Dell DVS Enterprise for Windows Server 2012 RDS – Reference Architecture

7/18/2013
Phase 2
Version 1.6

Table of Contents

1 Introduction

1.1 Purpose of this document

This document describes:

1. Dell DVS Reference Architecture for Windows Server 2012, scaling from 50 to 750 VDI users or up to 900 RDSH sessions.
2. A VDI Experience Proof of Concept (POC) Solution, an entry level configuration supporting 10 VDI users.
3. A remote office/ branch office solution supporting 250-500 pooled or session desktops.

This document addresses the architecture design, configuration and implementation considerations for the key components of the architecture required to deliver virtual desktops via RDS on Windows Server Hyper-V 2012.

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 design 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 Physical Architecture Overview

The core architecture design consists of the Local Tier1 solution model. "Tier 1" in the DVS context defines from which disk source the VDI sessions execute. Local Tier1 applies to rack servers only while Shared Tier 1 can be rack or blade. Tier 2 storage is utilized for user profile/data and Management VM execution.

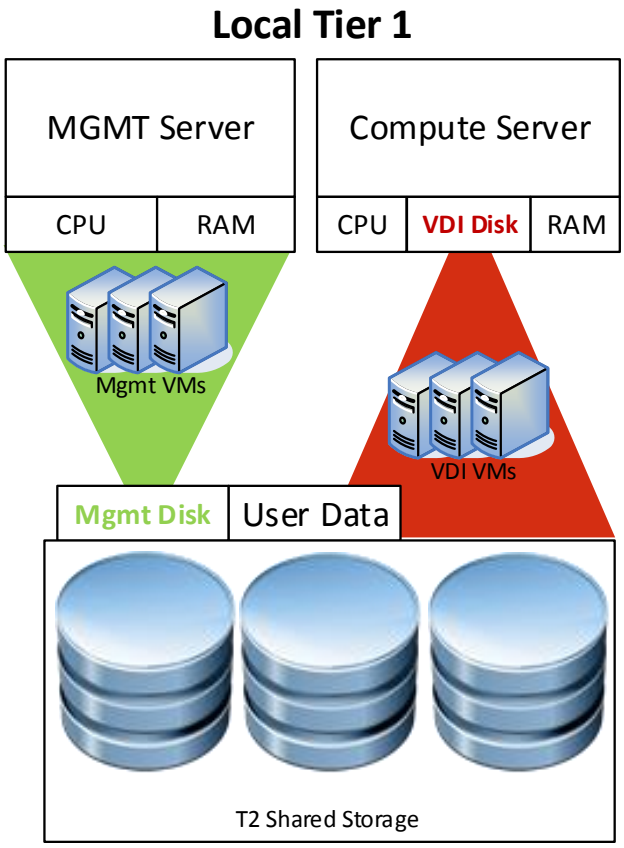

2.2 Layouts and Breakpoints

The solution architecture will follow the traditional DVS distributed design model consisting of 4 primary layers: Network, Compute, Management, and Storage. The Network and Storage layers can be optionally provided by the customer if suitable infrastructure is already in place. The Compute layer contains the hosts that serve the VDI sessions and the Management layer contains the components required to support the VDI broker infrastructure.

The following highlights the key layout and scaling elements of the solution.

2.3 Local Tier 1 – Solution Layers

Only a single high performance Force10 S55 48-port switch is required to get started in the Network layer. This switch will host all solution traffic consisting of 1Gb iSCSI and LAN sources.

Additional switches can be added and stacked as required to provide High Availability for the Network layer.

The Compute layer consists of the server resources responsible for hosting the user sessions, whether shared via RDSH (formerly Terminal Services) or pooled via RDVH (see section 4.5 for a detailed explanation of each role). The RDVH role requires Hyper-V as well as hardware assisted virtualization so must be installed into the parent partition of the Hyper-V instance. The RDSH role is enabled within dedicated VMs on the same or dedicated hosts in the Compute layer.

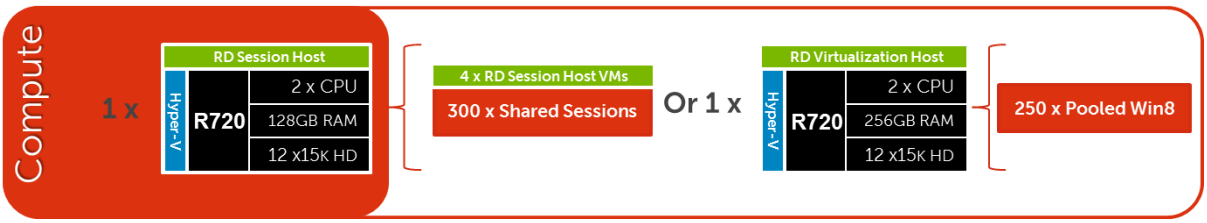

Management components are dedicated to their own layer so as to not negatively impact the user sessions running in the Compute layer. This physical separation of resources provides clean, linear, and predictable scaling without the need to reconfigure or move resources within the solution as you grow. The Management layer will host all the RDS VMs necessary to support the infrastructure as well as a file server to host SMB shares for user Profile Disks or data.

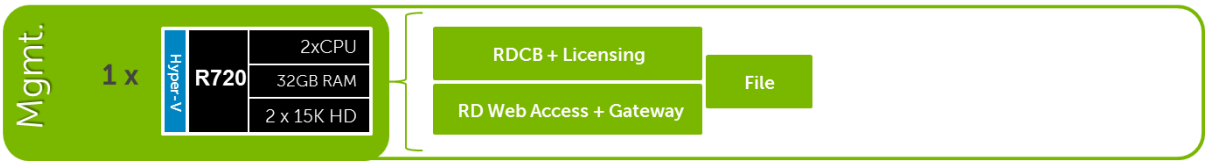

The Storage layer is made up by the capacity dense and performance capable EqualLogic 4100E iSCSI array. 12TB is provided in base form that can scale as high as 36TB to suit your capacity requirements. A second 4100E can be added to group the two arrays to provide greater capacity or tier 2 performance.

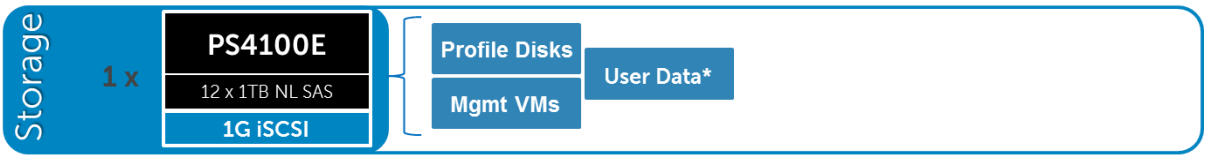

2.4 Dell DVS 10-Seat Trial Kit

To get up and running as quickly as possible with pooled VDI, Dell is offering an extremely affordable solution capable of supporting 10 concurrent users for a minimal investment. This architecture leverages an inexpensive single server platform intended to demonstrate the capabilities of VDI for a small environment or focused POC/ trial. Networking is provided optionally in this solution and all VDI roles/ sessions are hosted on a single server.

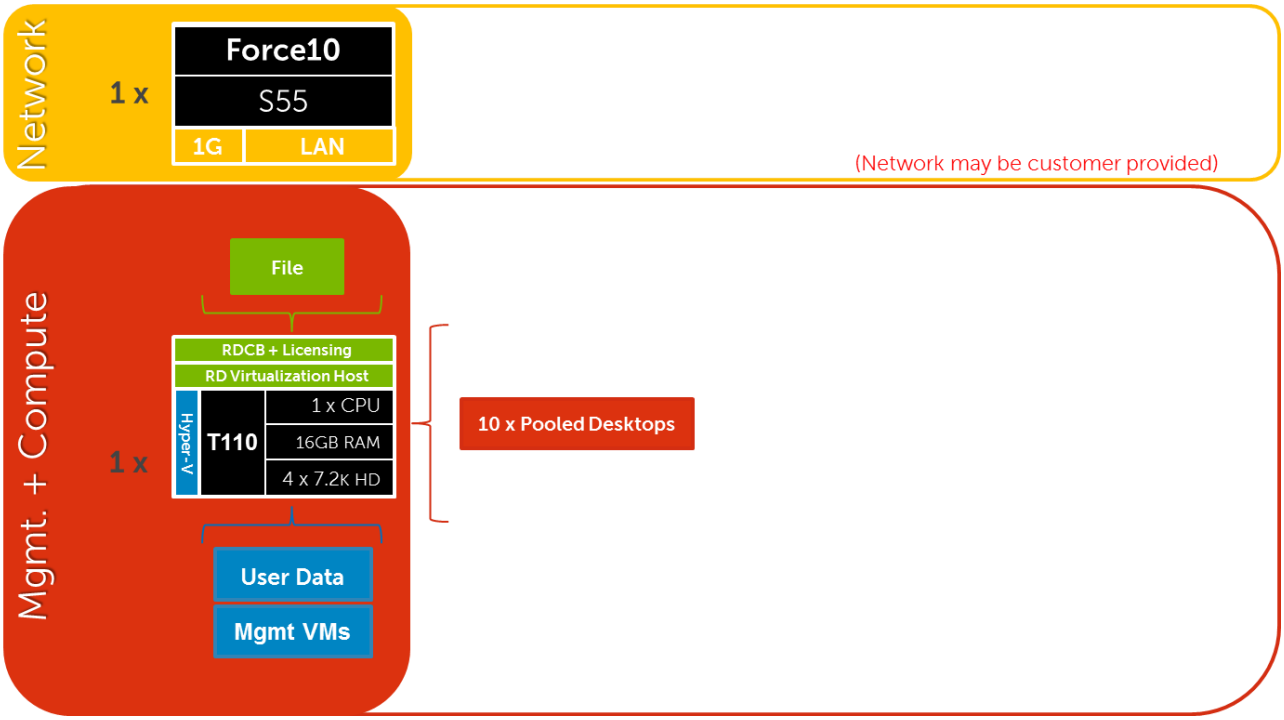

For more information on the 10-seat Trial Kit, please see Appendix A.

2.5 Local Tier 1 – 250 User POC

For small scale deployments or pilot efforts intended to familiarize your organization with the Dell enterprise VDI, we offer a 200 user combined pilot solution. This architecture is non-distributed following a SAN-less model with all VDI, management and storage functions running on a single host. If additional scaling is desired, you can grow into a larger distributed architecture seamlessly with no loss on initial investment.

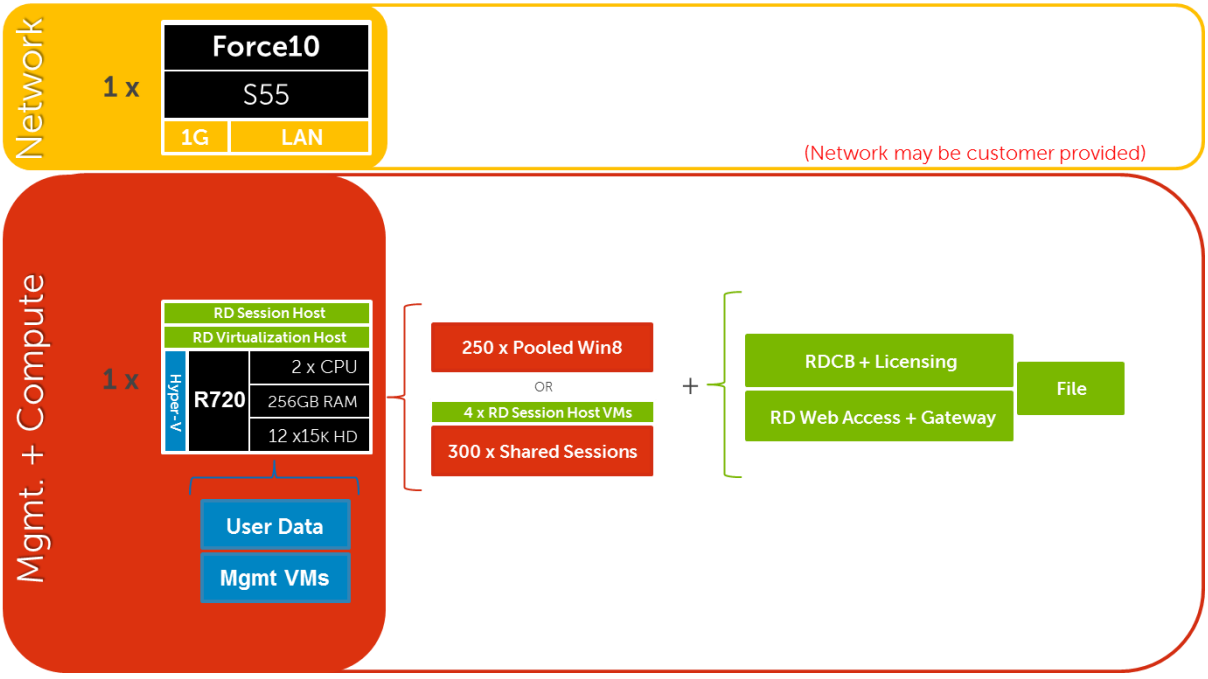

2.6 Local Tier 1 – 250 User Pilot (distributed)

In addition to the 250 user POC combined offering we also offer a scale ready version that includes Tier 2 storage. The basic architecture is the same but customers looking to scale out quickly will benefit by building out into Tier 2 storage initially.

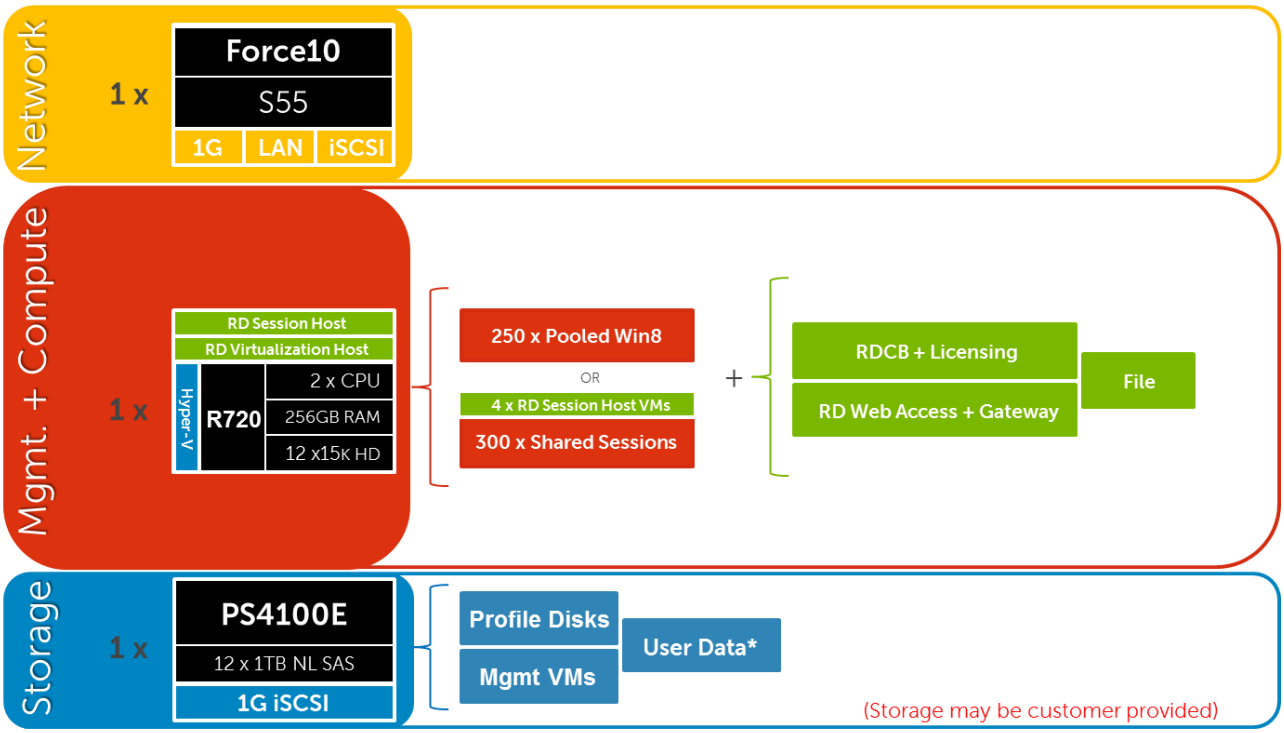

2.7 Local Tier 1 – Combined

As a logical entry point to the distributed RDS solution stack, a combined architecture is offered to host both the RD Virtualization Host (RDVH) and RD Session Host (RDSH) roles within the same physical Compute host while separating the Management layer. This will enable users requiring either shared RDP or pooled VDI sessions to be hosted on the same physical server. The value of this solution is a minimum infrastructure investment with maximum VDI flexibility easily tailored to shared and pooled user types. Horizontal scaling is achieved simply by adding additional Compute hosts. Additional information on the hardware components can be found in section 3 below.

