaneously from the available CPU for sessions as this is the reservation set for the CVM on each node in the cluster.

The final host in the cluster is a dedicated management host which is running the vCenter & Citrix XenDesktop management virtual machines. As expected its CPU utilization is significantly lower than the compute hosts in the cluster. As you can see from the below image the CPU utilization for this host does not exceed 25% at any point in the test cycle. The main CPU activity is during the 60minute session launch period and once this is complete the utilization drops off to circa 15% before logoffs begin.

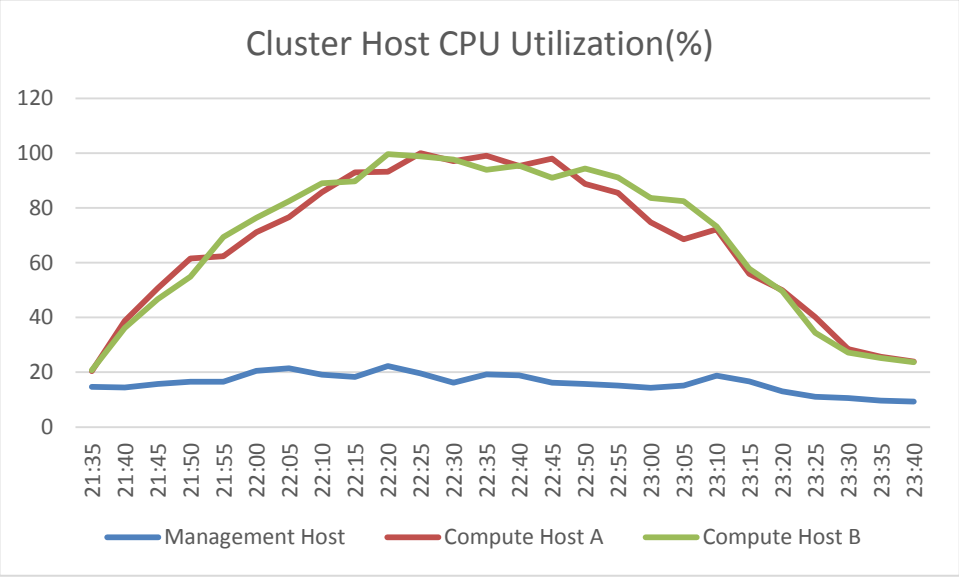

*** This does not account for the additional 10% of available CPU capacity provided by Turbo Boost.**

Below is a combined graph illustrating the CPU utilization breakdown for the management host for this test run. All virtual machines with the exception of storefront used less than 10% of their available CPU. Storefront processed all of the desktop connections and performed really well at fewer than 50% of its allocated CPU shares.

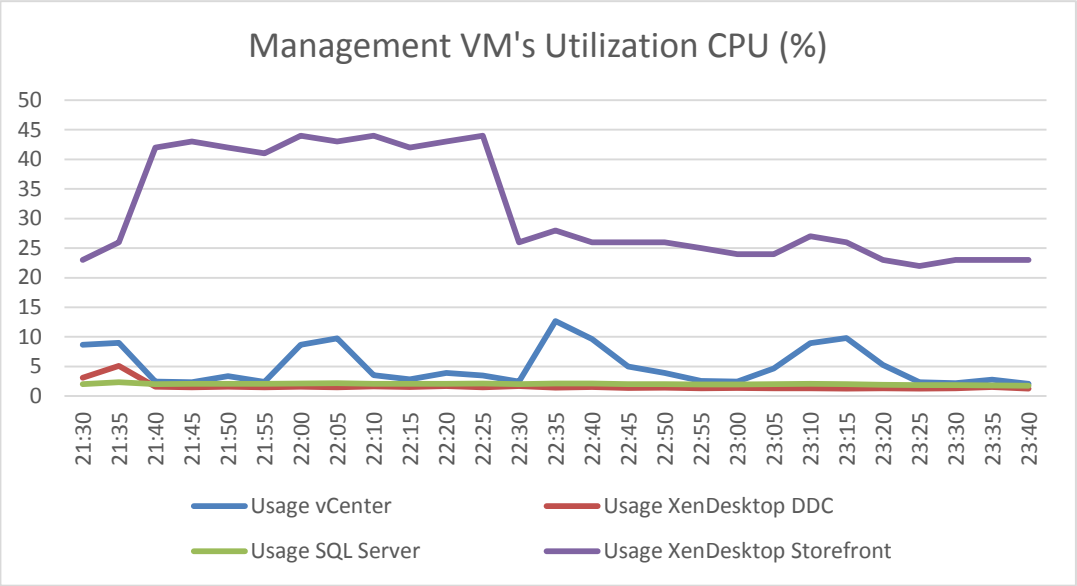

*** This does not account for the additional 10% of available CPU capacity provided by Turbo Boost.**

Next, we look at Memory consumption for the entire cluster. Starting with the management cluster out of a total of 384GB of available memory, there were clearly no constraints. Active memory was extremely low for the management host.

The below graphs represent the memory utilization for the compute nodes. Consumed memory did not exceed 70% of the available memory and active memory was extremely low with no ballooning occurring.

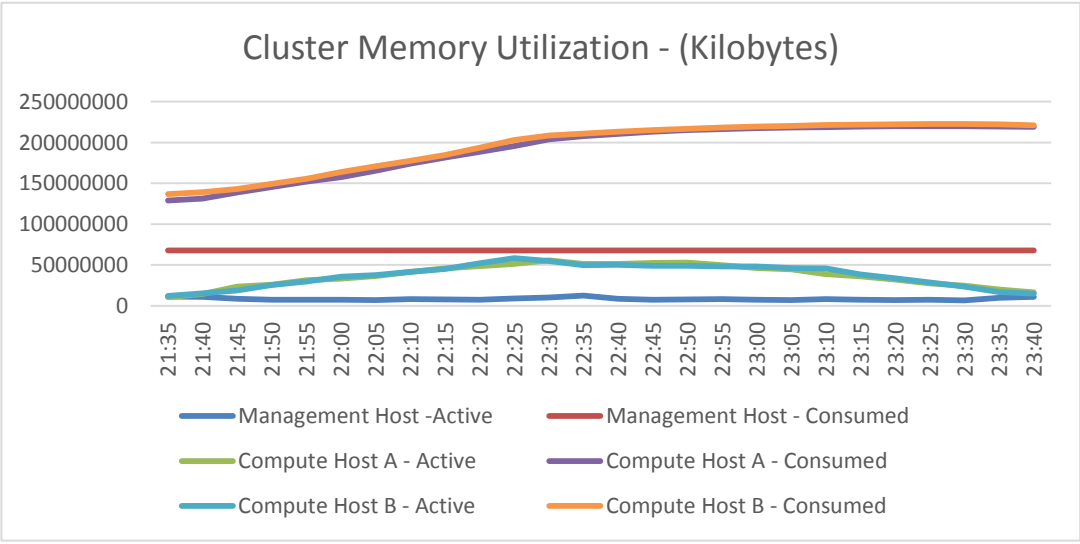

Next, we look at the Network performance for the cluster. As expected the management consumed less network bandwidth that both compute hosts however based on the Nutanix architecture the CVM on the management host may perform operations for the other hosts in the cluster.

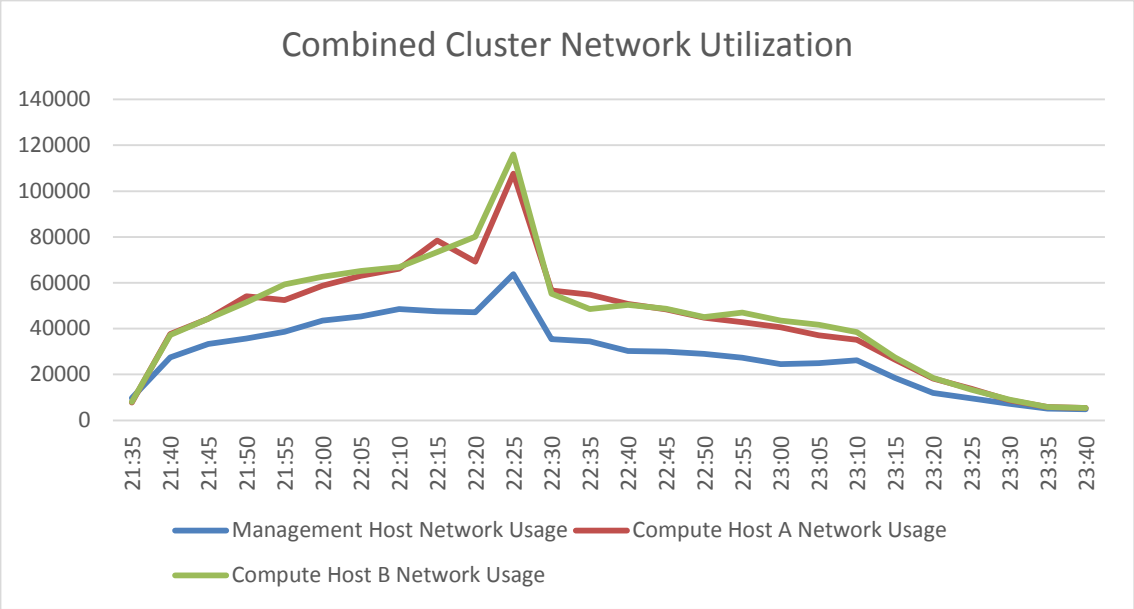

We also measure the user experience for each test run. LoginVSI was used for both load generation and to measure the user experience. VSI Max was not reached for this test as can be seen below.