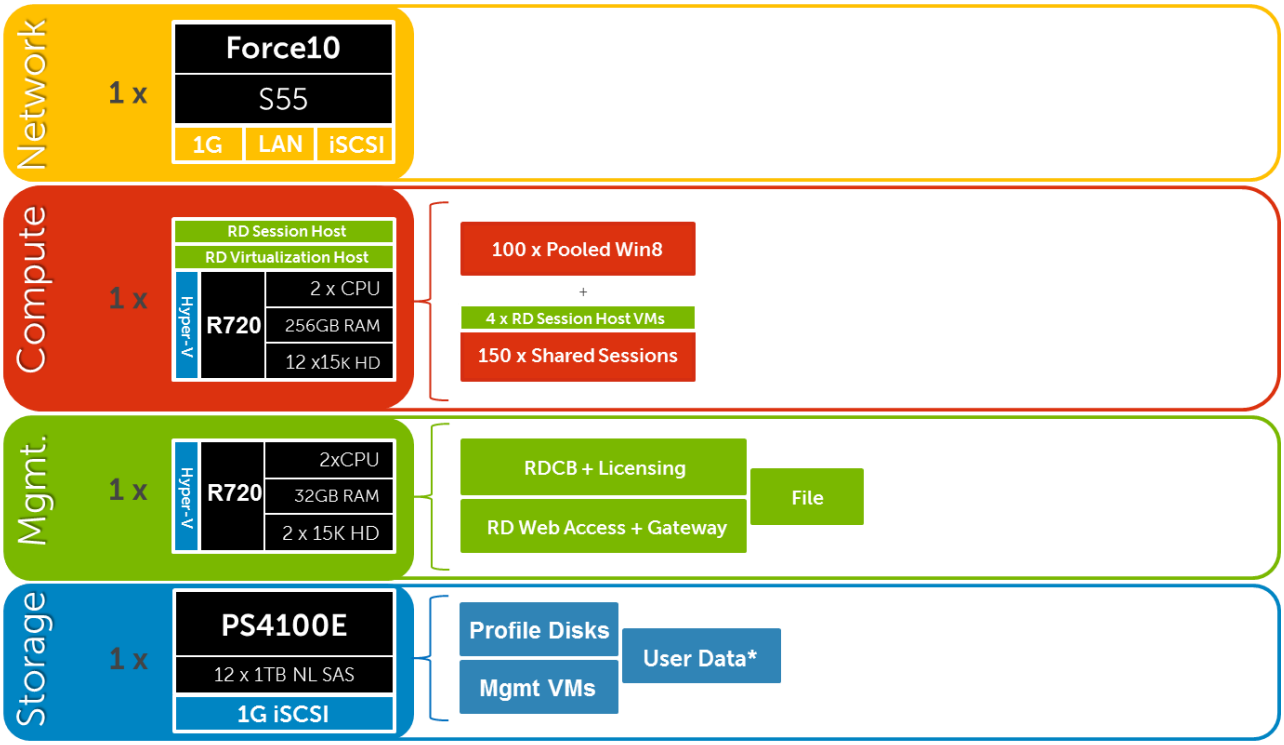

2.8 Local Tier 1 – Base

In the base distributed architecture the RDVH or RDSH roles are assigned to a dedicated Compute host. This architecture can support either a single RDVH or RDSH Compute host or one of each. This solution provides maximum Compute host user density for each broker model and allows clean linear upward scaling. You'll notice that the hardware spec is slightly different for the two Compute host types, giving additional RAM to the virtualization host. This of course can be adjusted to suit your specific needs. Additional information on the hardware components can be found in section 3 below.

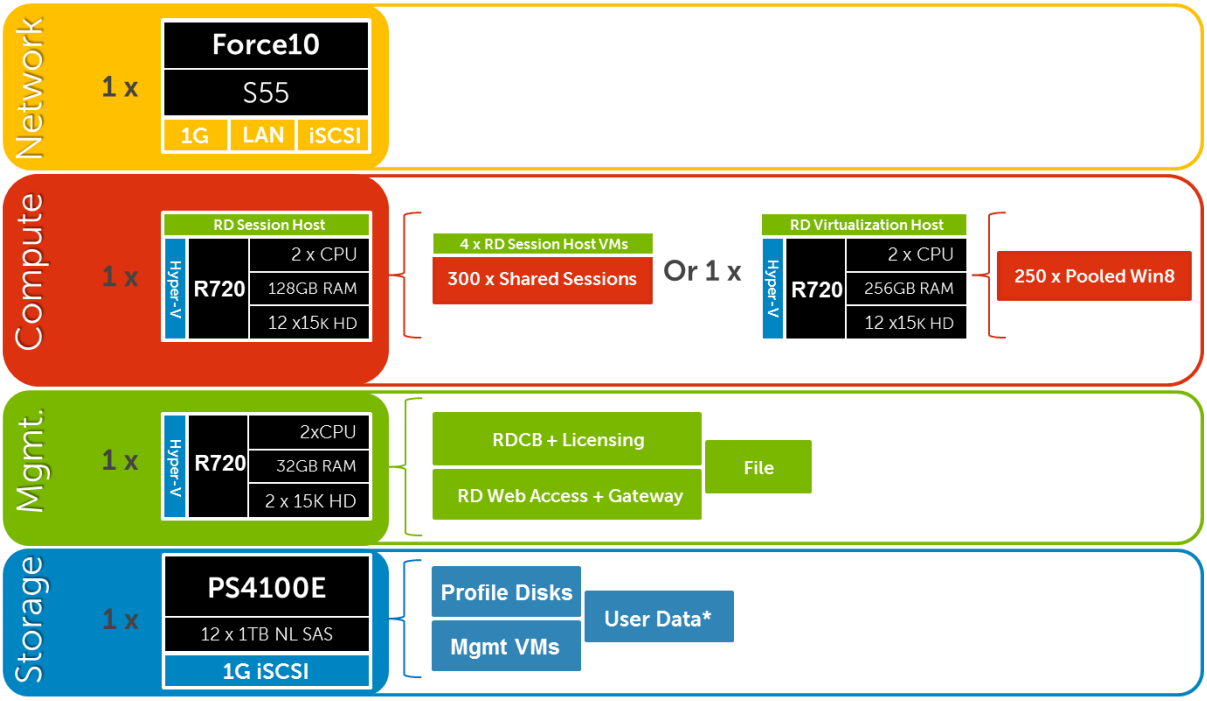

2.9 Local Tier 1 – Fully Expanded

The fully expanded architecture provides linear upward scale for both the RDVH and RDSH roles optimized for 750 pooled VDI sessions or 900 shared. This solution supports up to 3 Compute hosts of any combination running either RDVH or RDSH roles to meet the needs of the enterprise. Additional information on the hardware components can be found in section 3 below. The overall infrastructure is capable of supporting a much higher number of sessions but additional management infrastructure will be required. The native RDS management infrastructure was designed by Microsoft to support in the neighborhood of 600 sessions so the scale of this architecture is limited accordingly.

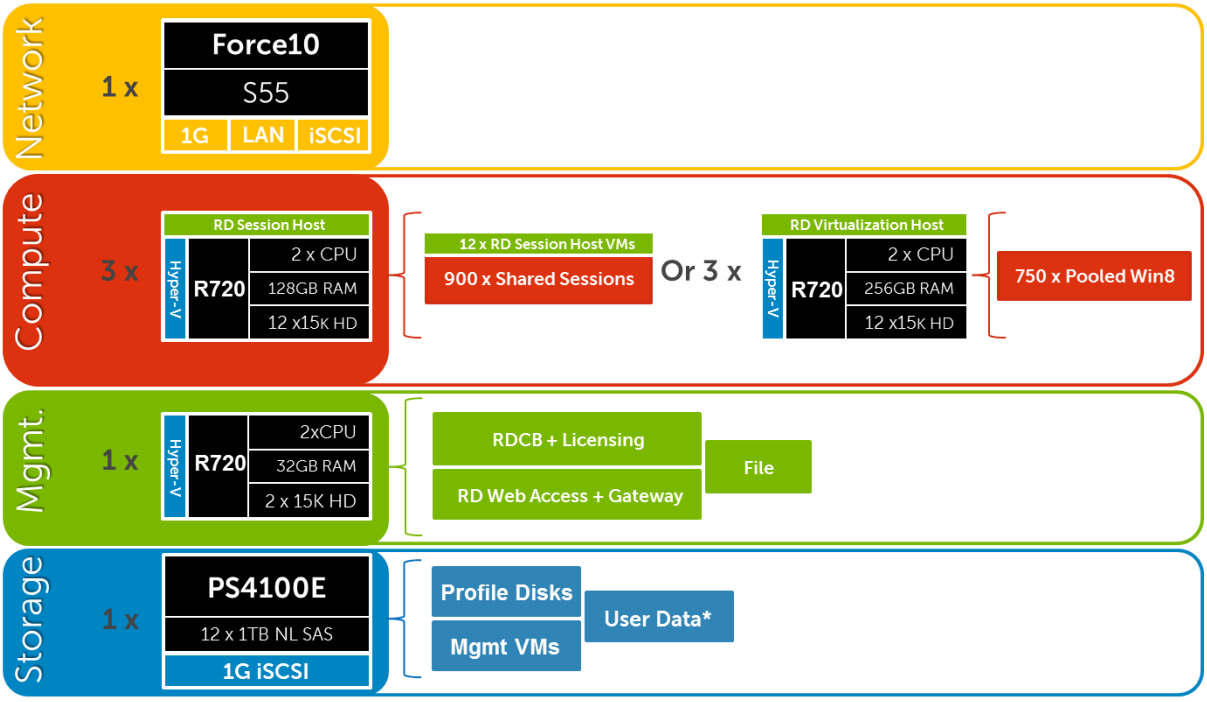

2.10 Local Tier 1 – Fully expanded (Graphics acceleration only)

For solutions requiring a higher amount of graphical acceleration, additional GPU-based hosts can extend the overall scale of the compute layer. In keeping with the limitations of the native RDS management infrastructure, up to 700 graphical user desktops can be supported on 7 compute hosts. Please see section 3.2.2 for more information on the graphics accelerated compute hosts.

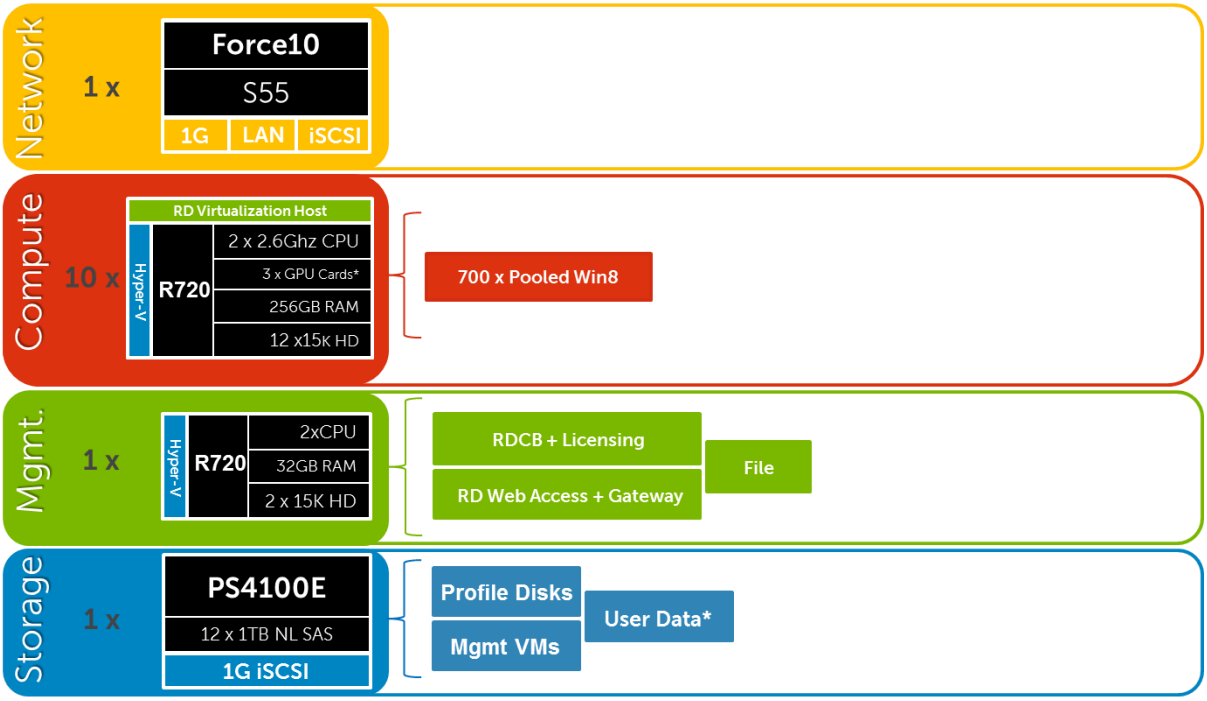

2.11 High Availability

High availability (HA) is currently offered to protect all layers of the solution architecture. An additional ToR switch is added to the Network layer and stacked to provide redundancy, additional Compute and Mgmt hosts are added to their respective layers, and Hyper-V clustering is introduced in the Management layer. Please see section 4.8 for more detailed information regarding HA in this architecture.

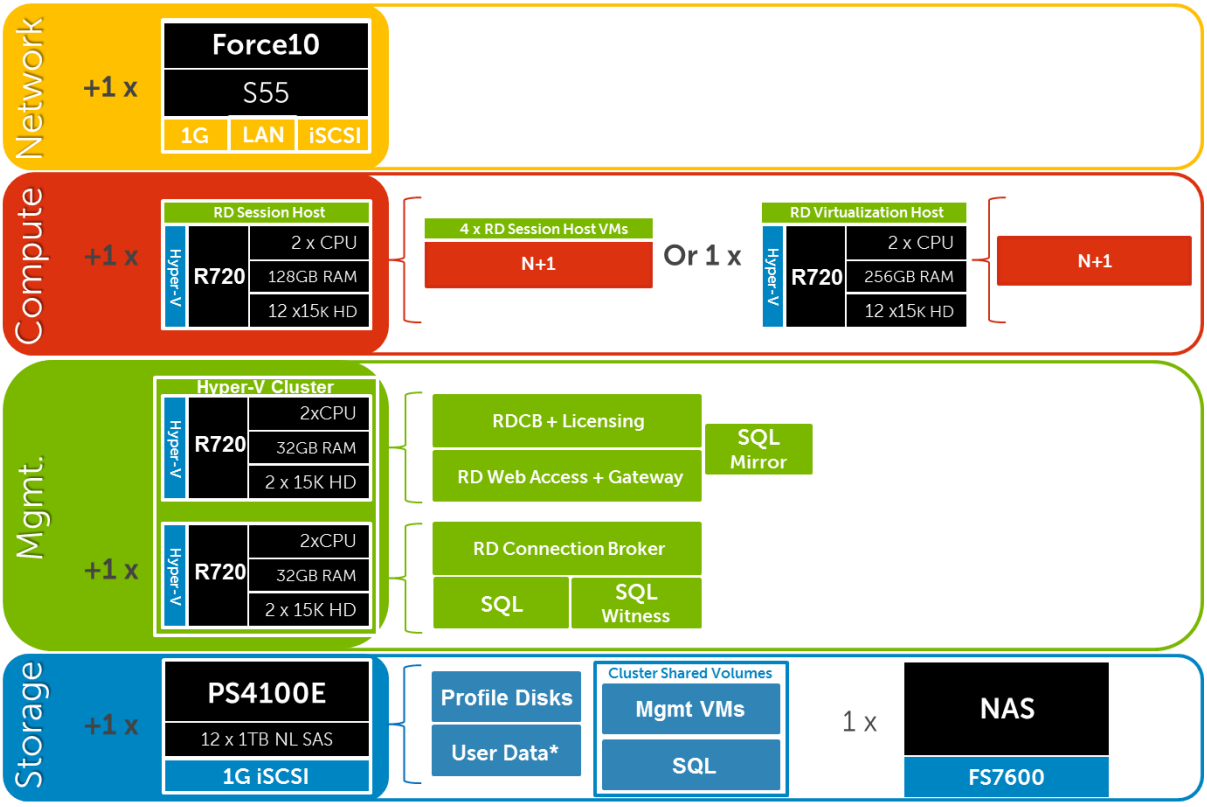

2.12 Local Tier 1 – Network Architecture

In the Local Tier 1 architecture, a single Force10 S55 switch can be shared among all network connections for both Management and Compute layer components, for the upper limit of the stack. Only the Management servers connect to iSCSI shared storage in this model. All ToR traffic has been designed to be layer 2/ switched locally, with all layer 3/ routable VLANs trunked from a core or distribution switch. The following diagrams illustrate the logical data flow in relation to the core switch.

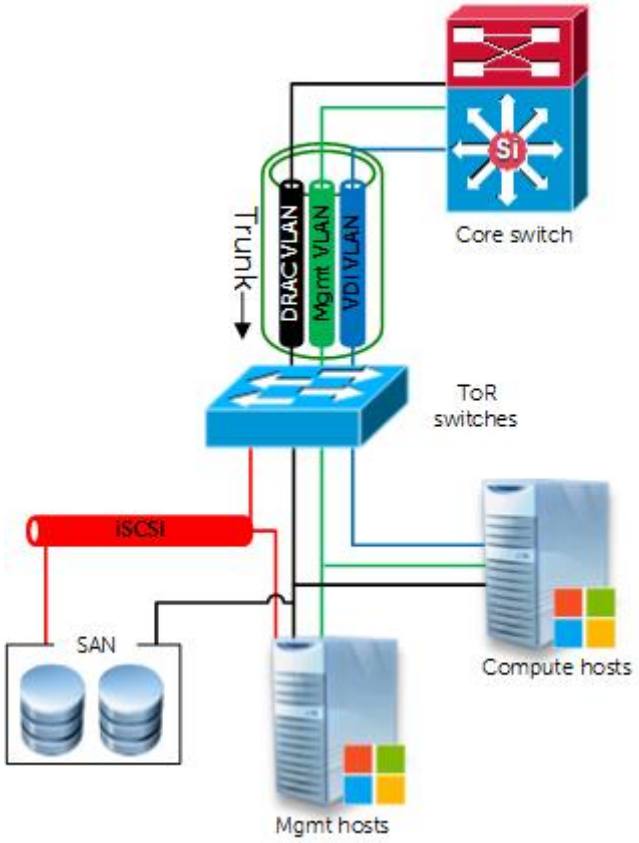

2.12.1 Cabling Diagram

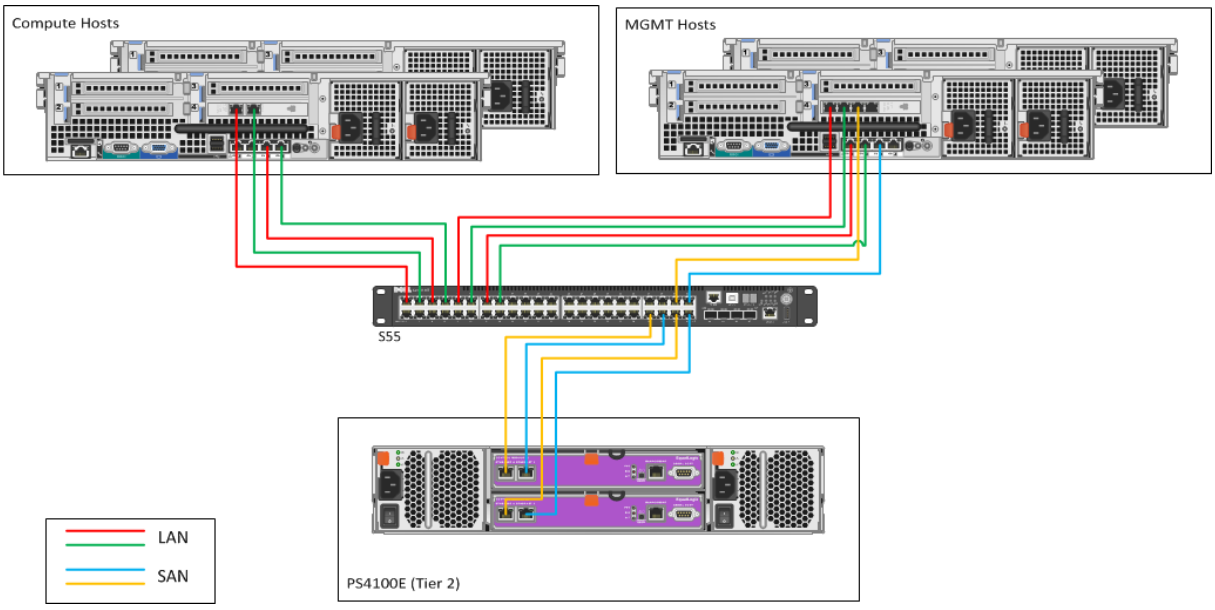

2.13 Local Tier 1 – Active Infrastructure

The solution architecture can also be configured using Dell Active Infrastructure methodology if desired. This architecture consists of a pre-integrated pod approach for scaling up to 400 pooled users. 1Gb switching for LAN and iSCSI are separated and stacked, 2 x 1U compute hosts as well as a single smaller management host are included. The EqualLogic array is also upgraded to provide more performance and storage capacity. The PC7024 switch stack used for LAN traffic is also used for out-of-band management traffic. RDS management roles and the use of Tier2 shared storage persist between the solutions.