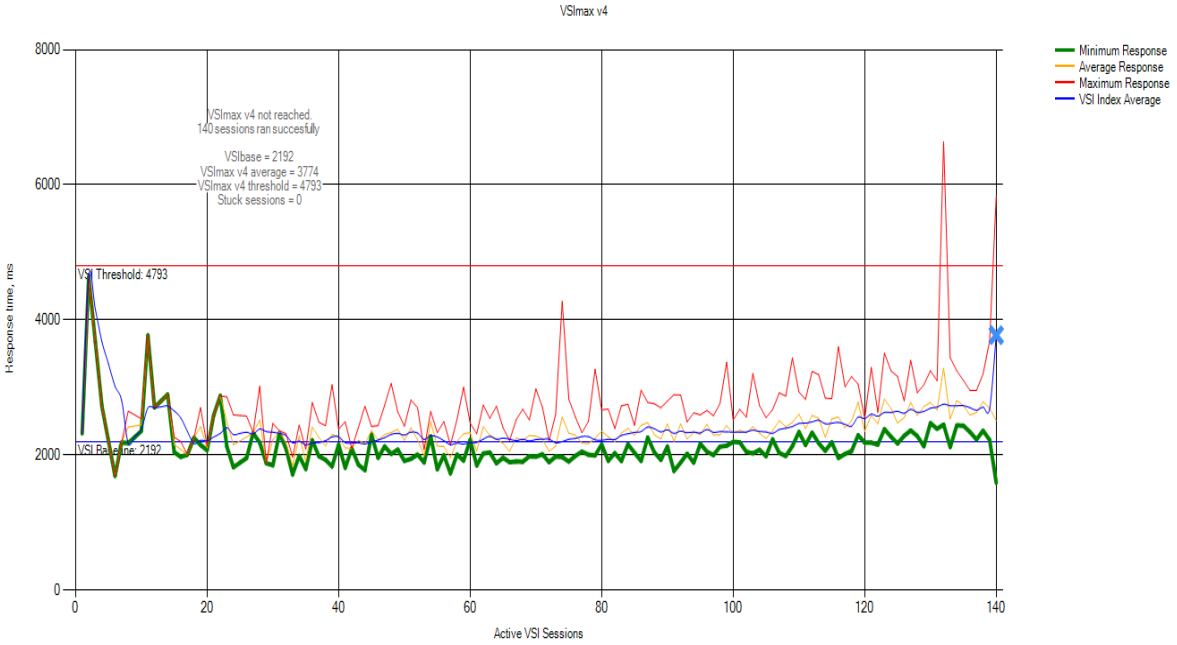

The storage platform for this testing was the Dell/Nutanix XC appliance, below are some of the graphs showing the CPU, IOPS, Network and Memory consumption for the cluster.

The cluster CPU utilization measured on the Nutanix management appliances shows as maximum of 72% during the test run.

The below combined graph shows the IOPS consumption per host at peak of the test. At this point all sessions are just after logging in.

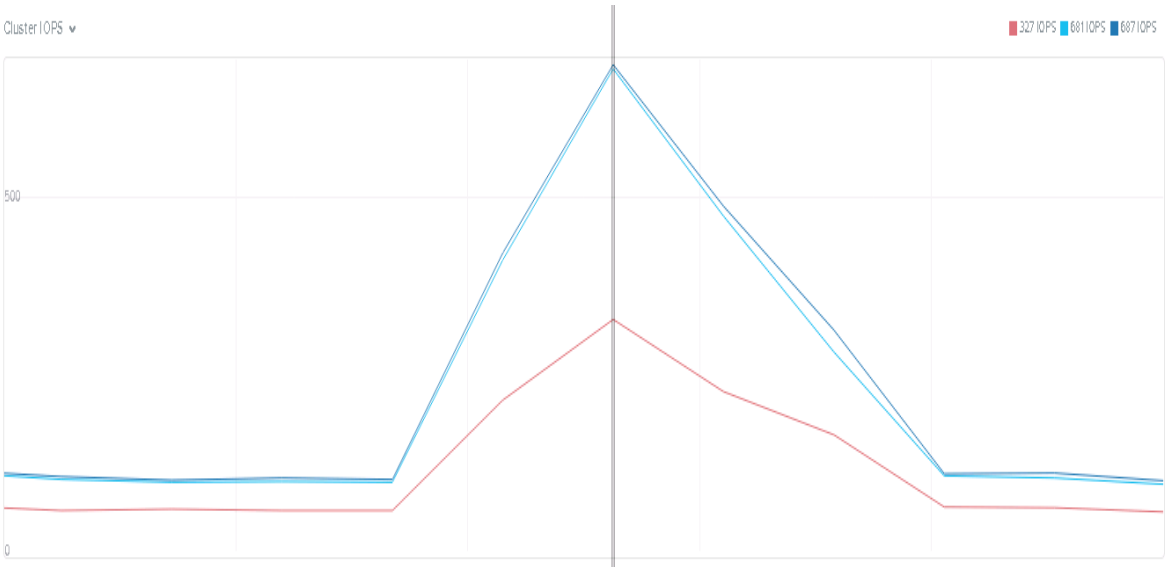

Cluster Memory did not exceed 44% during the steady state of the test contributing to the excellent storage performance

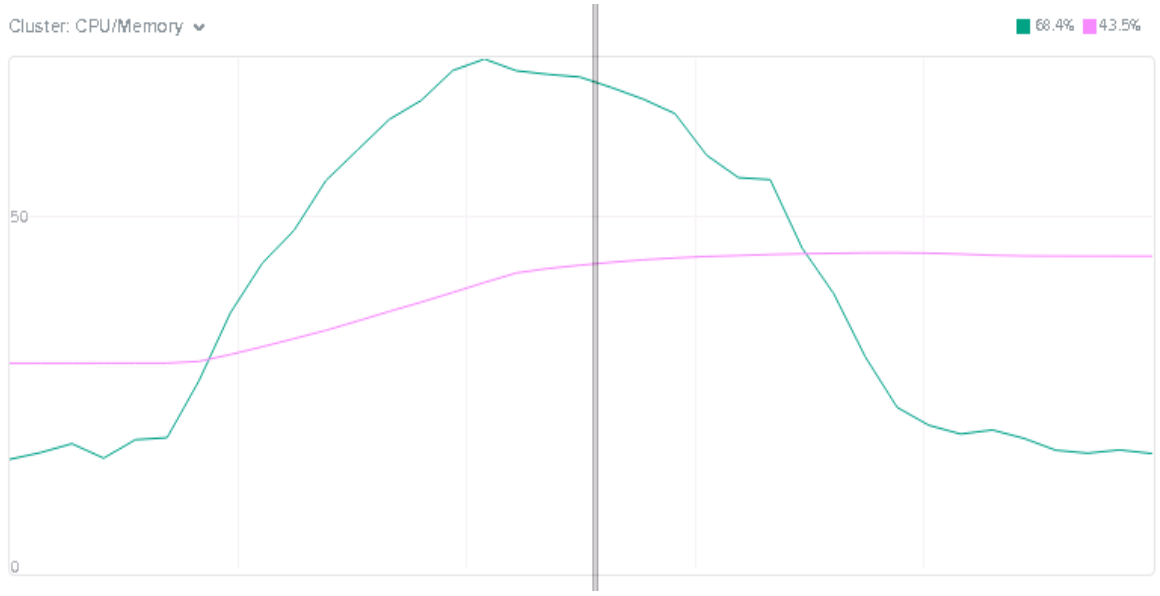

In summary, it would not be recommended to exceed 70 Users per host as CPU is the bottleneck. The Nutanix CVM accounts for almost 20% of the resources due to its reservation. The cluster performance was excellent with low IOPS consumption. User experience was also excellent with VSI Max not being reached.