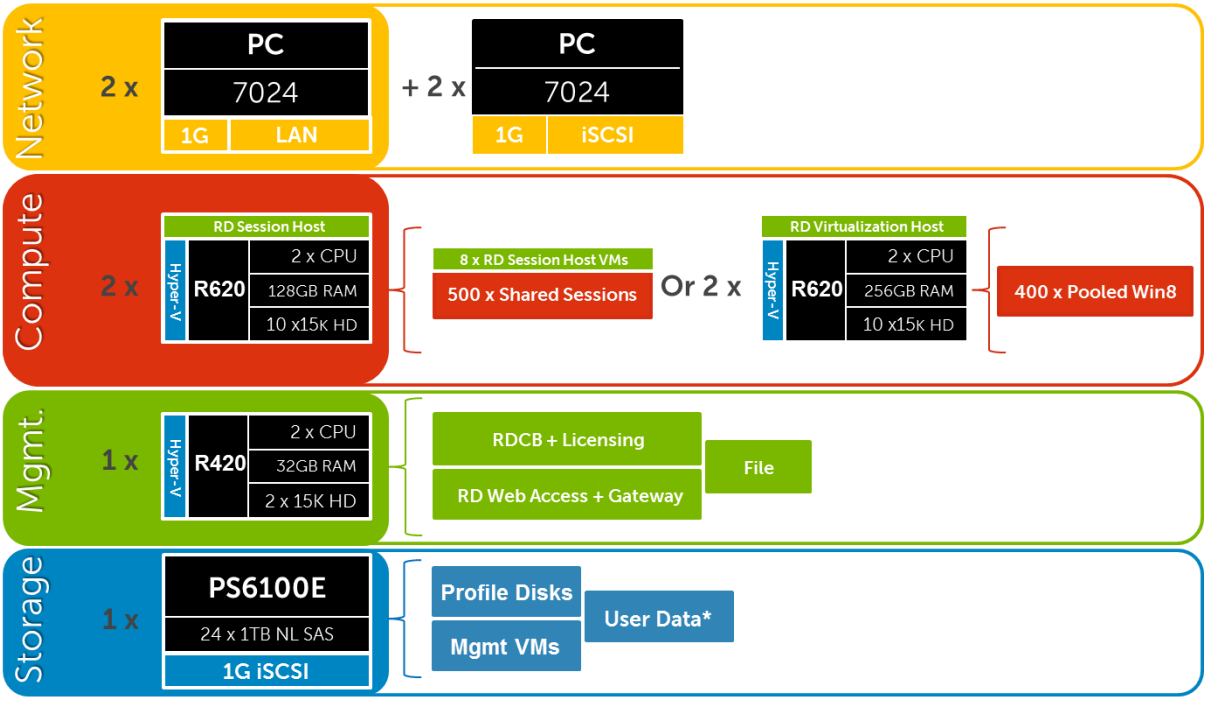

2.14 Remote Office Branch Office (ROBO)

For remote or branch office deployment scenarios, Dell DVS offers a 2 or 4 blade plus 25-disk Direct Attached Storage (DAS) solution, all contained within a single 5U chassis. All switching, compute, management, and storage is included. This solution can support up to 500 pooled VDI users in an incredibly efficient, small, and cost effective platform. Additional ToR switching is available if required.

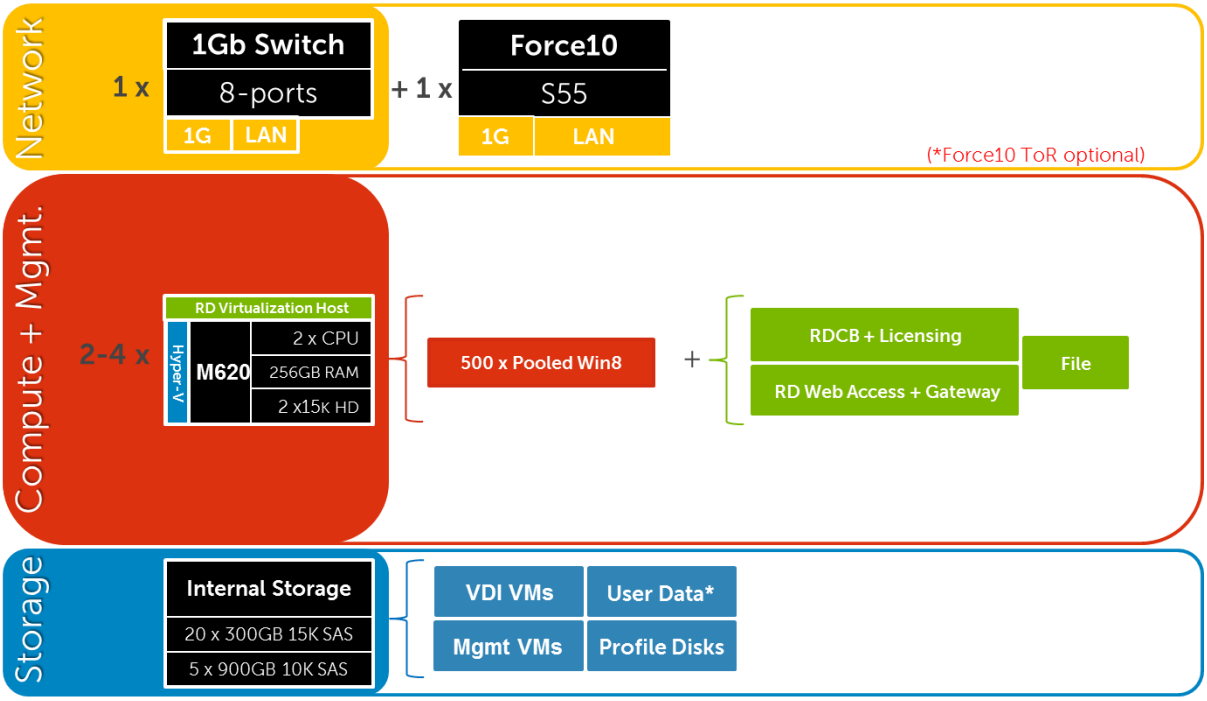

For more information on the ROBO solution offering, please see Appendix C.

2.14.1 Solution Density Summary

Design Scale	Management Hosts	Compute Hosts	RDSH Sessions		RDVH Sessions	HA
10 User POC	0	1	0		10	-
250 User Pilot	0	1	300	*or*	250	-
Combined	1	1	150	*and*	100	+ 1 Compute +1 Mgmt
Base Distributed	1	1	300	*or*	250	+ 1 Compute +1 Mgmt
Expanded	1	2	600	*or*	500	+ 1 Compute +1 Mgmt
Fully Expanded	1	3	900	*or*	750	+ 1 Compute +1 Mgmt
Fully Expanded (GPU)	1	10	0		700	+1 Compute +1 Mgmt
Active Infrastructure	1	2	500	*or*	400	-
ROBO	2 or 4 (shared)		500	*or*	250-500	-

3 Hardware Components

3.1 Network

The following sections contain the core network components for the DVS local and shared Tier 1 solutions. General cabling guidance to consider in all cases is that TwinAx is very cost effective for short 10Gb runs and for longer runs fiber with SFPs should be used.

3.1.1 Force10 S55 (ToR Switch)

Model	Features	Options	Uses
Force10 S55	44 x BaseT (10/100/1000) + 4 x SFP	Redundant PSUs	ToR switch for LAN and iSCSI in Local Tier 1 solution
		4 x 1Gb SFP ports the support copper or fiber	
		12Gb or 24Gb stacking (up to 8 switches)	
		2 x modular slots for 10Gb uplinks or stacking modules	

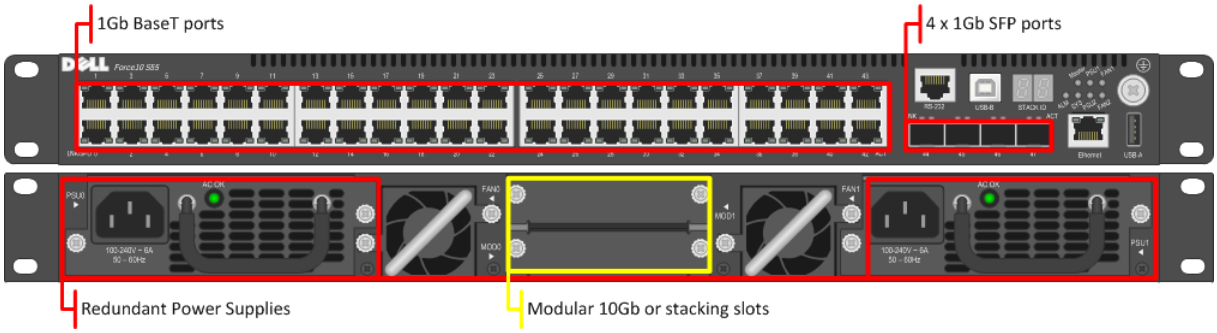

Guidance:

- 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 can be used.
- The front 4 SFP ports can support copper cabling and can be upgraded to optical if a longer run is needed.

For more information on the S55 switch and Dell Force10 networking, please visit: LINK

3.1.1.1 Force10 S55 Stacking

The Top of Rack switch in the Network layer can be optionally stacked with a second switch, 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. Switch stacks greater than 2 should be cabled in a ring configuration with the last switch in the stack cabled back to the first. Uplinks should be configured on all switches in the stack back to the core to provide redundancy and failure protection.

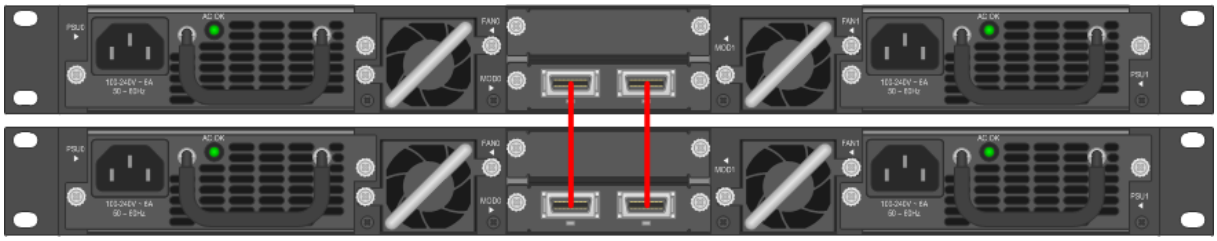

Please reference the following Force10 whitepaper for specifics on stacking best practices and configuration: LINK

3.2 Servers

3.2.1 Local Tier 1 Rack

The server platform for the Windows Server 2012 RDS enterprise solution is the best-in-class Dell PowerEdge R720. 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 16 2.5" SAS disks. Uncompromising performance and scalability in a 2U form factor.

In the local Tier 1 model, VDI sessions execute on the local storage of each Compute server. Due to the local disk requirement in the compute layer, this model supports rack servers only. In this model only the Management server hosts access shared storage to support the solution's Management role VMs. Because of this, the Compute and Management servers are configured with different add-on NICs to support their pertinent network fabric connection requirements. Refer to section 2.3.1 for cabling implications. The management server host has reduced RAM, CPU and fewer disks, since its VMs execute on shared Tier 2 storage. Both servers require a pair of hard disks configured in RAID1 to support the host operating system.

Local Tier 1 Compute Host – PowerEdge R720
2 x Intel Xeon E5-2690 Processor (2.9Ghz)
256GB Memory (16 x 16GB DIMMs @ 1600Mhz) (**RDVH**)
Or 128GB Memory (16 x 8GB DIMMs @ 1600Mhz) (**RDSH**)
Microsoft Windows Server 2012 Hyper-V
12 x 300GB SAS 6Gbps 15k Disks (OS + VDI)
PERC H710 Integrated 1GB RAID Controller – RAID10
Broadcom 5720 1Gb QP NDC (LAN)
Broadcom 5720 1Gb DP NIC (LAN)
iDRAC7 Enterprise w/ vFlash, 8GB SD
2 x 750W PSUs

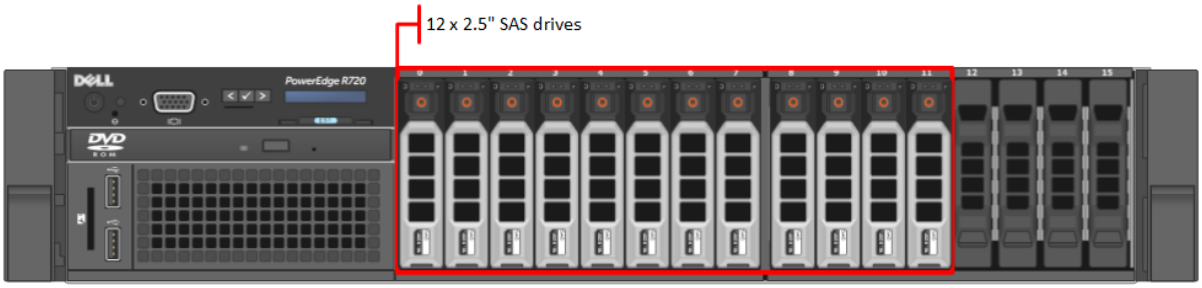

Local Tier 1 Management Host – PowerEdge R720
2 x Intel Xeon E5-2680 Processor (2.7Ghz)

32GB Memory (4 x 8GB DIMMs @ 1600Mhz)
Microsoft Windows Server 2012 Hyper-V
2 x 300GB SAS 6Gbps 15k Disks (OS)
PERC H710 Integrated 1GB RAID Controller – RAID1
Broadcom 5720 1Gb QP NDC (LAN/iSCSI)
Broadcom 5719 1Gb QP NIC (LAN/iSCSI)
iDRAC7 Enterprise w/ vFlash, 8GB SD
2 x 750W PSUs

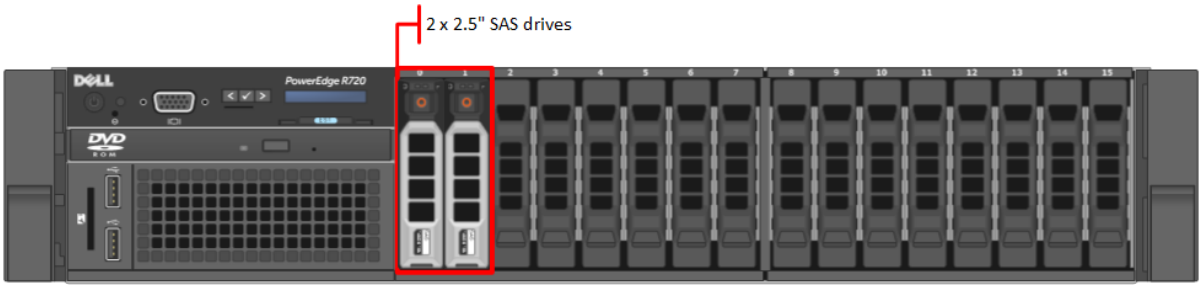

For more information on the Dell PowerEdge R720 server and other servers from Dell, please visit: LINK

3.2.2 Local Tier 1 Graphics Acceleration

Graphics acceleration is offered as an option for compute hosts to serve 100 shared GPU user desktops per server. This compute host configuration can be used in any RDVH scenario at any supported scale to meet these specific user requirements. As shown below, 3 of the AMD S7000 cards are required to achieve a 75 user per host density. Additional cards from both AMD and nVidia can be used depending on the workload and performance requirements.