The main observation from the tests run was the CVM's need time to recover from intense I/O operations such as a boot storm. Typically up to 60minutes were allocated during this testing prior to commencement of each test.

6.5.2 High Workload ESXi/XenDesktop MCS (160 Users)

The architecture for this solution includes a dedicated Management host & two compute hosts. Each of the compute hosts we populated with 80 Login VSI Medium users.

The below graph shows the performance data for 80 Users per host on a pair of Compute hosts. CPU reaches an AVG of 99% at peak of the test cycle. This does not account for the additional 10% of available CPU capacity provided by Turbo Boost. Each host will show a total MHz availability of 55980 MHz however with Turbo boost the available MHz value is 61578MHz. You can remove 10000MHz from the available CPU instantaneously for sessions as this is the reservation set for the CVM on each node in the cluster.

Below is the CPU utilization for both Compute hosts in the cluster. Utilization hit the high 90 percentage without calculating the additional 10% from Turbo boost which is enabled by default on the R720 PowerEdge Servers.

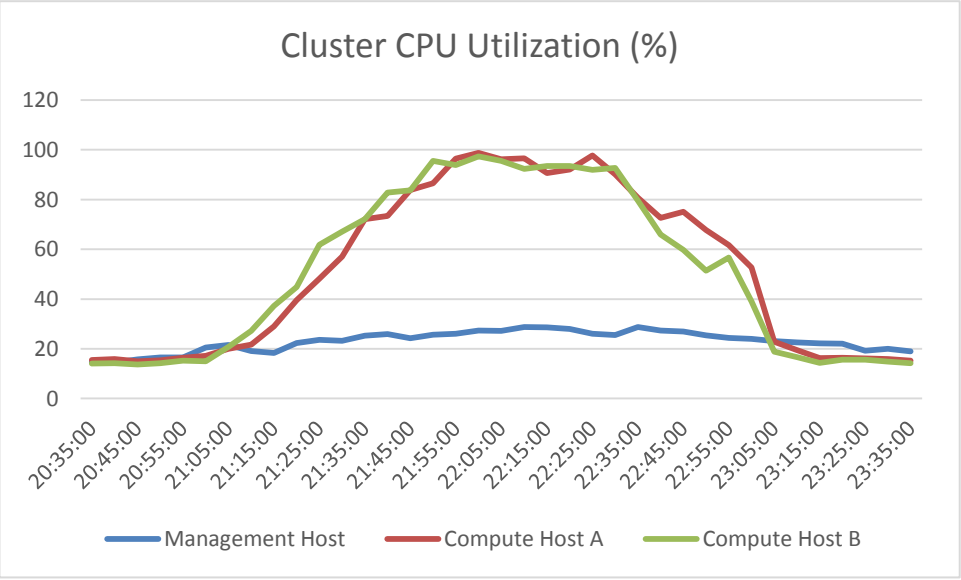

*** This does not account for the additional 10% of available CPU capacity provided by Turbo Boost.**

Next, we look at the Memory performance for this test. We begin with the management host which recorded a low level of utilization as expected.

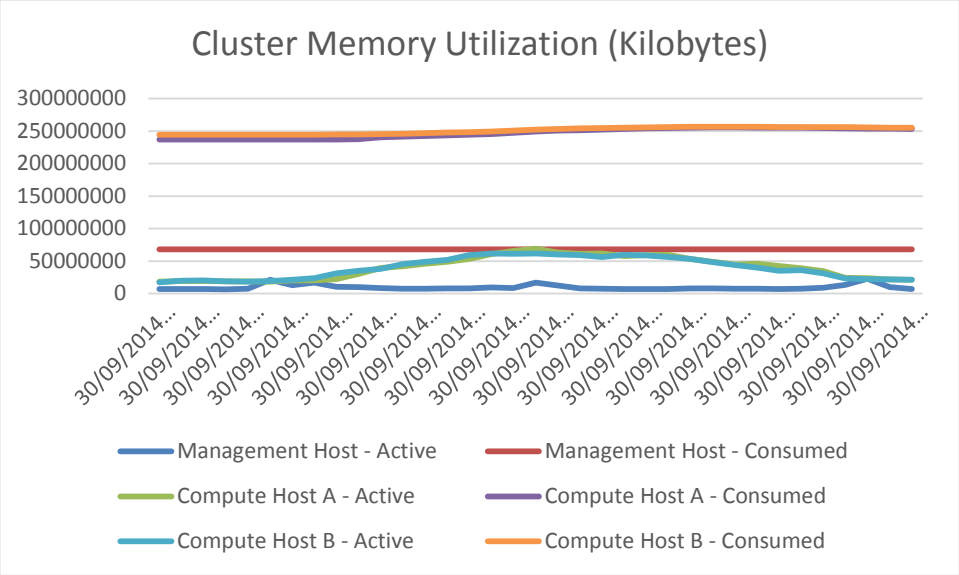

As expected for the compute hosts no memory ballooning or swapping occurred during the testing. Active memory utilization did not exceed 30% of the available memory for both hosts.

Network utilization again was low on each of the hosts in the cluster. The below graph is a combined cluster network utilization graph based on 10GB networking.

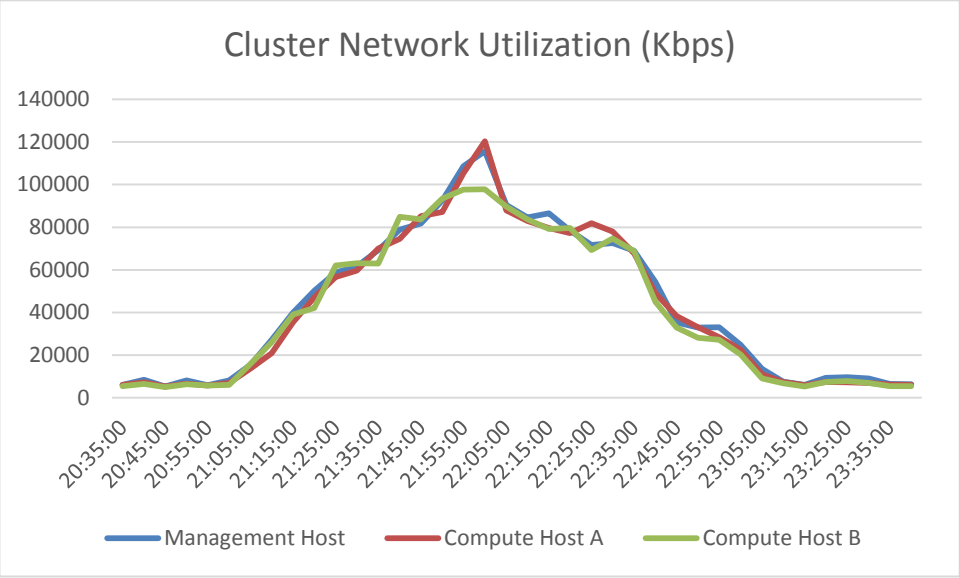

Next, we look at the VSI Scoring for this test. Based on the VSI score the max score was not reached however CPU did not reflect this on each of the compute hosts with CPU reaching 990% plus on each compute node.