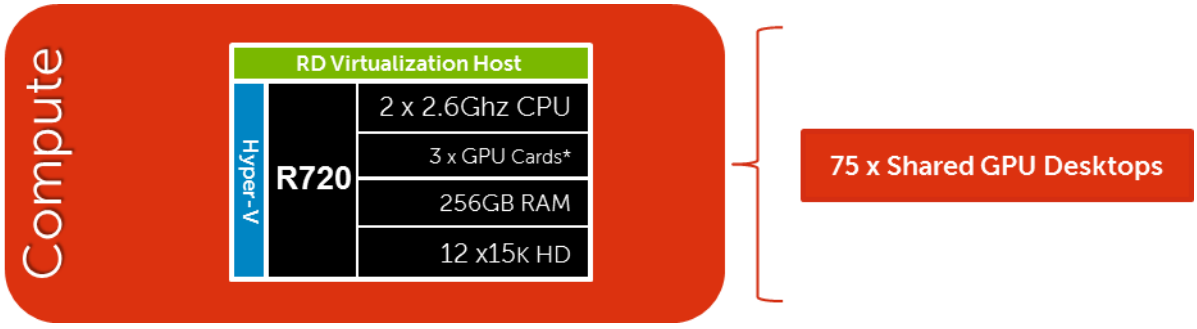

	AMD FirePro™ S7000	AMD FirePro S9000
GPU/ card	1 GPU	1 GPU
Cores	1280 Cores	1792 Cores
Memory/GPU	4GB	6GB

# of cards / R720 Server	3	2

Local Tier 1 Graphics Compute Host – PowerEdge R720
2 x Intel Xeon E5-2670 Processor (2.6Ghz)
256GB Memory (16 x 16GB DIMMs @ 1600Mhz) (**RDVH**)
Microsoft Windows Server 2012 Hyper-V
3 x S7000 or 2 x S9000 GPU Cards
12 x 300GB SAS 6Gbps 15k Disks (OS+VDI)
PERC H710 Integrated 1GB RAID Controller – RAID10
Broadcom 5720 1Gb QP NDC (LAN)
Broadcom 5720 1Gb DP NIC (LAN)
iDRAC7 Enterprise w/ vFlash, 8GB SD
2 x 1100W PSUs

3.2.2.1 AMD

AMD FirePro™ S7000 is the industry's most powerful single-slot server graphics processor available for powering Virtual Desktop Infrastructure, Workstation Graphics or Compute deployments in the data center. It features 2.4 TFLOPS of peak single-precision floating-point performance horsepower and is up to 2.3 times as fast as the competitive solutions in peak single-precision floating-point performance, due in part to its innovative all-new Graphics Core Next (GCN) architecture.

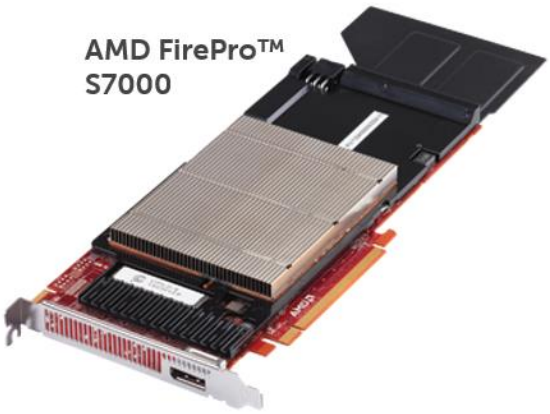

AMD FirePro™ S7000 server graphics are an ideal solution for VDI deployments. The S7000 supports Microsoft® RemoteFX, which allows it to support multiple concurrent users from the data center with local-like desktop experiences.

AMD FirePro™ S9000 is the industry's most powerful single GPU server graphics card available for powering Compute, Virtual Desktop Infrastructure and Workstation Graphics deployments in the data center. It features a blazing 3.23 TFLOPS of peak single-precision compute performance and is up to 2.4 times as fast as competitive solutions,1 due in part to its innovative all-new Graphics Core Next (GCN) architecture, which allows it to effortlessly process challenging compute workloads.

Equipped with a massive 6GB GDDR5 frame buffer, S9000 delivers up to 1.49 times the memory bandwidth of the competing solution2 — a scorching 264 GB/s. With high performance computing (HPC) features like error correcting code (ECC) memory and 806 GFLOPS of dual-precision floating-point performance, it's the ideal choice for HPC workflows.

3.2.2.2 GPU memory utilization for vGPU enabled VMs

Maximum Resolution	Maximum number of monitors in virtual machine setting			
	1 monitor	2 monitors	4 monitors	8 monitors
1024 x 768	48 MB	52 MB	58 MB	70 MB
1280 x 1024	80 MB	85 MB	95 MB	115 MB
1600 x 1200	120 MB	126 MB	142 MB	
1920 x 1200	142 MB	150 MB	168 MB	
2560 x 1600	252 MB	268 MB		

*Memory values equate to startup vRAM for a given VM.

3.3 EqualLogic Storage

EqualLogic shared storage is used for Tier2 in this solution architecture.

3.3.1 PS4100E

Model	Features	Options	Uses
EqualLogic PS4100E	12 drive bays (NL-SAS/ 7200k RPM), dual HA controllers, Snaps/Clones, Async replication, SAN HQ, 1Gb	12TB – 12 x 1TB HDs	Tier 2 array for 1000 total users or less in local Tier 1 solution model (1Gb)
		24TB – 12 x 2TB HDs	
		36TB – 12 x 3TB HDs	

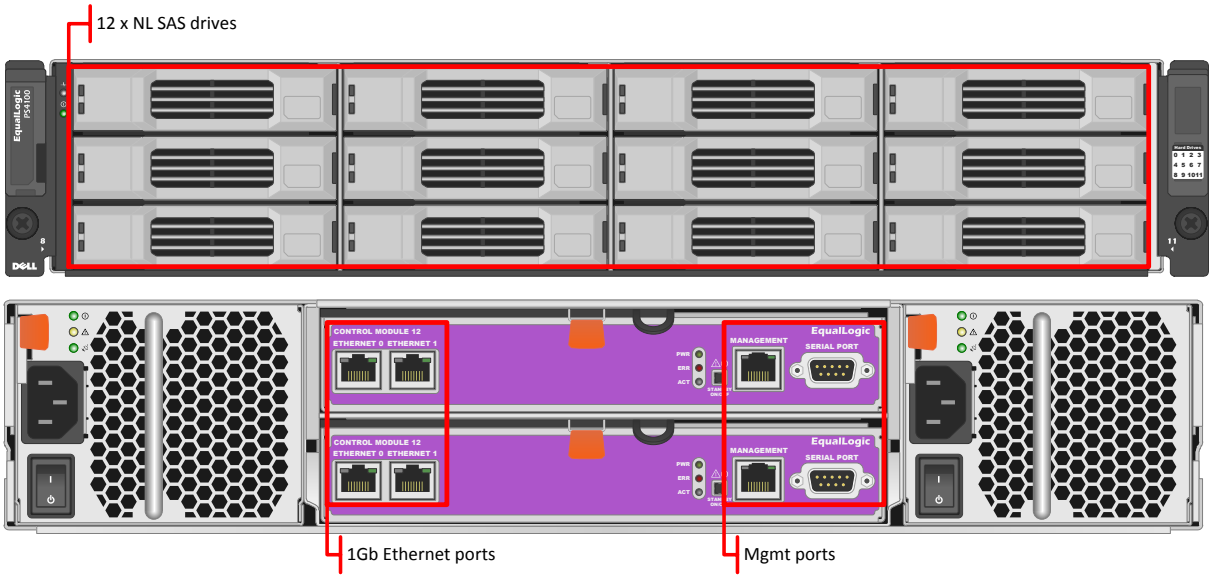

For more information on the Dell EqualLogic PS4100E and other networked storage options from Dell, please visit: LINK

3.3.2 FS7600

Model	Features	Scaling	Uses
EqualLogic FS7600	Dual active-active controllers, 24GB cache per controller (cache mirroring), SMB & NFS support, AD-integration. Up to 2 FS7600 systems in a NAS cluster (4 controllers). 1Gb iSCSI via 16 x Ethernet ports.	Each controller can support 1500 concurrent users, up to 6000 total in a 2 system NAS cluster.	Scale out NAS for Local Tier 1 to provide file share HA.

3.4 Dell Wyse End Points

3.4.1 Display Choices for Dell Wyse Endpoints

Good	**E2213*** 1680x1050 VGA, DVI	
Better	**P2212H** 1920x1080 VGA, DVI, USB	
Best	**U3011*** 2560x1600 VGA, DVI, DP, USB	

*U3011 Not Available In Mainland China

*E2213 Not Available In South Korea

3.4.2 Dell Wyse T10

The T10 handles everyday tasks with ease and also provides multimedia acceleration for task workers who need video. Users will enjoy integrated graphics processing and additional WMV & H264 video decoding capabilities from the Marvell ARMADA™ PXA 510 v7 1.0 GHz System-on-Chip (SoC). In addition, the T10 is one of the only affordable thin clients to support dual monitors with monitor rotation, enabling increased productivity by providing an extensive view of task work. Designing smooth playback of high bit-rate HD video and graphics in such a small box hasn't been at the expense of energy consumption and heat emissions either. Using just 7 watts of electricity earns this device an Energy Star V5.0 rating. In addition, the T10's small size enables discrete mounting options: under desks, to walls, and behind monitors, creating cool workspaces in every respect.

3.4.3 Dell Wyse D10D

The Dell Wyse D10D is a high-performance and secure ThinOS 8 thin client that is absolutely virus and malware immune. The D10D features an advanced dual-core AMD processor that handles demanding multimedia apps with ease and delivers brilliant graphics. Powerful, compact and extremely energy efficient, the D10D is a great VDI end point for organizations that need high-end performance but face potential budget limitations.

3.4.4 Dell Wyse D90D7

The Dell Wyse D90D7 is a high-performance Windows Embedded Standard 7 thin client for virtual desktop environments. Featuring a dual-core AMD processor and a revolutionary, unified engine that eliminates performance constraints, the D90D7 achieves outstanding speed and power for the most demanding VDI and embedded Windows applications, rich graphics and HD video. Take a unit from box to productivity in minutes. Just select the desired configuration and the D90D7 does the rest automatically—no need to reboot. And with Microsoft Windows Embedded Device Manager connectivity, you can leverage your existing Microsoft System Center Configuration Manager platform. The D90D7 is an ideal thin client for demanding virtual desktop or cloud applications.

3.4.5 Dell Wyse Z90D8

The versatile Z90D8 gives people the freedom to mix and match a broad range of legacy and cutting edge peripheral devices. Ports for parallel, serial, and USB 3.0 offer fast, flexible connectivity. Like all Dell Wyse cloud clients, the new Dell Wyse Z90D8 is one cool operator. Its energy efficient processor – which out-performs other more power hungry alternatives – and silent fan-less design, all contribute to lowering an organization's carbon footprint through power usage and emissions that are a fraction of traditional PC desktops.

4 Solution Architecture for Microsoft Remote Desktop Services

4.1 Overview

This solution architecture follows a distributed model where solution components exist in layers. The Compute layer is where VDI desktop VMs execute, the Management layer being dedicated to the broker management role VMs. Both layers, while inextricably linked, scale independently.

4.1.1 RDS Options

Server 2012 RDS provides a number of VDI options to meet your needs, all within a single, simple, wizard-driven environment that is easy to set up and manage.

- **Sessions**, hosted by the RDSH role (formerly Terminal Services), provide easy access to a densely shared session environment. Each RDP-based session shares the total available server resources with all other sessions logged in concurrently on the server. This is the most cost effective option and a great place to start with Server 2012 RDS. An RDS CAL is required for each user or device accessing this type environment.
- **Pooled VMs** are the non-persistent user desktop VMs traditionally associated with VDI. Each user VM is assigned a dedicated slice of the host server's resources to guarantee the performance of each desktop. The desktop VM is dedicated to a single user while in use then returned to the pool at logoff or reboot and reset to a pristine gold image state for the next user. Applications can be built into gold images or published via RemoteApps. An RDS CAL is required for each user or device accessing this type environment.
- **Personal VMs** are persistent 1-to-1 desktop VMs assigned to a specific entitled user. All changes made by Personal VM users will persist through logoffs and reboots making this a truly personalized computing experience. An RDS CAL is required for each user or device accessing this type environment.

Please contact Dell or Microsoft for more information on licensing requirements for VDI.

4.2 Compute Server Infrastructure

The Compute host configuration for the enterprise solution varies slightly as to whether it will be hosting RDSH or RDVH roles, or both. The RDVH role must be enabled in the Hyper-V parent partition thus providing one RDVH role per Compute host if pooled or personal VMs are required.

For scenarios requiring pooled desktops, the RDVH role is enabled in the parent and supports up to 250 pooled VDI sessions in child partitions.

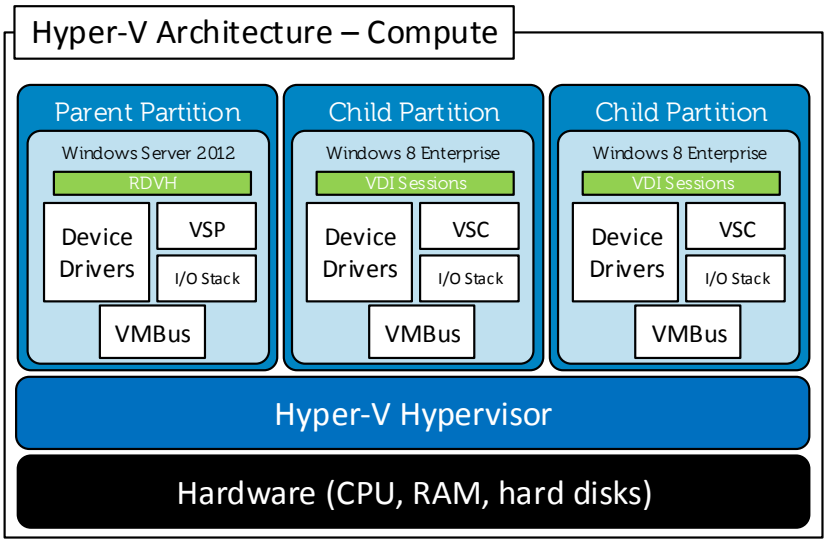

The requirements for RDVH scenarios per our workload testing scenarios are outlined below:

User Type	vCPU	Startup RAM (GB)	Dynamic Memory			NIC	OS + Data vDisk (GB)
			Min|Max	Buffer	Weight		
Basic	1	512	512MB | 2GB	20%	Med	1	40
Standard	1	512	512MB | 3GB	20%	Med	1	40
Premium	1	512	512MB | 4GB	20%	Med	1	40

The RDSH role can be enabled in up to 4 VMs on a single Compute host to support up to 300 session-based users. The requirements for RDSH VMs are outlined below.

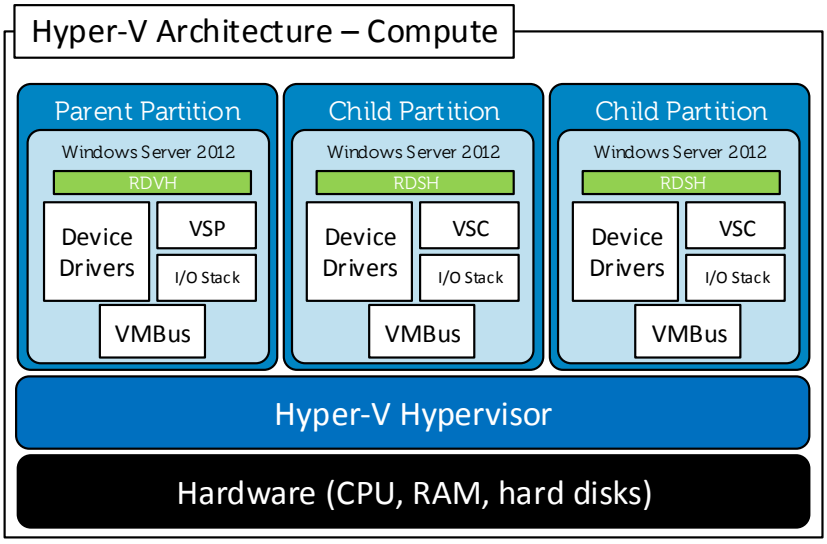

With 4 RDSH VMs on a single compute host, 4GB is reserved for the Management OS. All application and non-OS related files should be installed in the data disks as applicable:

Role	vCPU	Startup RAM (GB)	Dynamic Memory			NIC	OS + Data vDisk (GB)	Tier 2 Volume (GB)
			Min|Max	Buffer	Weight			
RD Session Host	8	16	512MB | 31GB	20%	Med	1	40 + 20	-

4.3 Management Server Infrastructure

The Management host configuration consists of VMs running in Hyper-V child partitions with the pertinent RDS roles enabled. No RDS roles need to be enabled in the root partition for Management hosts.

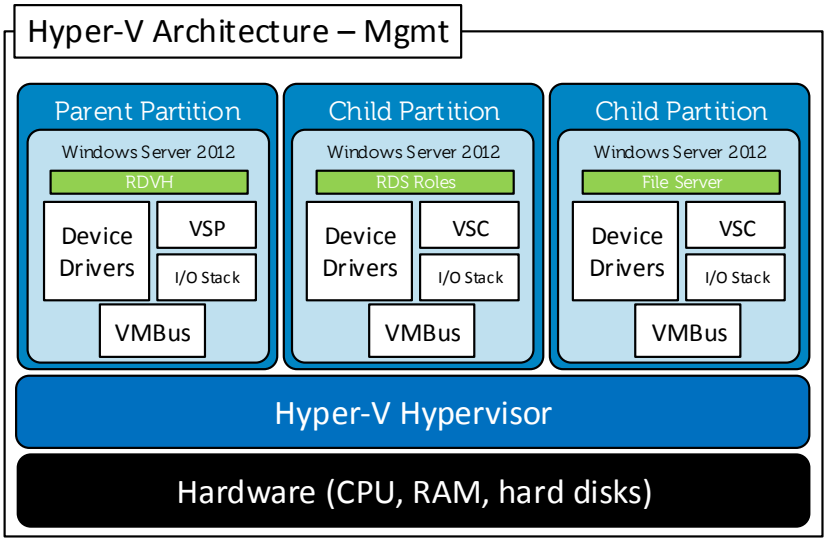

Management role requirements for the base solution are summarized below. Data disks should be used for role-specific application files, data, logs, IIS web files, and so on, and should exist in the Management volume on the 4100E array. Please note that the Tier2 volume presented to the file server is designated as a pass-through disk (PTD).