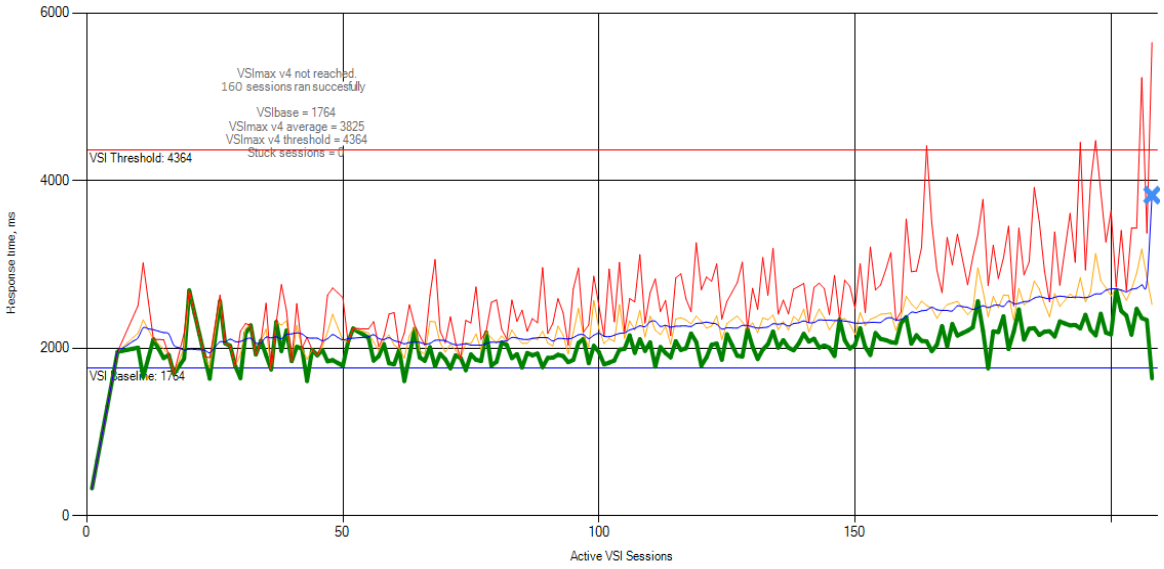

The storage platform for this testing was the Dell/Nutanix XC appliance, below are some of the graphs showing the CPU, IOPS, Network and Memory consumption for the cluster.

The maximum CPU value recorded for the Cluster was CPU 71.1% indicating sufficient capacity to increase the number of users on this storage pool.

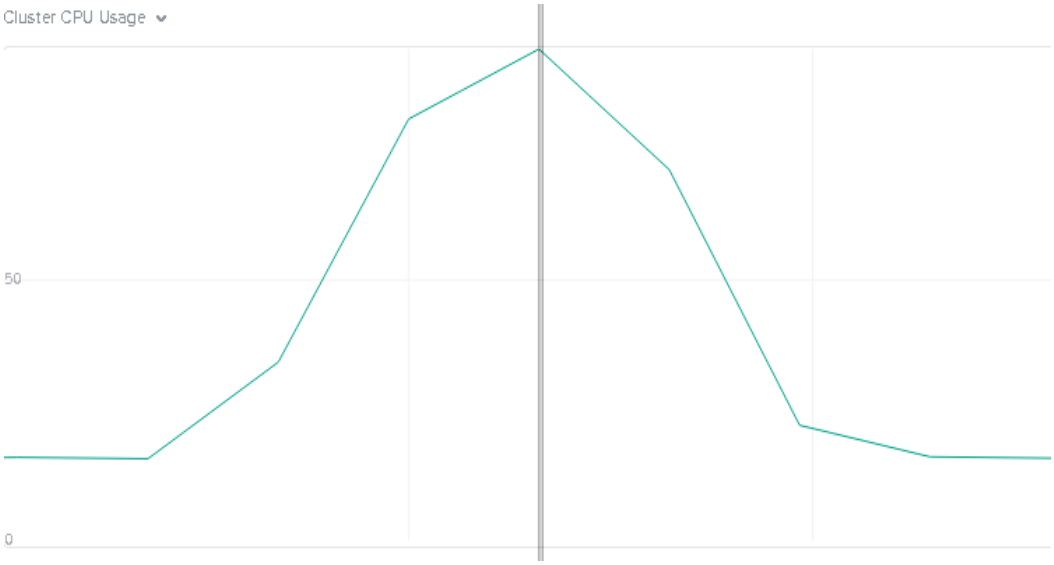

Cluster Compute IOPS can be observed below. As per previous tests the storage was showing low utilization with low Latency

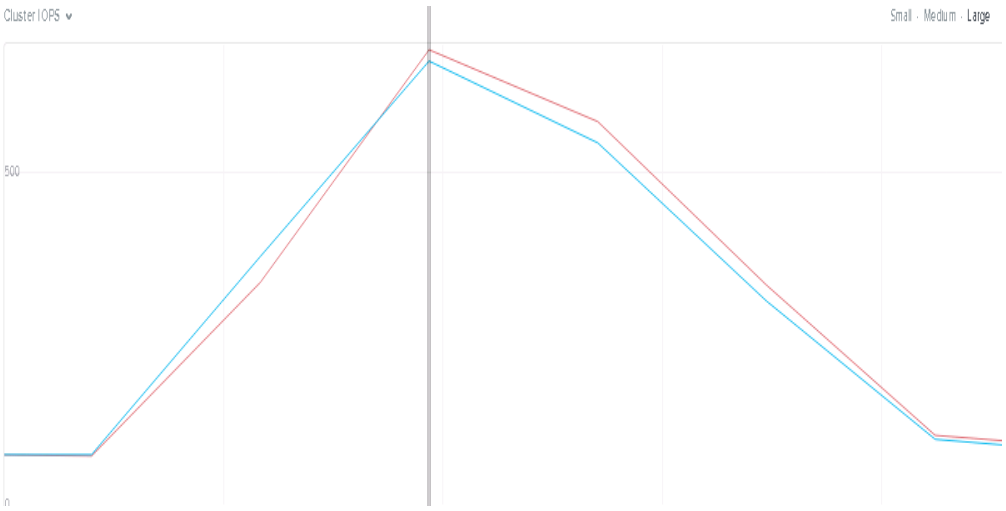

In summary based on the CPU performance it is not recommended to exceed 75-80 users per host with a dual vCPU configuration. The CVM accounts for almost 20% of the available CPU on each node. Also "prebooting" desktops is recommended as the CVM struggles to catch-up with the physical CPU requiring more time to settle on each cluster VM.

6.6 XenApp Test Results

6.6.1 Mid workload for 180 XenApp users on XenDesktop

The XenApp Session hosts were brokered by Citrix XenDesktop MCS 7.5 with a dedicated management host and two Compute hosts. Each compute host had a total of 6 VM's running Server2012 RDS. Each host had a total of 90 sessions running during testing. The standard LoginVSI application packages were installed and the Login VSI Medium workload was used for load generation.

The graph below shows the CPU utilization for the Management, Compute hosts and the XenApp VM's.

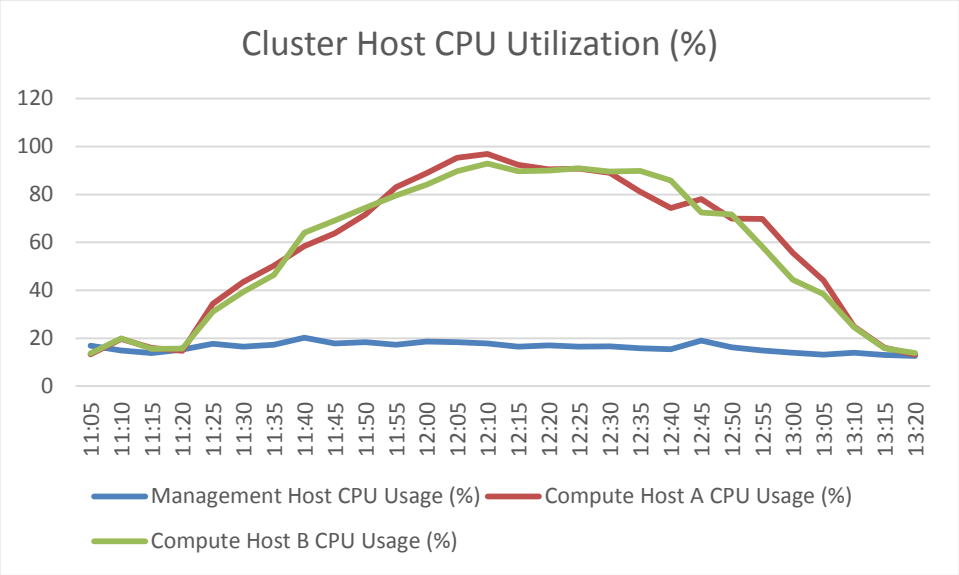

*** This does not account for the additional 10% of available CPU capacity provided by Turbo Boost.**

As you can see from the graph, CPU for the Management host did not exceed 20% while the Compute hosts both stayed below 100%, as outlined in section 1 turbo boost is enabled by default providing an additional 10% CPU availability.

Next, we look at the management VM's utilization during test. Similar to the MCS testing resource usage was extremely low.

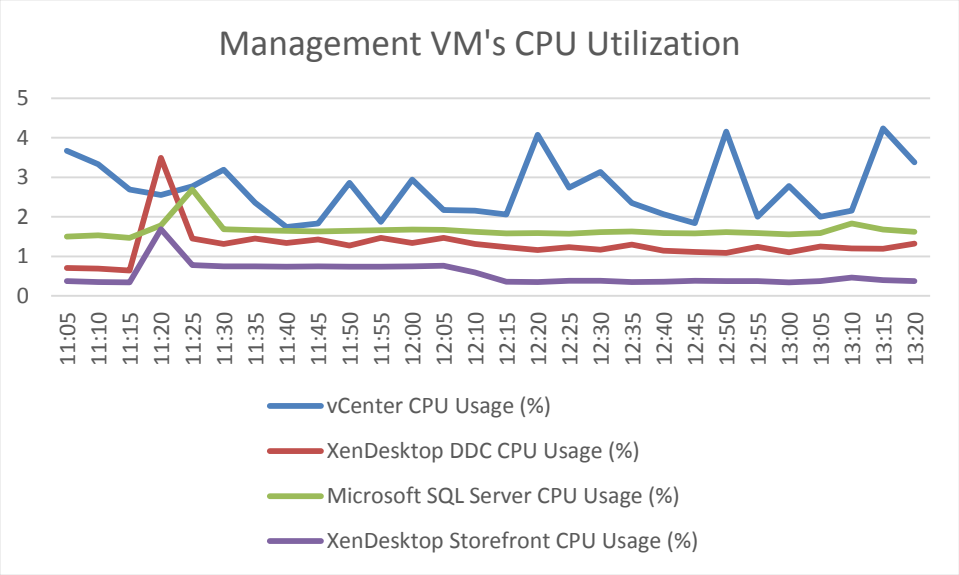

*** This does not account for the additional 10% of available CPU capacity provided by Turbo Boost.**

Below is the CPU utilization for the 12 XenApp VM's during the test cycle. None of the VM's exceeded 70% of their allocated resources.

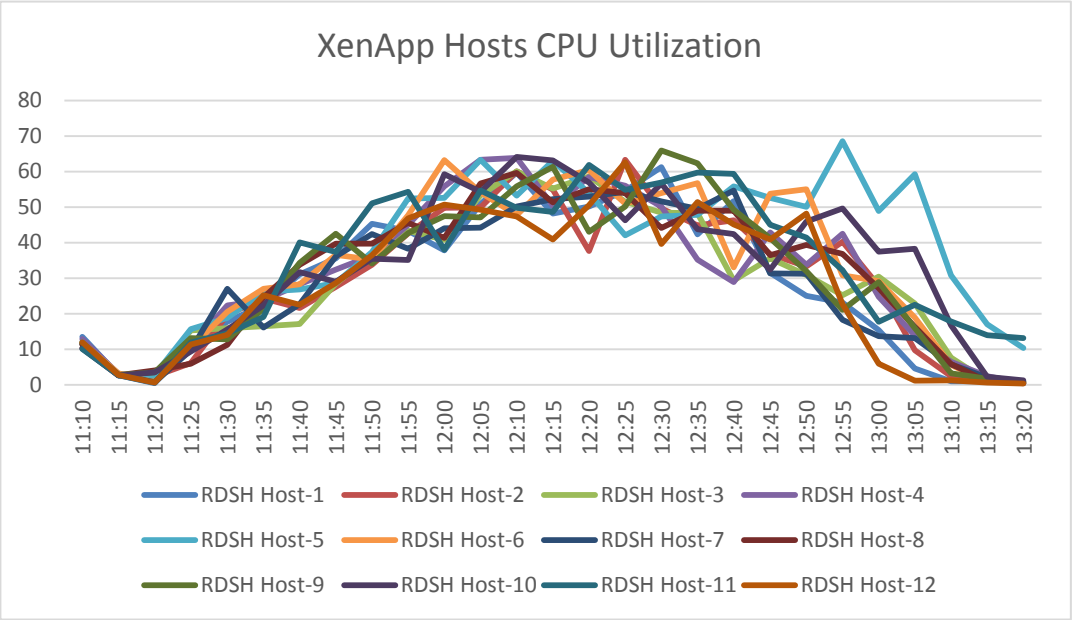

*** This does not account for the additional 10% of available CPU capacity provided by Turbo Boost.**

Below is the memory consumption for the management host. Less than 60GB of memory is consumed during the test run

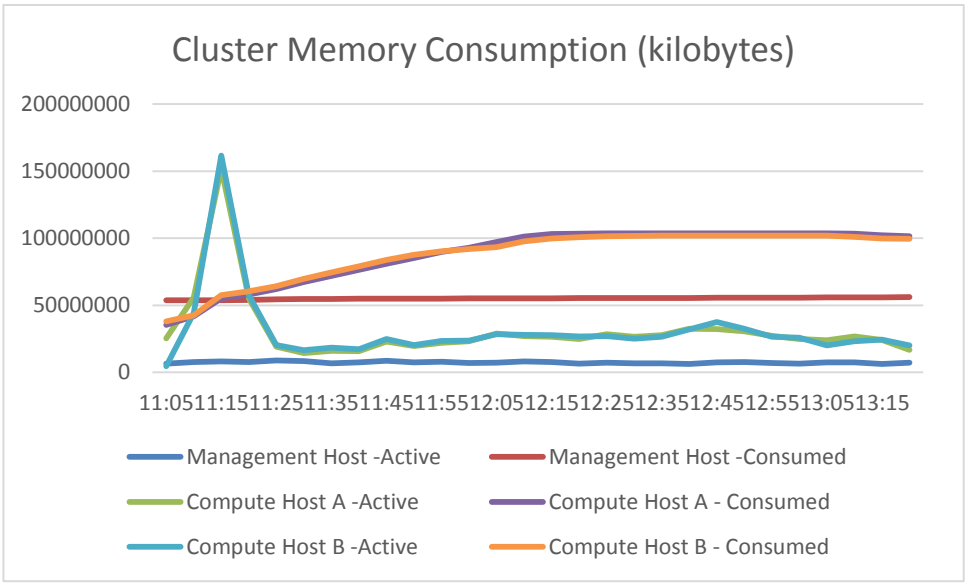

The above graphs show the memory consumption for the compute hosts in the cluster, at the beginning of the test there is a large memory activity before active memory drops right off for the remainder of the test cycle.

Next, we look at the Network utilization for the cluster. It was observed that compute host B registered a large network spike at the beginning of the test; this was due to a reboot of the host. At steady state of the test network consumption was low at less than 1GB for all three hosts.

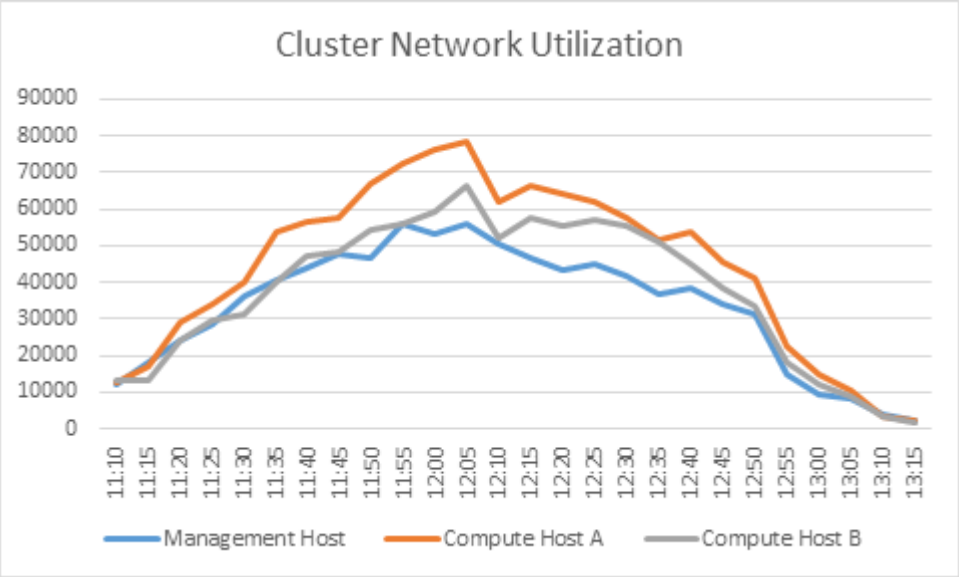

We also measure the user experience for each test run. LoginVSI was used for both load generation and to measure the user experience. The VSI max score of 180 suggests that it is at the upper density that should be specified for this host.