Role	vCPU	Startup RAM (GB)	Dynamic Memory			NIC	OS + Data vDisk (GB)	Tier 2 Volume (GB)
			Min|Max	Buffer	Weight			
RDCB + License Server	1	4	512MB|8GB	20%	Med	1	40 + 10	-
RDWA + RDG	1	4	512MB|8GB	20%	Med	1	40 + 10	-
File Server	1	4	512MB|8GB	20%	Med	1	40 + 10	2048 (PTD)
TOTALS	**3**	**12**				**3**	**120 + 30**	**2048**

4.4 Dell Open Manage Essentials

Dell OpenManage Essentials (OME) is a system management console that provides simple, basic Dell hardware management. Dell OME manages the host running Windows Server 2012. It supports agent-free hardware inventory and system update (e.g., BIOS, firmware) for the 11th and 12th generations of Dell PowerEdge servers. It also adds Dell SupportAssist extension for proactive "phone home" hardware issue resolution. Figure 11 shows the dashboard of Dell OME. Other features in Dell OME are:

- Discover, inventory, monitor health, and view logs of Dell clients (i.e., desktops, portables)
- Discover, inventory, and monitor health of Dell Force 10 Switches, Dell PDU and UPS devices
- Discover and collect limited inventory Linux servers via Secure Shell (SSH)
- Enhanced hardware inventory for ESXi servers with Dell OpenManage Server Administrator (OMSA)
- Alerts for hardware health status changes (not just when a hardware trap is sent)
- Ability to forward alerts in original format to support tiered event management

- A new command line interface (CLI) for automation of server discovery and grouping
- Calendar style task schedule view
- Group filtering on portal pages
- More configurable portal pages
- Grouped discovery ranges for common protocol settings
- Custom URL launch points to allow user-defined right-click actions on devices

For more information on OpenManage Essentials, see [www.delltechcenter.com/ome.](http://www.delltechcenter.com/ome)

Figure 1: OME Dashboard

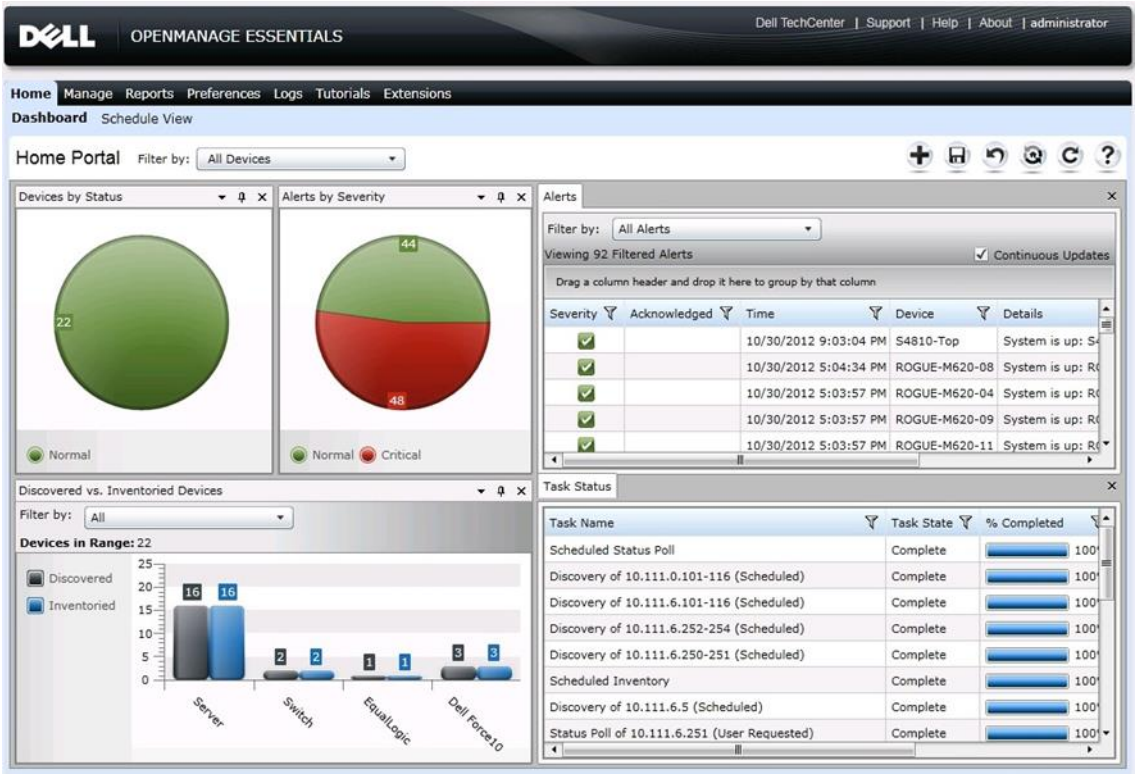

Dell Lifecycle Controller: Helps to reduce operating costs by simplifying server deployment and management. Key features include diagnostics, self-update (UEFI, Driver Pack update), firmware updates (BIOS, NIC FW, RAID Controllers), and hardware configuration. Also, the integrated iDRAC provides the mechanism for OME agentless management.

Dell OpenManage Server Administrator (OMSA): Dell OMSA provides a comprehensive, one-to-one systems management solution. OMSA provides the mechanism for the agent-based management for Dell OME. With agent-based management, OME can update the device drivers on the host in addition to other management capabilities provided by the agentless management. For more information, see http://www.delltechcenter.com/page/OpenManage+Server+Administrator+-+OMSA..

4.5 Storage Architecture

4.5.1 Local Tier 1

In the enterprise solution model, Tier 1 storage exists as local hard disks on the Compute hosts. To achieve the required performance level, RAID 10 must be used across all 12 local disks. A single volume per local Tier 1 compute host is sufficient to host the provisioned desktop VMs along with their respective write caches. Increased IO performance is provided via the 1GB cache module on the H710 RAID controller. The OS and VDI virtual disks should be specified in the PERC configuration.

Volumes	Size (GB)	RAID	Storage	Purpose	File System
OS	135	10	Tier 1	Host Operating System	NTFS
VDI	1600	10	Tier 1	Pooled + Shared VDI	NTFS

4.5.2 Shared Tier 2

Tier 2 is shared iSCSI storage used to host the Management server VMs and user data. The EqualLogic 4100 series arrays will be used for smaller scale deployments. The table below outlines the minimum volume requirements for Tier 2 storage. Larger disk sizes can be chosen to meet the capacity needs of the customer. The user data volume can be presented either via a VHDX or native NTFS pass-through disk to simplify a future upgrade to NAS. All VM disks should be presented as VHDX.

Volumes	Size (GB)	RAID	Storage Array	Purpose	File System
Management	500	50	Tier 2	RDS VMs, File Server	NTFS
User Data	2048	50	Tier 2	File Server	NTFS (PTD)
User Profiles	20	50	Tier 2	User profiles	NTFS
Templates/ ISO	200	50	Tier 2	ISO/ gold image storage (optional)	NTFS

4.5.3 Virtual Hard Disk Format

The VHDX disk format provides numerous advantages over the older VHD specification and should be used for all virtual disks in the solution. Larger vDisk support, up to 64TB, corruption protection during power failures, and larger sector disk alignment are a few of the new features.

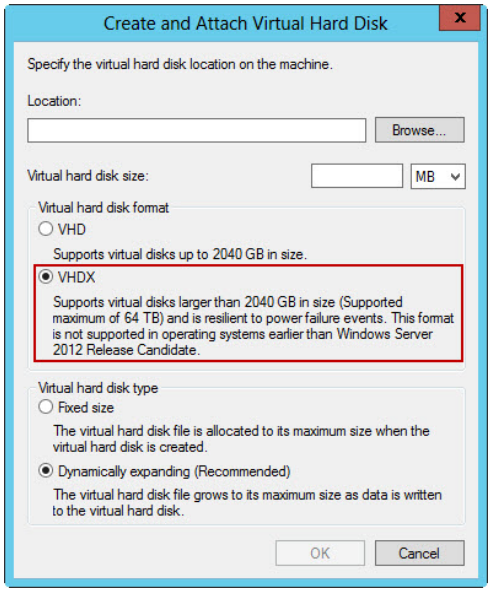

4.5.4 DNS

DNS plays a crucial role in the environment not only as the basis for Active Directory but will be used to control access to the various 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 should be adhered to.

To plan for eventual scaling, access to components that may live on one or more servers should be considered during initial deployment. The use of CNAMEs and the round robin DNS mechanism should be employed to provide a front-end "mask" to the back-end server actually hosting the service or data source.

4.5.5 File Services

The File Services role will be provided via a dedicated VM. In the interest of portability and providing a clean path to an optional HA upgrade, the volumes can be presented to the file server VM in the form of a Pass-Through Disk. This will ensure a cleaner transition for customers who upgrade to HA and add a NAS head to their environments by keeping the data on the storage array and not inside a VHDX that will need to be copied out.

4.5.6 User Profile Disks

User Profile Disks is a component of the Server 2012 RDS solution which is used to manage user profiles. Profile Disks provide a cohesive method to manage user documents and profile data in a VDI environment. Profile disks can be enabled in either session or virtual desktop collections and provide options to customize based on the need of the implementation. The solution file server will be used to host user profile disks and home drive data via SMB shares, which can be separated using an optional user data volume, if desired. Otherwise, all user data should be stored within a profile disk.

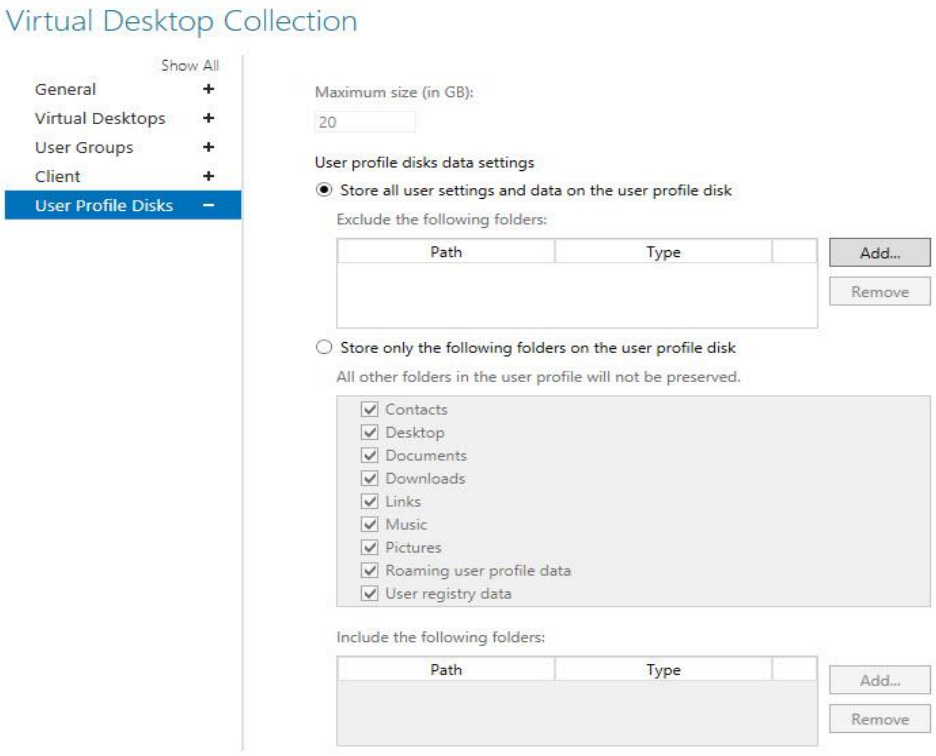

4.6 Hyper-V Configuration

The Local Tier 1 solution will be built upon the Server 2012 Hyper-V hypervisor. All Microsoft best practices and prerequisites should be adhered to (NTP, DNS, Active Directory, etc).

Solution Architecture Components	
Hypervisor	Microsoft Windows Server 2012 Hyper-V
VDI Broker	Remote Desktop Services 2012
Server OS	Microsoft Windows Server 2012 Standard edition for Management and Compute hosts.
Desktop OS	Microsoft Windows 8 Enterprise (x86)

Virtual Machine Configuration	
Hard Drive	SCSI Controller
Disk Type	VHDX – Dynamically Expanding
Smart Paging File	Store with VM
Virtual CPU	1 per VDI VM
Dynamic Memory (VDI)	512MB – Minimum, 2GB – Maximum

4.6.1 Core Components

Each Compute and Management host will run the full GUI version of Server 2012 in this solution. All RDS component roles, except for RDVH, will exist as VMs yielding 100% virtualized architecture in both Compute and Management server layers. RD Session Hosts will be enabled in dedicated VMs on the Compute hosts, while the RDS infrastructure components will be enabled in dedicated VMs on the Management hosts. The RDVH role will be enabled in the parent partition of selected Compute hosts as discussed in section 4.2.

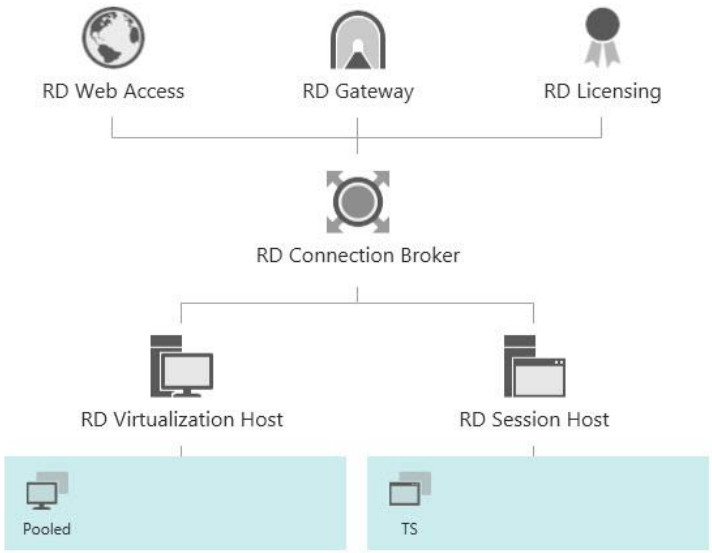

4.6.2 Hyper-V Networking (Local Tier 1)

The network configuration in this model will vary slightly between the Compute and Management hosts. The Compute hosts will not need access to iSCSI storage since they are hosting the VDI sessions on local disk. The following outlines the VLAN requirements for the Compute and Management hosts in this solution model:

- Compute hosts (Local Tier 1)
  - Management VLAN: Configured for Hyper-V infrastructure traffic – L3 routed via core switch
  - VDI VLAN: Configured for VDI session traffic – L3 routed via core switch
- Management hosts (Local Tier 1)
  - Management VLAN: Configured for Hyper-V Management traffic – L3 routed via core switch
  - iSCSI VLAN: Configured for iSCSI traffic – L2 switched only via ToR switch
  - VDI Management VLAN: Configured for VDI infrastructure traffic – L3 routed via core switch
- An optional DRAC VLAN can be configured for all hardware management traffic, which should be L3 routed via core switch

In this solution architecture, LAN and iSCSI traffic will be segmented in dedicated VLANs but combined within a single switch to minimize the initial network investment. Following best practices and in solutions that may desire larger scales, this traffic should be separated into discrete switches. Each Local Tier 1 Compute host will have a quad port NDC as well as an add-on 1Gb dual port PCIe NIC. The LAN traffic from the server to the ToR switch should be configured as a LAG to maximize bandwidth. The Compute hosts will require 2 NIC teams: 1 for LAN and the other for management of the Hyper-V parent OS. The LAN team should be connected to a Hyper-V switch, and the parent OS can utilize the Mgmt team directly.

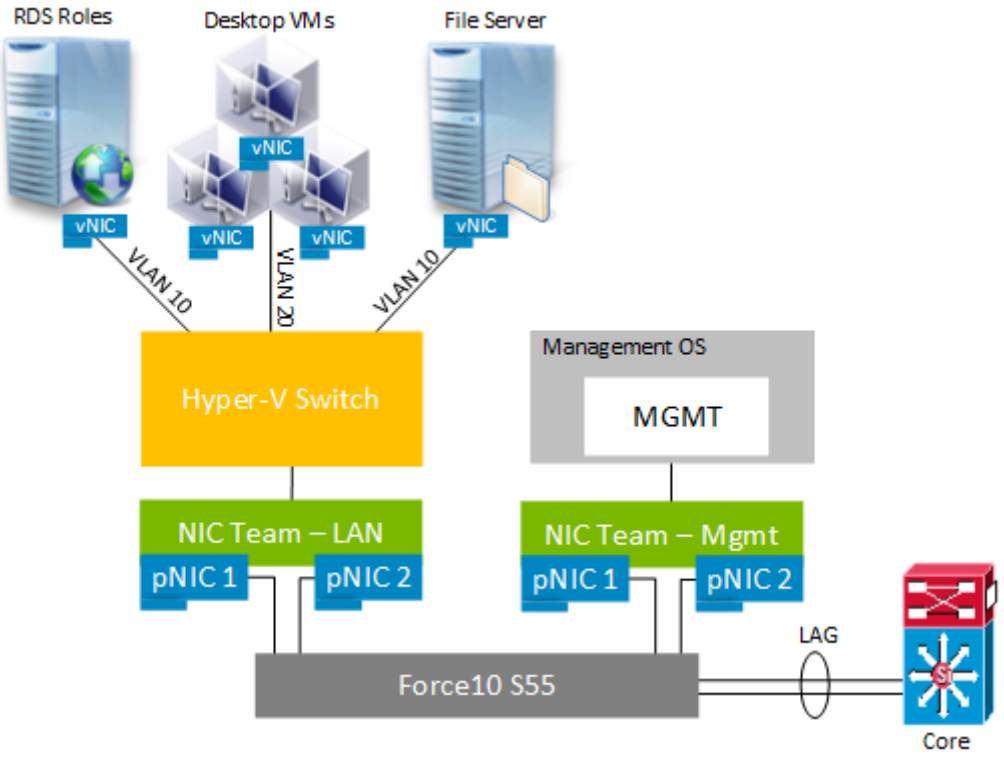

The Management hosts have a slightly different configuration since they will additionally access iSCSI storage. The add-on NIC for the Management hosts will be a 1Gb quad port NIC. 3 ports of both the NDC and add-on NIC will be used to form 2 NIC teams and 1 MPIO pair.

iSCSI should be isolated onto the NIC pair used for MPIO and connections from all 3 NIC pairs should pass through both the NDC and add-on NIC. The LAN traffic from the server to the ToR switch should be configured as a LAG. VLAN IDs should be specified in the Hyper-V switch and vNICs designated for management OS functions.