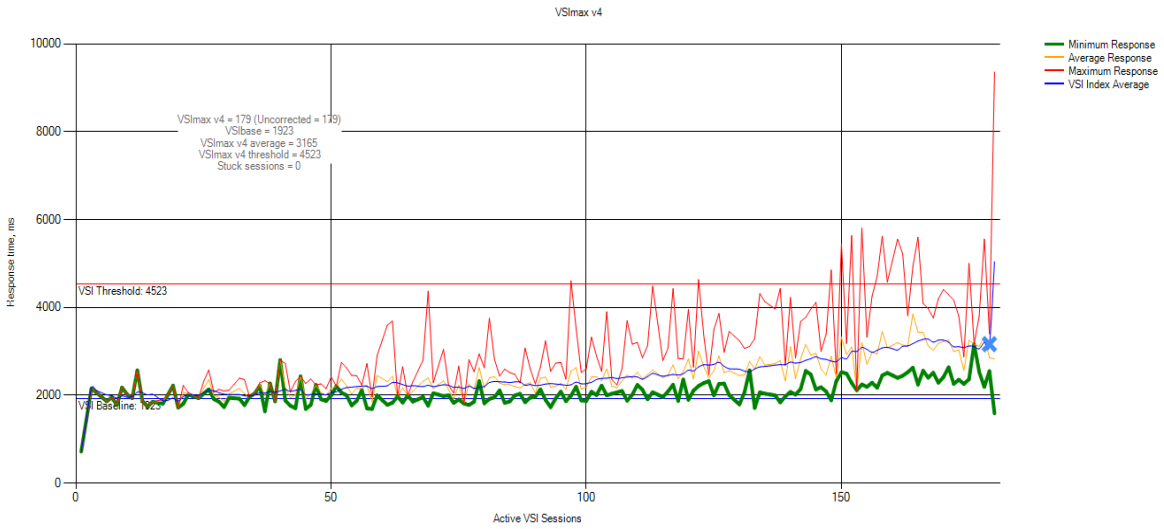

The storage platform for this testing was the Dell/Nutanix XC appliance, below are some of the graphs showing the CPU, IOPS, Network and Memory consumption for the cluster.

The cluster CPU utilization measured on the Nutanix management appliances shows as maximum of 67% during the test run.

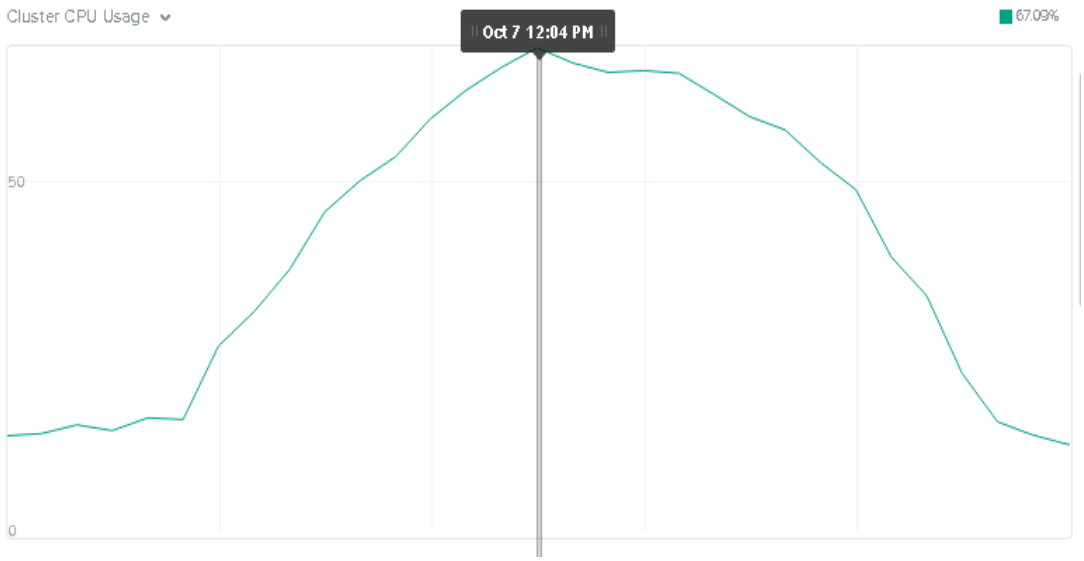

The below graph shows the IOPS per host at the highest point of test. This is when all sessions have just logged in.

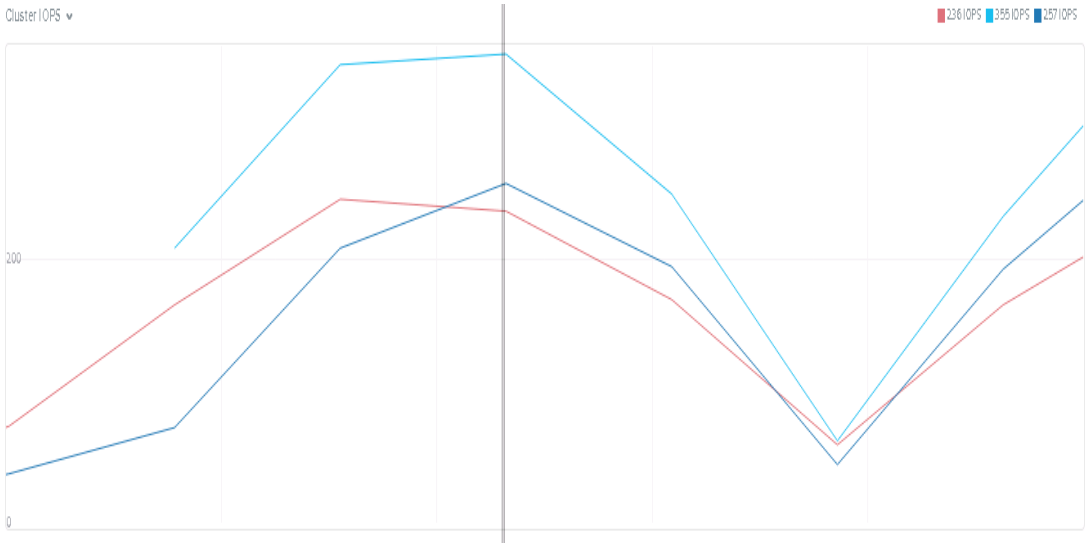

In summary, the numbers are in line with expected figures for the 2650 v2 Processor, CPU continues to be the bottleneck for VDI solutions. Based on the resource consumption for the test runs Memory and network are completely under-utilized. User experience was excellent throughout the test runs with no jitter or delay observed on the Video content.

The main observation is the resting CPU consumption of the physical host CPU, If DRS, HA and vMotion were used then a lower density would be expected. The additional 10% CPU that Turbo boost provides gives some additional headspace for sessions.

6.6.2 High Workload 250 XenApp Users on XenDesktop

The XenApp Session hosts were brokered by Citrix XenDesktop MCS 7.5 with a dedicated management host and two Compute hosts. Each compute host had a total of 6 VM's running Server2012 RDS. Each host had a total of 90 sessions running during testing. The standard LoginVSI application packages were installed and the Login VSI Medium workload was used for load generation.

The graph below shows the CPU utilization for the Management, Compute hosts and the XenApp VM's.

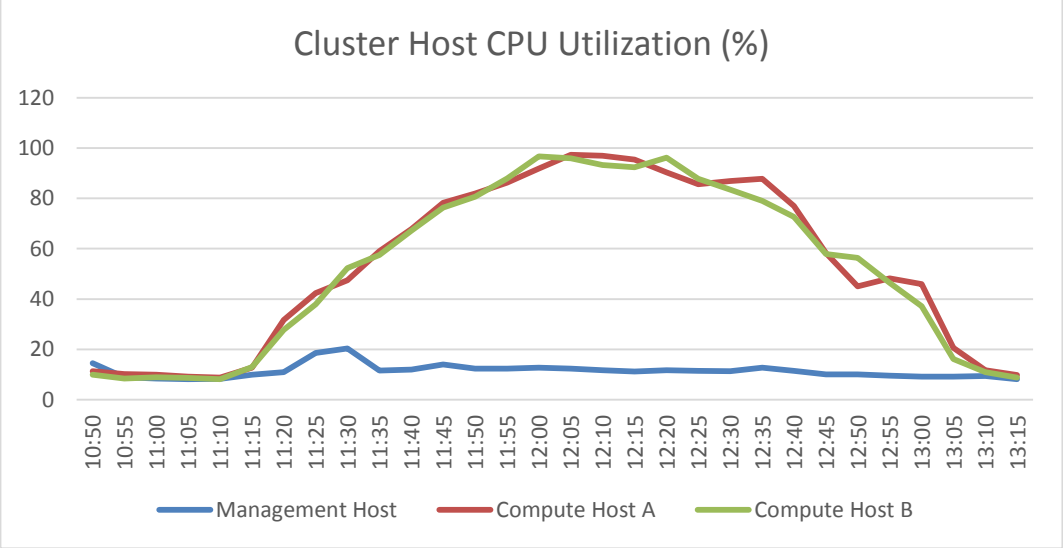

*** This does not account for the additional 10% of available CPU capacity provided by Turbo Boost.**

Above are the CPU performance charts for both Compute hosts with neither host exceeding 95% during the test. These figures do not include the Turbo boost additional 10% CPU available.

Next, we look at the performance Metrics for the Management VM's during this test. As expected there was a low level of I/O and resource generation for the XenApp test.

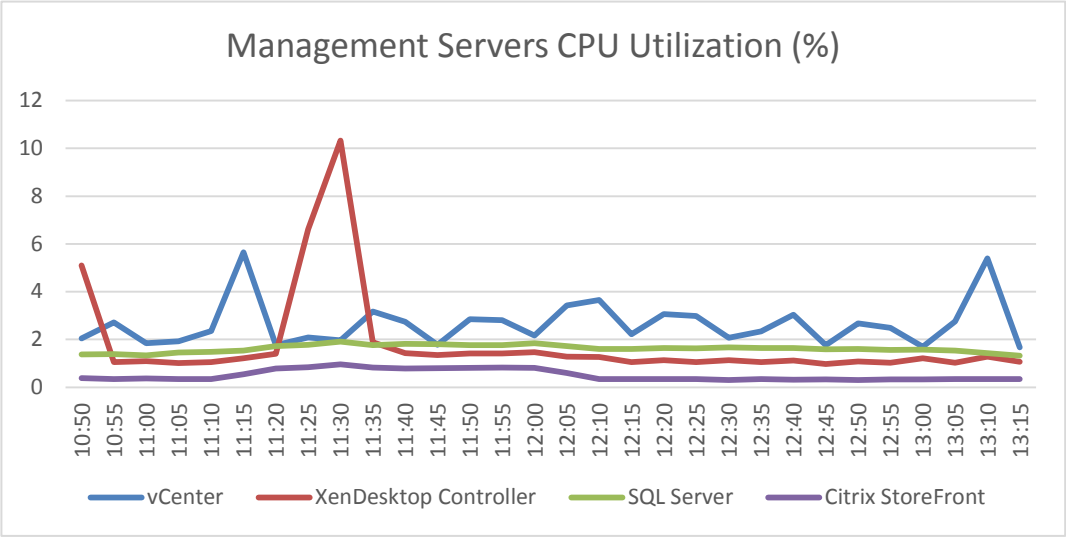

*** This does not account for the additional 10% of available CPU capacity provided by Turbo Boost.**

Below is a combined graph for the XenApp hosts showing the CPU utilization for each of these virtual machines. Average CPU did not exceed 70% across each host however the VSI Max score was close to the limit.

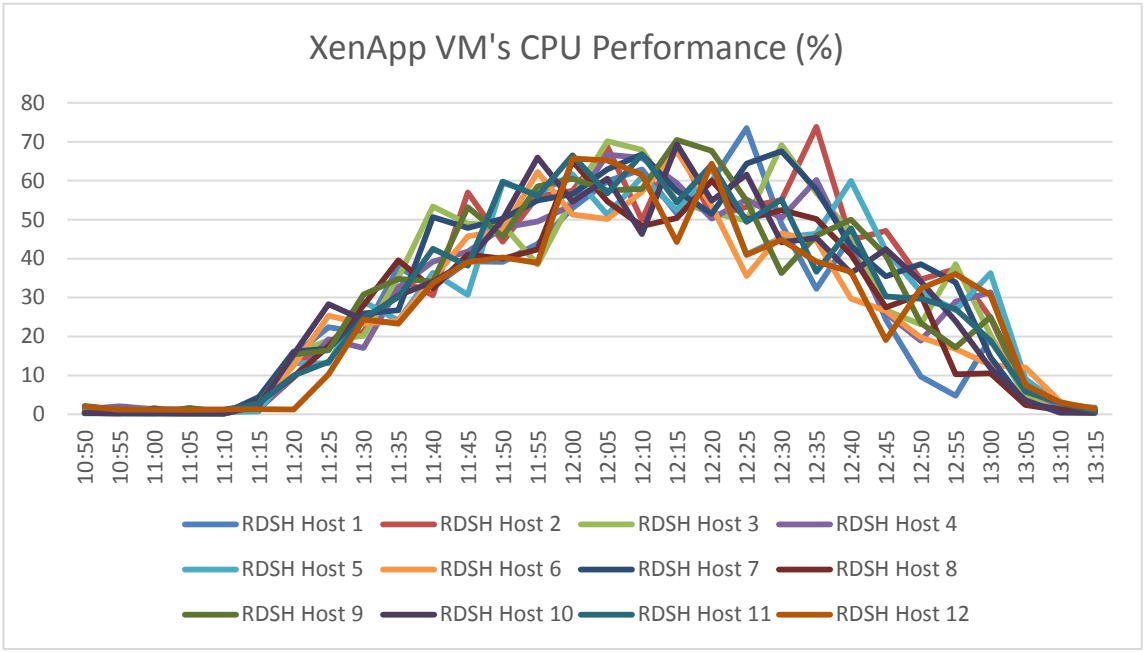

*** This does not account for the additional 10% of available CPU capacity provided by Turbo Boost.**

Memory utilization again was as expected with Active Memory extremely low suggesting that there is an over allocation of memory on this configuration

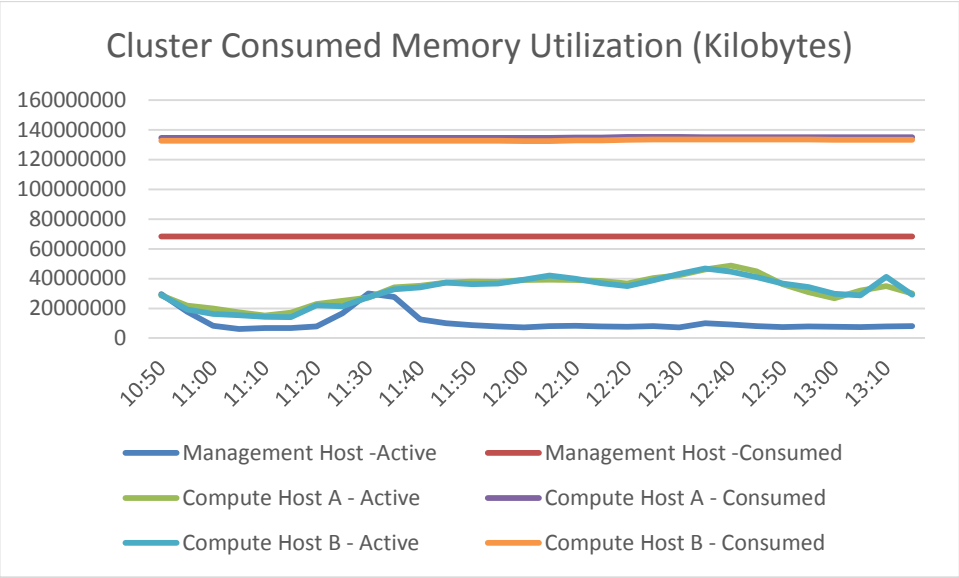

Above is the memory consumption for the management host. This host was not processing any sessions as part of the tests.

Now we look at network utilization which was less than 1Gb throughout the test with both compute hosts generating the most traffic.