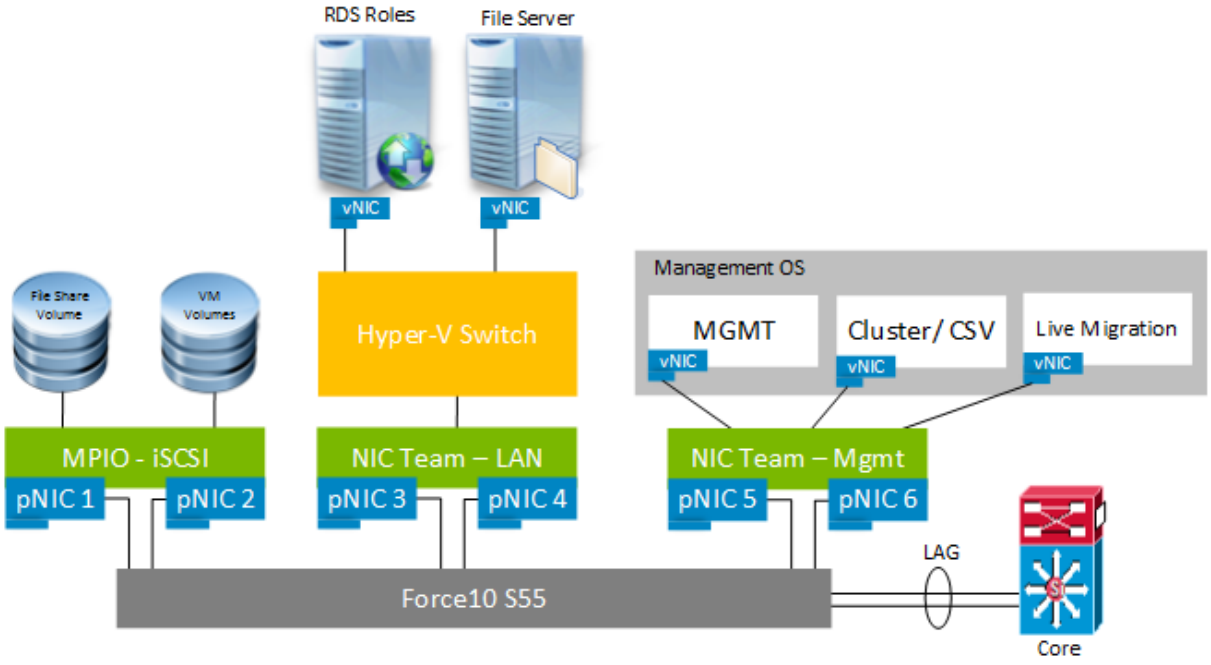

NIC teaming should be configured in the Hyper-V host using the native Windows server software and assigned to the Hyper-V switch to be used by the mgmt VMs. iSCSI should be configured to use MPIO drivers which does not need to be teamed or presented to any VMs. Team interfaces from the Mgmt team should be specified using VLANs so that the individual Management OS components can communicate individually. All NICs and switch ports should be set to auto negotiate. Additional vNICs for other Management OS functions should be attached to the Mgmt Hyper-V switch. If connecting a NIC team to a Hyper-V switch, ensure that the "Hyper-V port" load balancing mode is selected.

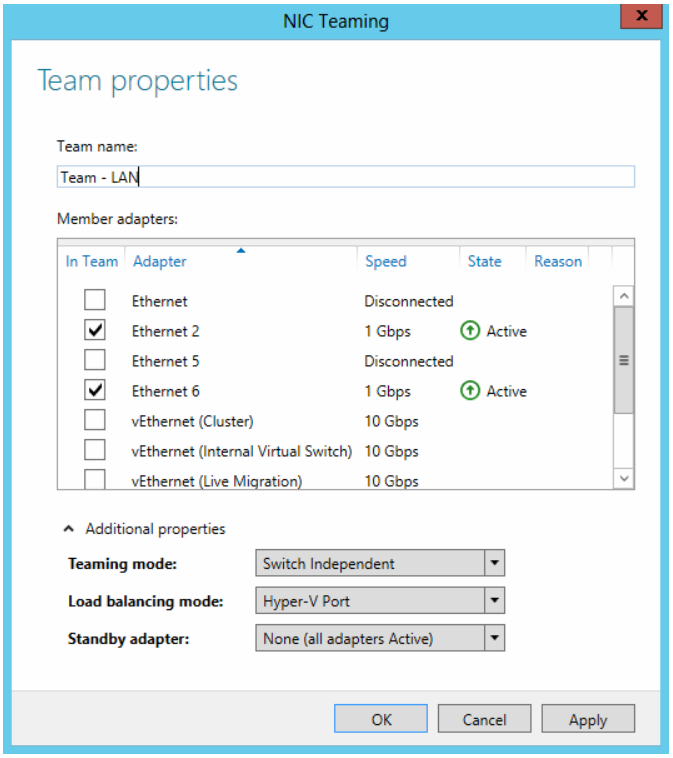

4.7 Solution High Availability

Each layer in the solution architecture can be individually protected to prevent an extended service outage. The Network layer only requires an additional switch configured in a stack with the first. Please refer to section 3.1.1.1 that covers Force10 switch stacking.

Protecting the Compute layer for RDSH and RDVH is provided by adding an additional host to a collection, thus effectively increasing the hosting capacity of a given collection. Session requests will be fulfilled by all hosts in the collection and as a result, each will have reserve capacity to insure against a host failure. **Care needs to be taken to ensure that user provisioning does not exceed the overflow capacity provided by the additional node.** A simple fail-safe measure would be to ensure that the appropriate number of users entitled to connect to the environment be tightly controlled via Active Directory. In a failure scenario users working on a failed host would simply reconnect to a fresh session on a surviving Compute host.

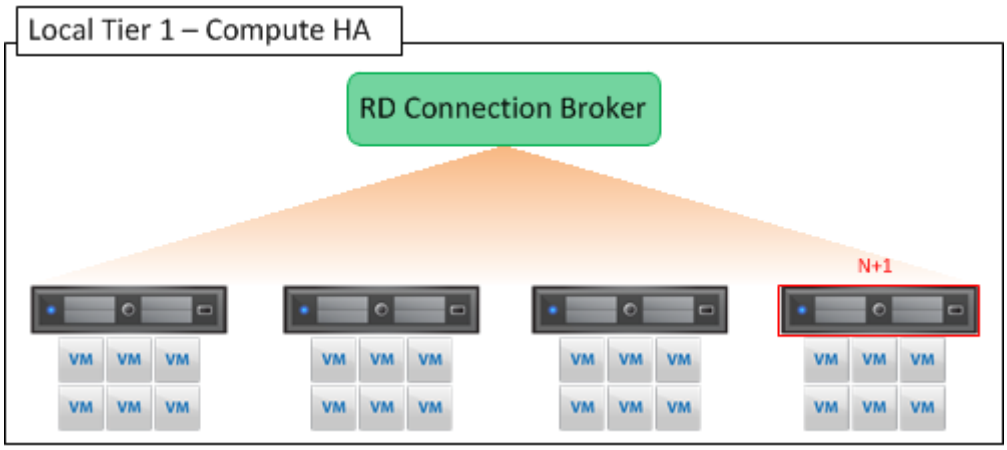

To implement HA for the Management layer, we will also add an additional host but will add a few more layers of redundancy. The following will protect each of the critical infrastructure components in the solution:

- The Management hosts will be configured in a Hyper-V cluster (Node and Disk Majority).
- The storage volume that hosts the Management VMs will be upgraded to a Cluster Shared Volume (CSV).
- SQL Server will be added to the environment to support RD Connection Broker HA.
    - Optionally SQL mirroring can be configured to further protect SQL.
- The RD Connection Broker will be configured for HA with a second RDCB introduced.

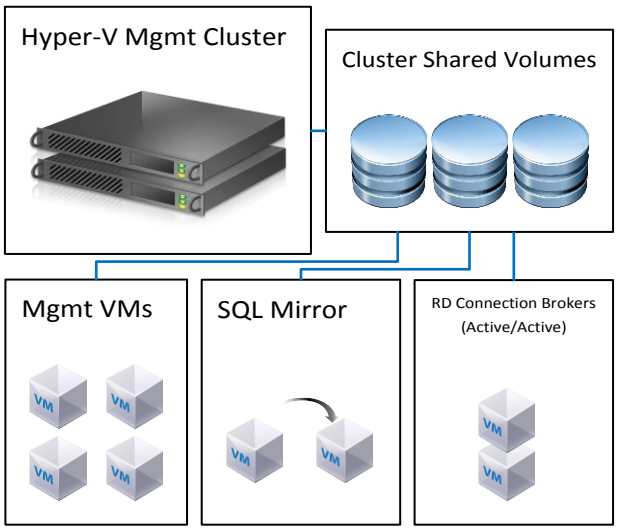

The following storage volumes are applicable in a 2-node Management layer HA scenario:

Volumes	Host	Size (GB)	RAID	Storage Array	Purpose	File System	CSV
Management	1	500	50	Tier 2	RDS VMs, File Server	NTFS	Yes
Management	2	500	50	Tier 2	RDS VMs, File Server	NTFS	Yes
SQL Data	2	100	50	Tier 2	SQL Data Disk	NTFS	Yes
SQL Logs	2	100	50	Tier 2	SQL Logs Disk	NTFS	Yes
SQL TempDB Data	2	5	50	Tier 2	SQL TempDB Data Disk	NTFS	Yes
SQL TempDB Logs	2	5	50	Tier 2	SQL TempDB Logs Disk	NTFS	Yes

Volumes	Host	Size (GB)	RAID	Storage Array	Purpose	File System	CSV
SQL Witness	1	1	50	Tier 2	SQL Witness Disk	NTFS	Yes
Quorum 1	-	500MB	50	Tier 2	Hyper-V Cluster Quorum	NTFS	Yes
User Data	-	2048	50	Tier 2	File Server	NTFS	No
User Profiles	-	20	50	Tier 2	User profiles	NTFS	No
Templates/ ISO	-	200	50	Tier 2	ISO/ gold image storage (optional)	NTFS	Yes

For more information on building redundancy for the RD Connection Broker please visit: LINK

4.7.1 SQL Database HA

HA for SQL will be provided via an optional 3-server synchronous mirror configuration that includes a witness (High safety with automatic failover). This configuration will protect all critical data stored within the database from physical server as well as virtual server problems. The principal VM that will host the primary copy of the data should exist on the first Mgmt host. The mirror and witness VMs should exist on the second or later Mgmt hosts. All critical databases should be mirrored to provide HA protection.

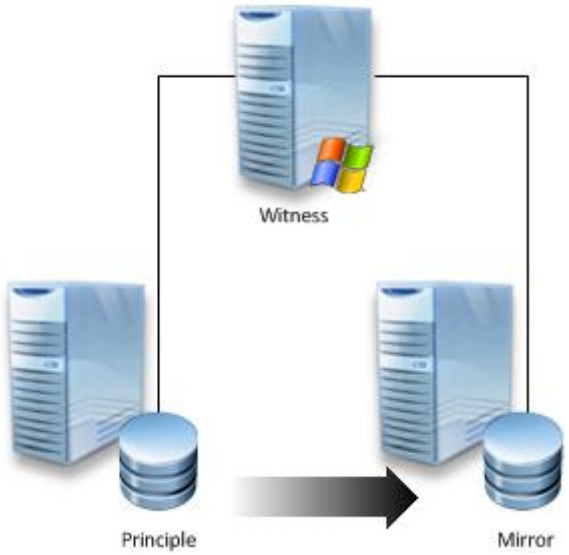

There are a number of steps required to successfully set up SQL mirroring per Microsoft best practices.

The following article details the step-by-step mirror configuration: LINK

Additional resources can be found in TechNet: LINK1 and LINK2

4.8 Application Virtualization

Microsoft Application Virtualization (App-V) provides multiple methods to deliver virtualized applications to RDS environments, virtual desktops, physical desktops, connected as well as disconnected clients. App-V can help reduce the costs and time associated with managing gold master VM and PC images with integrated applications. App-V also removes the problems of application conflicts since virtualized applications are never installed on an end point. Once an application has been packaged using the Microsoft Application Virtualization Sequencer, it can be saved to removable media, streamed to desktop clients or presented to session-based users on a

RDSH host. App-V provides a scalable framework that can be managed by System Center Configuration Manager for a complete management solution.

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

4.9 Disaster Recovery and Business Continuity

DR and BC can be achieved natively via Hyper-V Replicas. This technology can be used to replicate VMs from a primary site to a DR or BC site over the WAN asynchronously. Hyper-V Replicas are unbiased as to underlying hardware platform and can be replicated to any server, network, or storage provider. Once the initial replica is delivered from the primary site to the replica site, incremental VM write changes are replicated using log file updates. Multiple recovery points can be stored and maintained, using snapshots, to restore a VM to a specific point in time.

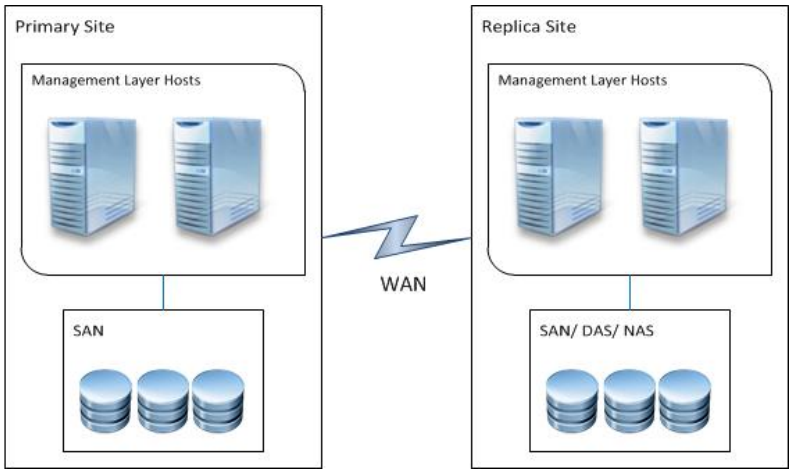

4.10 RDS 2012 Data Flow

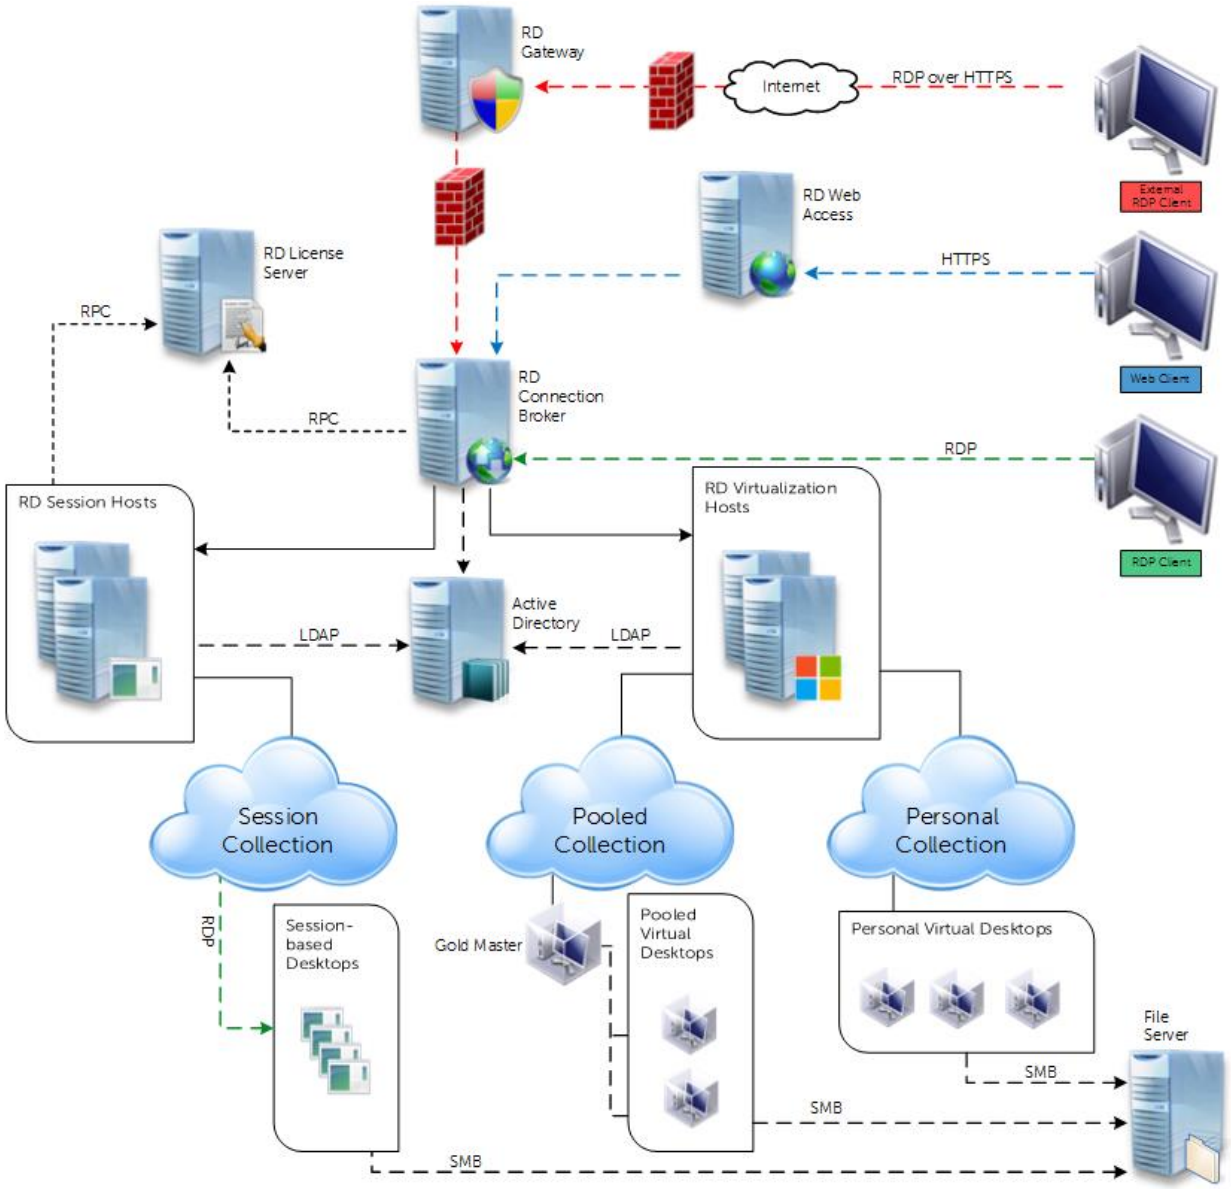

4.11 Summary

The Dell Windows Server 2012 RDS solution provides a robust and scalable VDI platform for pooled, personal and Session host deployments. Using VDI-optimized hardware in a configuration that has been validated and proven by Dell DVS Engineering, you can deploy Microsoft based VDI that is both cost effective and high performing. Our layered architecture provides flexibility to maximize your infrastructure investment with the capability to expand and contract where necessary.