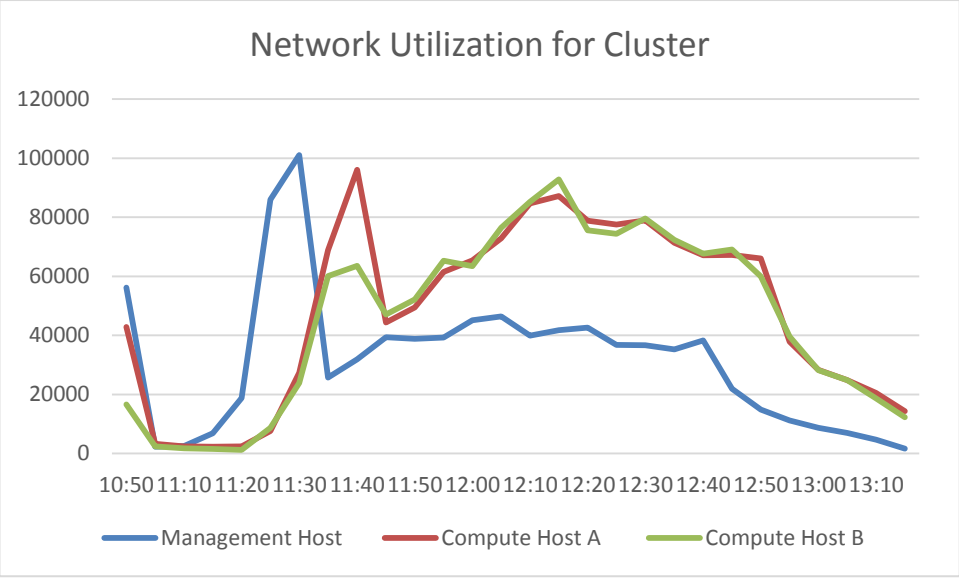

Next, we look at the VSI Max Score for this test. The scoring gives a good indication that 250 users is the maximum number of sessions that should be run

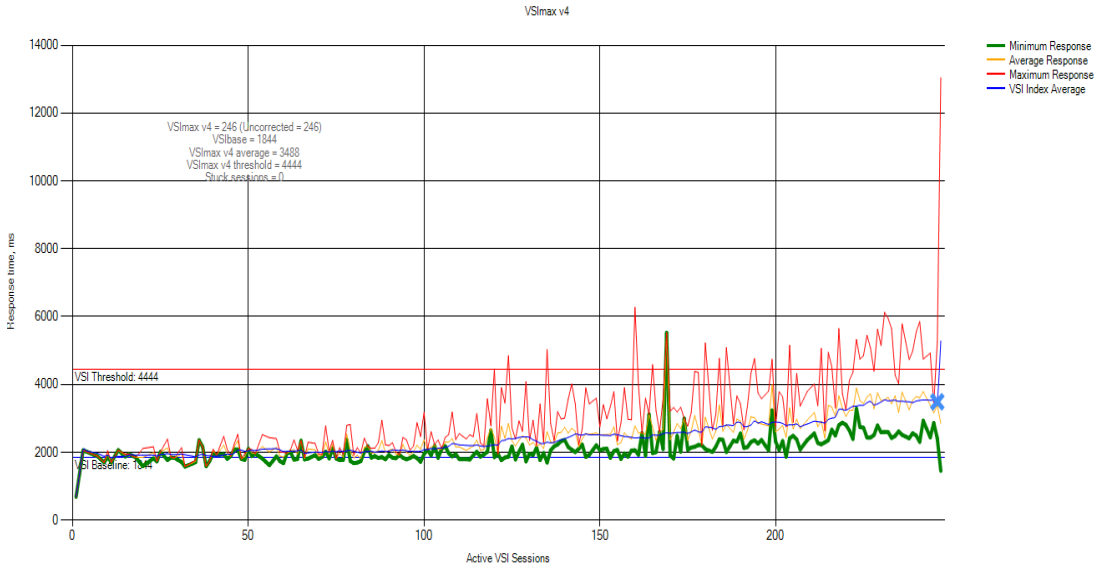

The storage platform for this testing was the Dell/Nutanix XC appliance, below are some of the graphs showing the CPU, IOPS, Network and Memory consumption for the cluster.

On the Prism interface the observed CPU for the cluster was 66%CPU.

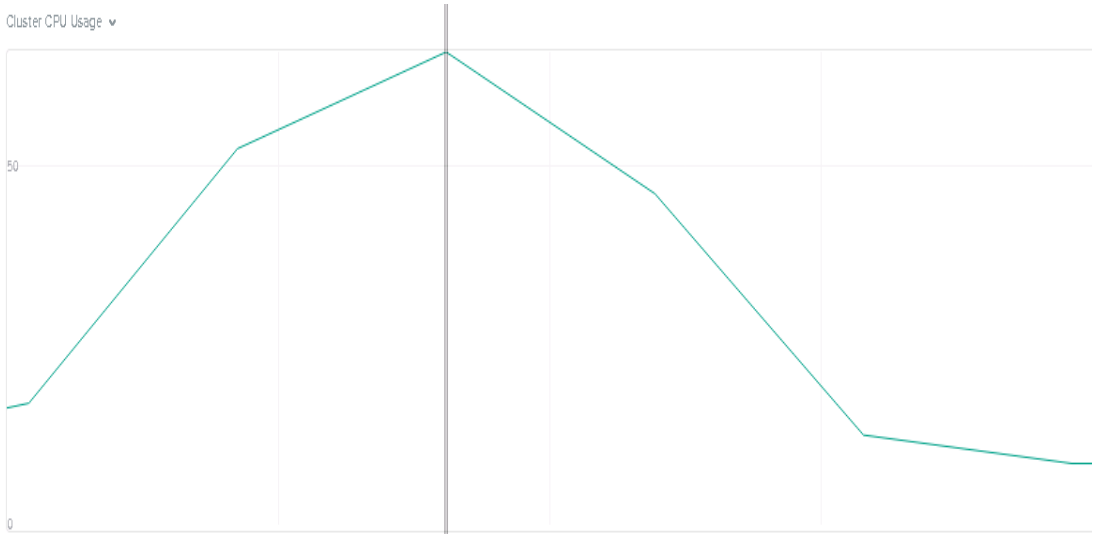

The cluster IOPS again were low as expected with a low level of latency observed on the storage appliance.

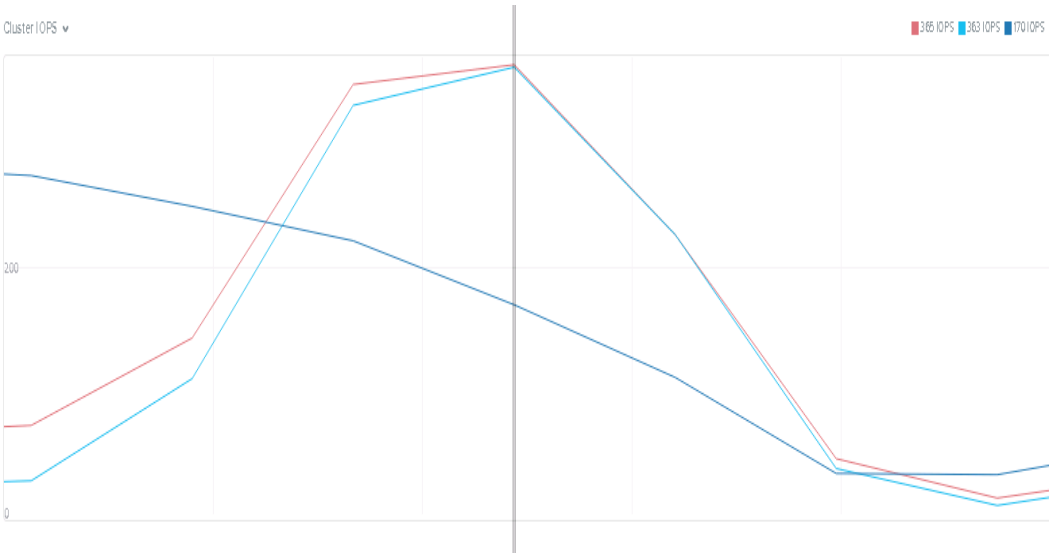

In summary, the bottleneck with the 250 XenApp users is Compute host CPU as expected. The Nutanix cluster performed well with low IOPS consumption and only 66% of the cluster CPU utilized. The user experience for each session was excellent with all sessions logging on and off during the test. The additional 10% of CPU provided by Turbo boost on each host gives some additional headroom for more sessions. Log on for users was fast and video quality within each session monitored was excellent.

6.7 Test Result Summary

- The CVM on each node consumes 20% of the CPU limiting the number of desktops that can be accommodated.
- Prebooting is recommended as the CVM CPU does not handle I/O storms adequately.
- The Cluster CPU, I/O and Memory performance was excellent.
- Based on utilization there is an over allocation of memory.
- The dedicated host model limits the number of users in a 3 node cluster making it more attractive to design the solution as a 4 node cluster.

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.

Jerry Van Blaricom is a Systems Principal Engineer in the Desktop Virtualization Solutions 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.

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 Citrix-based solutions at Dell. Before writing about technology, Manish has spent time designing, implementing and supporting technology- in IT, Systems Engineering & Network Performance & Monitoring. Manish has been a long-time Dell customer & Advocate before becoming a Dell employee.