Appendix A – Dell DVS 10-Seat Trial Kit

The 10 User POC bundle was purpose-built to provide high performance VDI using a modicum of infrastructure. Only 11 1Gb Ethernet ports are required (1 x server + 10 x end points) which can be provided using existing customer network infrastructure. If suitable network capacity is not in place, Dell recommends using a Force10 or PowerConnect model switch.

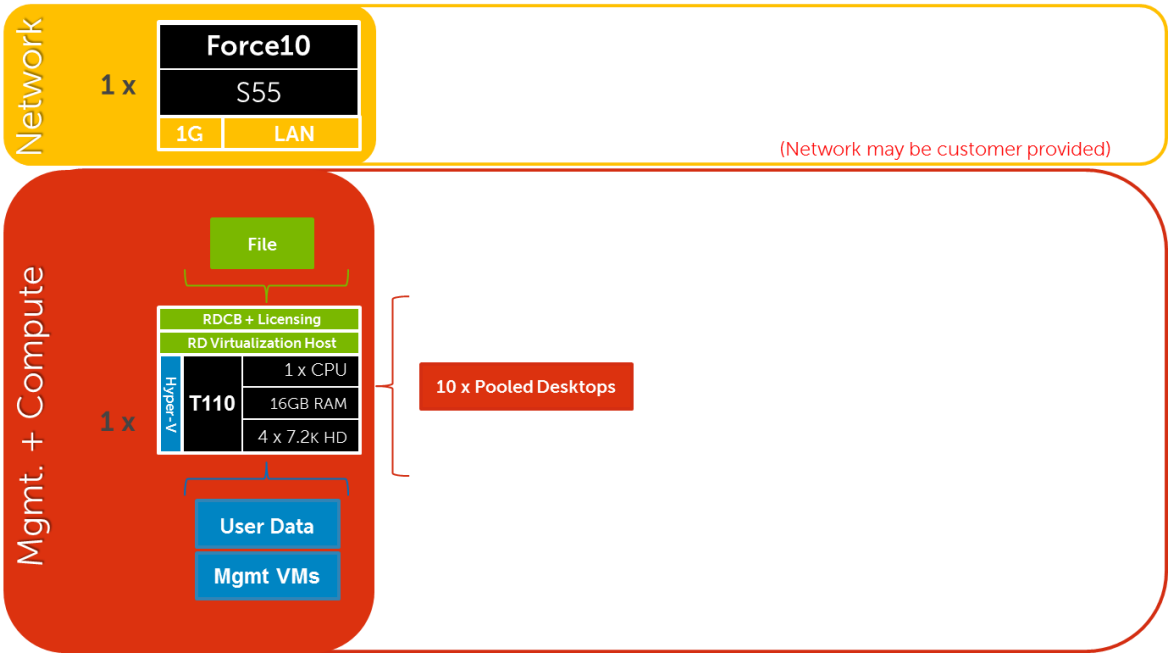

A.1 Server Configuration

The PowerEdge T110 II is the server platform of choice for this offering, providing high performance at an extremely low price of entry. Supporting the Intel Xeon E3-1200 series of CPUs and up to 32GB RAM, the T110 provides a solid server platform to get started with VDI.

All VDI server roles and desktop sessions are hosted on a single server in this model so there is no need for external storage. Higher scale and HA options are not offered with this bundle.

10 User Compute Host – PowerEdge T110 II
1 x Intel Xeon E3-1220 V2 (3.1Ghz)
16GB Memory (4 x 4GB DIMMs @ 1333Mhz) (**RDVH**)
Microsoft Windows Server 2012 Hyper-V
4 x 500GB SATA 7.2k Disks **RAID 10** (OS + VDI)
PERC H200 Integrated RAID Controller
Broadcom 5722 1Gb NIC (LAN)
305W PSU

Based on the server hardware configuration, 10 users will experience excellent performance with additional resource headroom available in reserve. The consumption numbers below are based on average performance:

Task Worker Users	CPU (%)	RAM (GB Consumed)	Disk (IOPS)	Network (Kbps)
10	40	7.5	50	262

A.2 Management and Compute Infrastructure

The solution architecture for the 10 user POC bundle combines the Compute, Management, and Storage layers onto a single server-based platform. To maximize server resources, the connection broker and license server roles are enabled within the Hyper-V parent partition, while the File server and VDI sessions exist as VMs within child partitions.

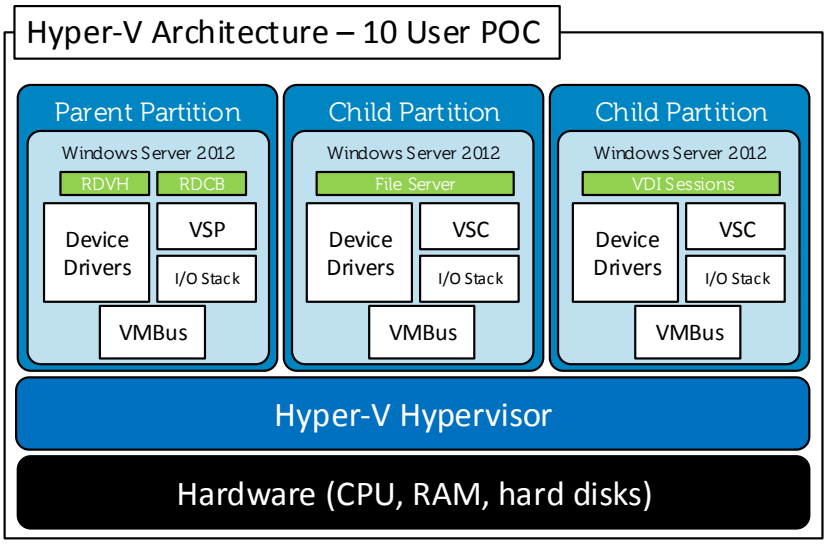

Since the RDCB and Licensing roles will be enabled within the Hyper-V parent partition, only the file server VM requires that specific physical resources be assigned.

Role	vCPU	Startup RAM (GB)	Dynamic Memory			NIC	OS + Data vDisk (GB)	Tier 2 Volume (GB)
			Min|Max	Buffer	Weight			
File Server	1	1	512MB|2GB	20%	Med	1	40 + 10	50
Pooled VDI VMs	1	512MB	512MB|2GB	20%	Med	1	20	-

A.3 Storage Configuration

The 10 User POC solution includes 4 total hard drives configured in RAID10 to host the Windows Server OS as well as VDI sessions. This configuration will maximize available performance and data protection.

Volumes	Size	RAID	Storage	Purpose	File System
OS + VDI	1TB	10	Tier 1	Host OS/ Mgmt roles + VDI Sessions	NTFS

Please see section D.4 below for performance and characterization results of the 10-Seat Trial Kit.

Appendix B – Remote Office/ Branch Office

PowerEdge VRTX by Dell is a design that reduces complexities of deployment in remote and branch offices by combining servers, storage, and networking into a single 5U chassis. The VRTX chassis supports a maximum of 4 blade servers and also includes an internally shared storage infrastructure using a SAS based RAID controller (PERC) across all the blades. An internal Ethernet switch also provides the external connectivity to the client network. This unique packaging eliminates the need for external storage arrays and fabrics for connectivity to the compute nodes.

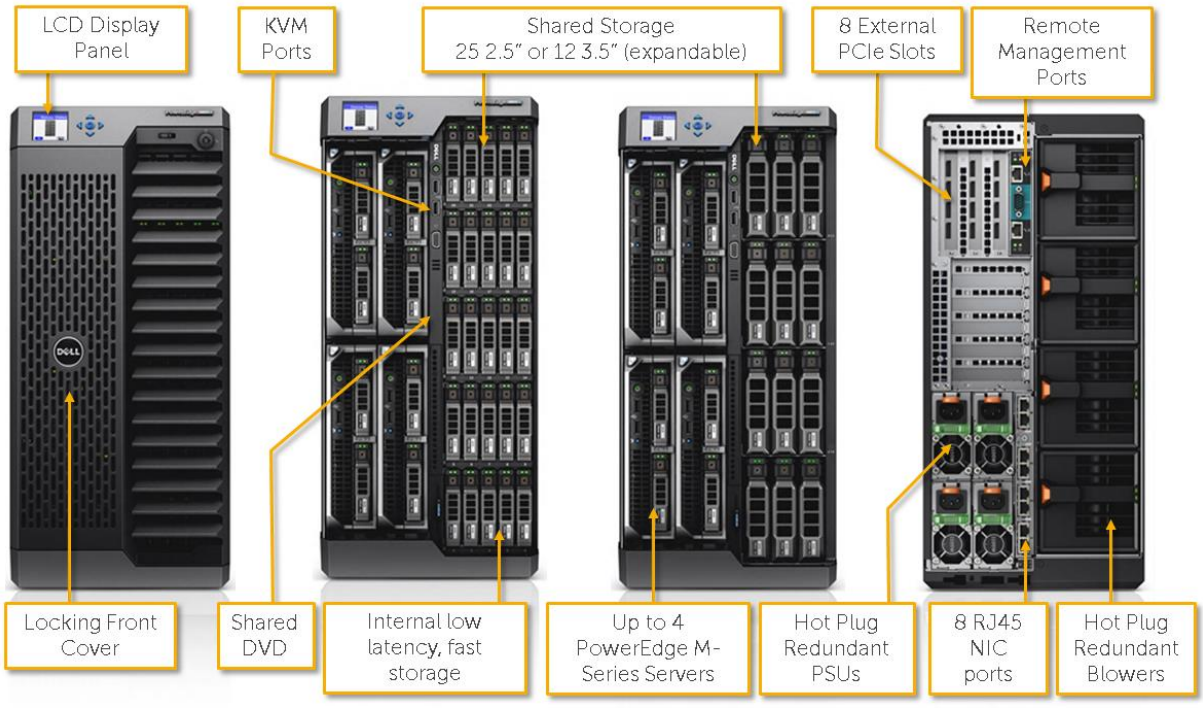

Two options for remote/branch office deployment are shown below. The solution leverages Microsoft's Failover Clustering technology to run the pertinent RDS VM roles across all blades. The

storage configuration is optimized to support the best performance and capacity utilization providing two tiers of shared internal storage. The use of Hyper-V failover clustering with Cluster Shared Volumes (CSV) enables complete mobility for the management and desktop Virtual Machines within the cluster during periods of maintenance and migration. This solution is available in a 2 blade or 4 blade configuration including the required storage. The 2 blade solution requires 15 total SAS disks, the 4 blade solution requires the full 25 disks.

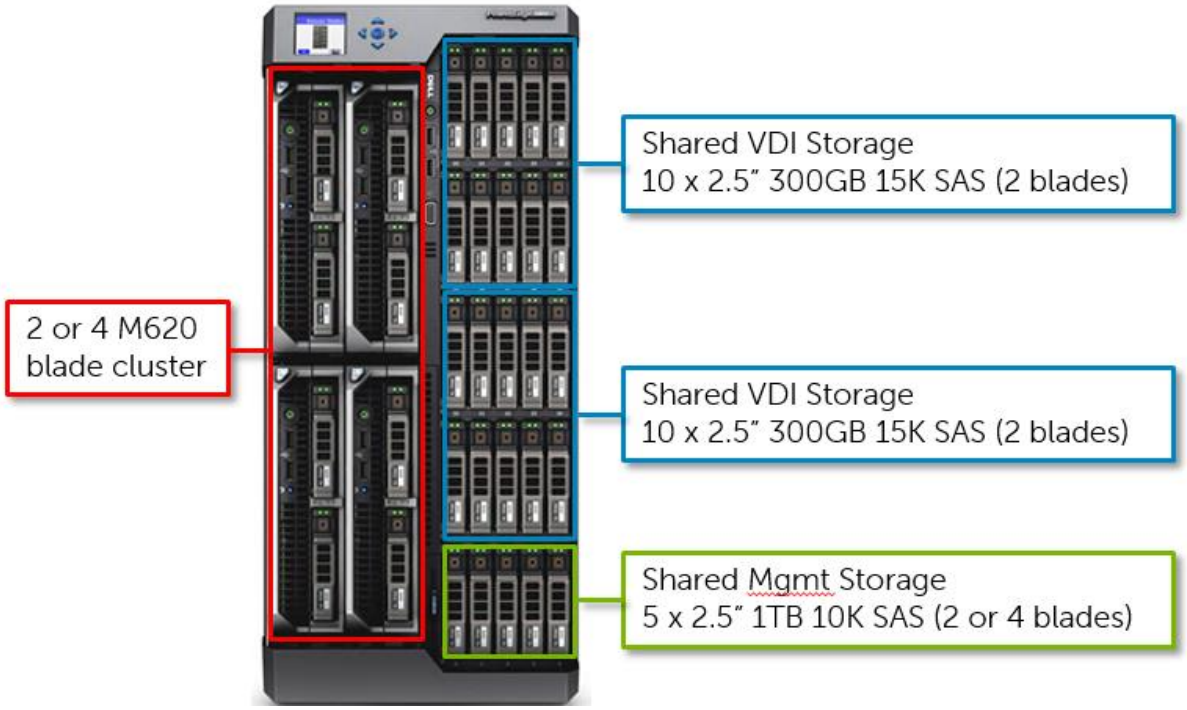

B.1 Network Configuration

The VRTX chassis can be configured with either switched (default) or pass-through modules. The switched method, shown below, supports up to 16 possible internal ports using dual fabrics with up to 8 external ports for uplinks. This solution configuration only makes use of the single A fabric in default form. External uplinks should be cabled and configured in a LAG to support the desired amount of upstream bandwidth.

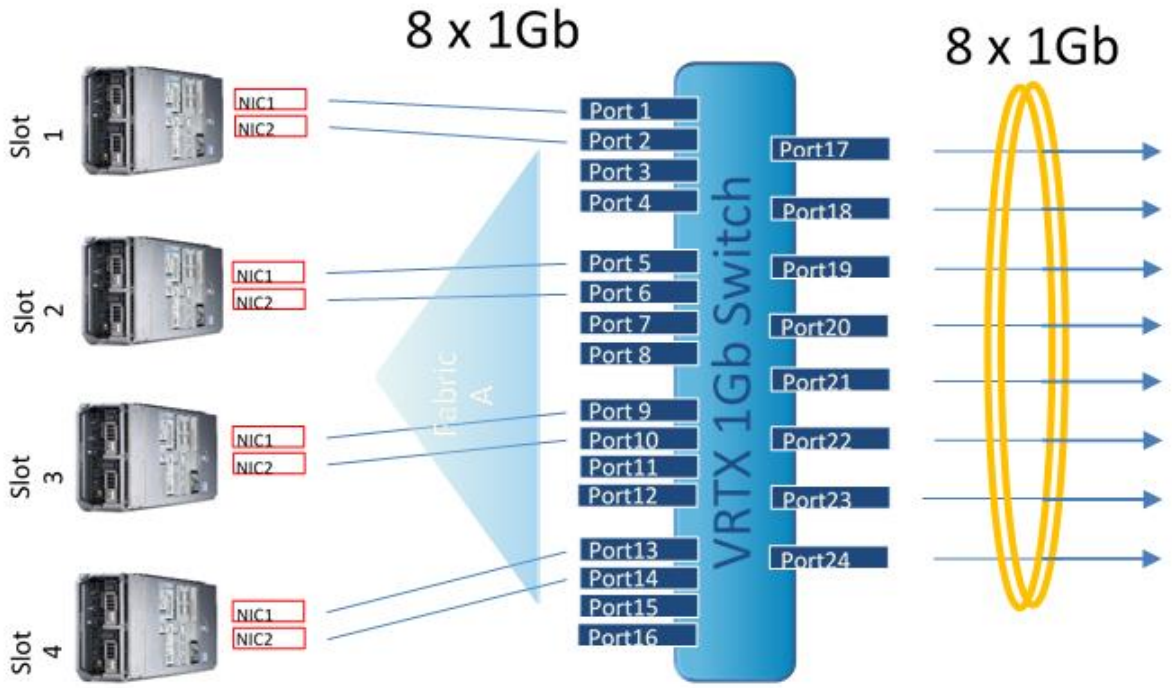

B.2 Server Configuration

The PowerEdge M620 is a feature-rich, dual-processor, half-height blade server which offers a blend of density, performance, efficiency and scalability. The M620 offers remarkable computational density, scaling up to 16 cores, 2 socket Intel Xeon processors and 24 DIMMs (768GB RAM) of DDR3 memory in an extremely compact half-height blade form factor.

VRTX Compute Host – PowerEdge M620
2 x Intel Xeon E5-2680 Processor (2.7Ghz)
256GB Memory (16 x 16GB DIMMs @ 1600Mhz)
Microsoft Windows Server 2012 Hyper-V
2 x 300GB SAS 6Gbps 15k Disks (OS)
Broadcom 57810-K DP 10Gb NDC (Fabric A)
iDRAC7 Enterprise w/ vFlash, 8GB SD

B.3 Management and Compute Configuration

The DVS Enterprise ROBO solution using VRTX consists of 2 or 4 nodes, clustered, sharing both management and VDI session hosting responsibilities across all available nodes. The RDVH role must be enabled within the Management OS and all subsequent RDS roles as well as VDI desktops will exist in child partitions running the appropriate OS (Server 2012 or Windows 8).

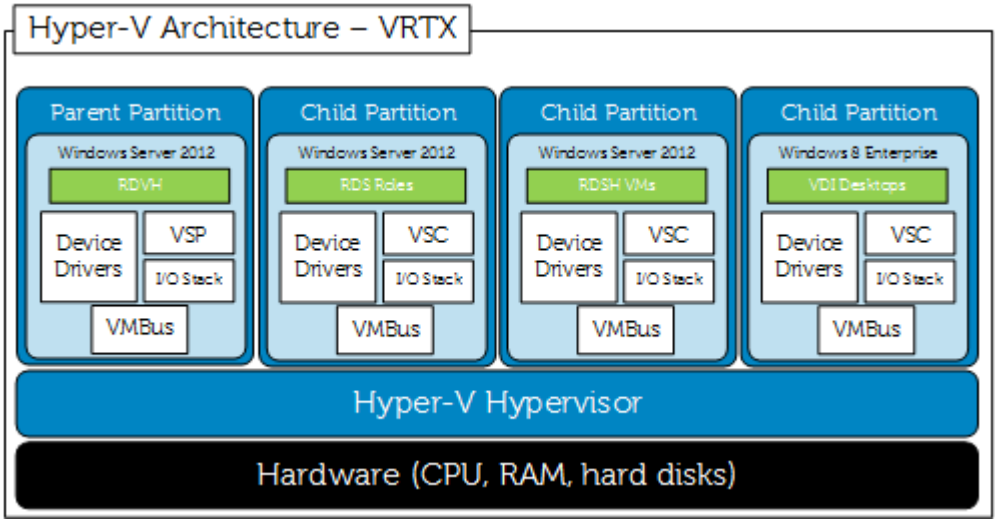

The following RDS management roles are required to support an RDVH or RDSH environment and should be spread amongst the available nodes:

Role	vCPU	Startup RAM (GB)	Dynamic Memory			NIC	OS + Data vDisk (GB)	Tier 2 Volume (GB)
			Min|Max	Buffer	Weight			
RDCB + License Server	1	4	512MB|8GB	20%	Med	1	40 + 10	-
RDWA + RDG	1	4	512MB|8GB	20%	Med	1	40 + 10	-
File Server	1	4	512MB|8GB	20%	Med	1	40 + 10	2048
TOTALS	**3**	**12**				**3**	**120 + 30**	**2048**

RD Session host configuration (up to 4 RDSH VMs per host):

Role	vCPU	Startup RAM (GB)	Dynamic Memory			NIC	OS + Data vDisk (GB)	Tier 2 Volume (GB)
			Min|Max	Buffer	Weight			
RD Session Host	8	16	512MB | 32GB	20%	Med	1	40 + 20	-

B.3.1 Virtual Networking

Native Windows Server 2012 NIC Teaming is leveraged in this solution design to maximize the available network connections and throughput. A single NIC team, by default, configured for use with a single Hyper-V switch is all that is required to serve the RDS role VMs, the VDI desktops and the various functions of the Management OS. The Management OS will share the Hyper-V Switch through the use of vNICs created via PowerShell. These individual vNICs can then be assigned specific VLANs used to segment the pertinent operations of the Management OS, such as failover cluster heartbeating and live migration. To guarantee sufficient bandwidth for a given function or vNIC, the QOS feature should be used as well.

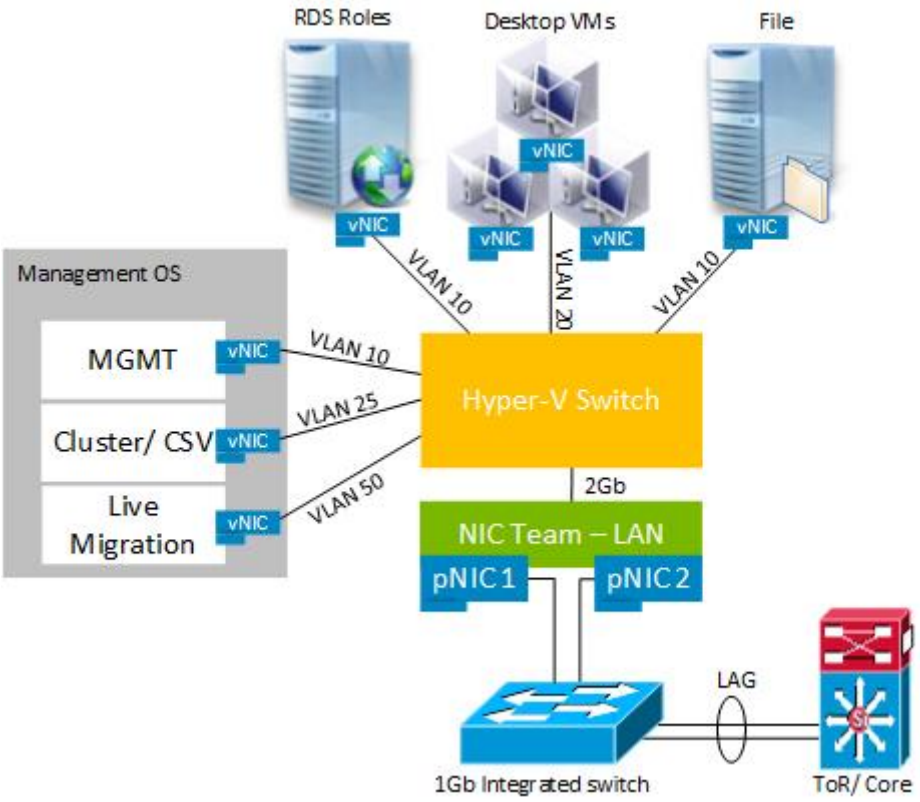

B.4 Storage Configuration

The VRTX chassis contains up to 25 available 2.5" SAS disks to be shared with each server blade in the cluster. Each blade houses its own 300GB disk pair for the Windows Server OS.

Solution Model	Features	Tier 1 Storage (VDI disks)	Tier Storage (mgmt. + user data)
2 Blade	Up to 250 desktops	10 x 300GB 2.5" 15K SAS	5 x 900GB 2.5" 10K SAS
4 Blade	Up to 500 desktops	20 x 300GB 2.5" 15K SAS	5 x 900GB 2.5" 10K SAS

RDS solution volume configuration:

Volumes	Size (GB)	RAID	Storage Array	Purpose	File System
VDI	1024	10	Tier 1	VDI Desktops	CSVFS
Management	500	10	Tier 2	RDS VMs, File Server	CSVFS
User Data	2048	10	Tier 2	File Server	NTFS
User Profiles	20	10	Tier 2	User profiles	NTFS
Quorum	500MB	10	Tier 2	Cluster	CSVFS
Templates/ ISO	200	10	Tier 2	ISO/ gold image storage (optional)	NTFS

Please see section D.5 below for performance and characterization results of the VRTX solution.

Appendix C – Testing Methodology

Performance analysis of the above architecture was carried out using Login VSI software. Login VSI is a widely used tool to generate workloads that are representative of typical corporate IT users of centralized desktop environments such as Server Based Computing (SBC) and Virtual Desktop Infrastructure (VDI). The workload produced by Login VSI for the current performance analysis effort was representative of a typical set of activities performed by a task worker (the basic workload), a knowledge worker (the standard workload), a power user (the premium workload), and a basic graphics user (the custom premium plus workload). Additionally, a custom workload using an application to simulate a high end/3D graphics test was used. Resource utilization on the compute and management nodes was monitored using Microsoft best practices for measuring performance on Hyper-V as detailed at:

http://technet.microsoft.com/en-us/library/cc768535.aspx

and

http://blogs.msdn.com/b/tvoellm/archive/2009/04/23/monitoring-hyper-v-performance.aspx

For all workloads, the performance analysis scenario used was to pre-boot all virtual desktops and then login the desired user count using a login interval of 30 seconds. Once all users have logged in, all run workload activities at steady-state for 60 minutes and then logoffs commence. Additionally, a real user logged in at the peak of resource utilization during performance analysis to ensure that the user experience was good.

It should be noted that in order to replicate a real corporate user environment, an enterprise-level anti-virus infrastructure was deployed, with McAfee VirusScan Enterprise 8.8 installed on all virtual desktops and exclusions configured locally following best practices for VDI environments.

C.1 Workload Definitions

Basic (Light) Workload / Task Worker

The basic/light workload runs a small number of applications that are representative of applications used by task workers (e.g. call center). The applications are closed immediately after use, resulting in relatively low memory and CPU consumption when compared to the standard workload. The applications used are Internet Explorer, Microsoft Word and Microsoft Outlook, with only 2 of these applications used simultaneously. The user idle time is approximately 1 minute and 45 seconds. **This workload can be used to test pooled or shared desktops.**

When testing pooled desktops with the basic workload, the following VM configuration parameters will be used:

Configuration Parameter	Setting
Number of vCPUs per VM	1
Max Memory on VM	1GB

Standard (Medium) Workload / Knowledge Worker

The standard/medium workload runs a number of applications that are representative of applications used by knowledge workers (e.g. accountants). The applications used are Internet

Explorer, a number of Microsoft Office applications (Excel, Outlook, PowerPoint and Word), Adobe Acrobat Reader, Bullzip PDF printer and 7-zip file compression software. About 2 minutes of idle time is used and a maximum of 5 applications are open simultaneously (compared to 2 for the basic workload). **This workload can be used to test pooled or shared desktops.**

When testing pooled desktops with the standard workload, the following VM configuration parameters will be used:

Configuration Parameter	Setting
Number of vCPUs per VM	1
Max Memory on VM	1.5GB

Premium (Heavy) Workload / Power User

The premium/heavy workload runs a number of applications that are representative of applications used by power users. The applications used are Internet Explorer, a number of Microsoft Office applications (Excel, Outlook, PowerPoint and Word), Adobe Acrobat Reader, Bullzip PDF printer and 7-zip file compression software. Idle time is about 40 seconds and a maximum of 8 applications are open simultaneously (compared to 5 for the standard workload). **This workload can be used to test pooled or shared desktops.**

When testing pooled desktops with the premium workload, the following VM configuration parameters will be used:

Configuration Parameter	Setting
Number of vCPUs per VM	1
Max Memory on VM	2.5GB

Premium Plus-Graphics (Custom Premium/Multimedia) Workload

The Premium Plus workload runs a number of applications that are representative of applications used by basic graphics users. The applications used are Internet Explorer, a number of Microsoft Office applications (Excel, Outlook, PowerPoint and Word), Adobe Acrobat Reader, Bullzip PDF printer and 7-zip file compression software. In addition, it runs a 30 sec 720P WMV file which is played in the loop and a HTML5 fishbowl workload runs for 60 sec in the beginning of the loop. Relative to the standard workload discussed above, idle time is slightly lower as a percentage of overall runtime and a maximum of 8 applications are open along with fishbowl activity in the beginning of the loop. **To successfully execute this workload, the virtual desktops must be GPU assisted and therefore can only be used to test pooled desktops.**

When testing pooled desktops with the premium plus workload, the following VM configuration parameters will be used:

Configuration Parameter	Setting
Number of vCPUs per VM	2
Max Memory on VM	2GB
Video RAM	80-125MB

3D Graphics – DX Heavy Workload

This custom workload is not based on Login VSI workloads. It consists of an application executing at full screen (1280 x 1024 resolution) while rendering a 3D object. It is designed to be representative of applications such as CAD and video games used by high end graphics users. **Like the Premium Plus workload, this workload also requires the use of GPUs and can only be used with pooled desktops.**

When testing pooled desktops with the premium plus workload, the following VM configuration parameters will be used:

Configuration Parameter	Setting
Number of vCPUs per VM	2
Max Memory on VM	2GB
Video RAM	Up to 384MB

C.2 Workload Summary Table

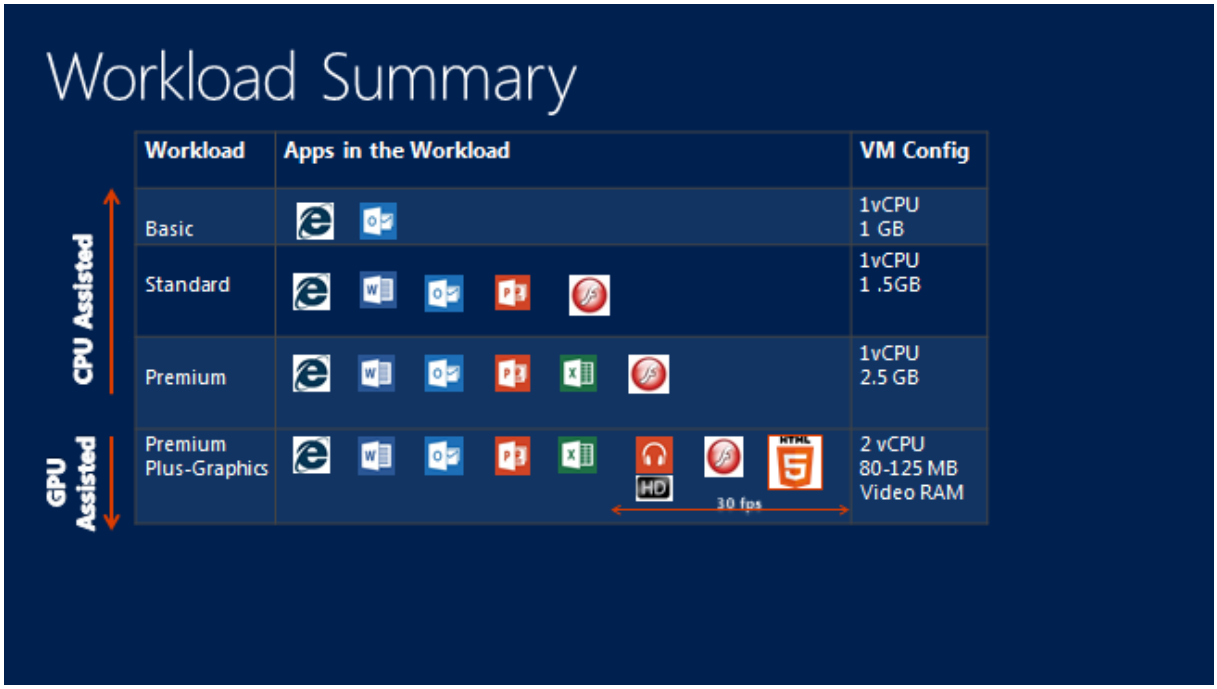

Appendix D – Pooled (RDVH) VDI Performance Analysis Results

D.1 Single Server – Compute and Management

Configuration Summary

The following validation was performed using hardware as described in the **Local Tier 1 – 250 User Pilot (distributed)** model. The compute and management roles were installed on a single R720 host with dual 2.9 GHz processors, 256GB RAM, and 12x300GB 15k local hard drives configured in a single RAID 10 array.

Microsoft Windows 8 was used as the desktop operating system.

Basic Workload Test

The CPU usage for this test reached 73.53% thus confirming that the server with this configuration can support up to 250 basic users and this will also be the baseline for testing with Windows 8 and RDSH.

CPU Resource Utilization Performance Analysis Results

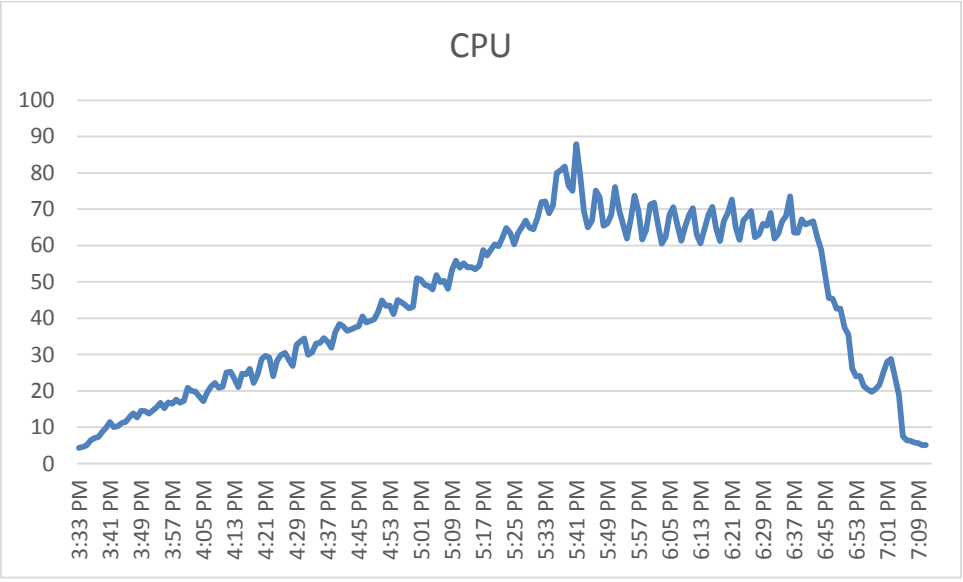

As seen from the above graph maximum CPU utilization was approximately 76.12%. After all users started logging in, CPU usage spiked from 0 to 88%. CPU became stable between 60-77% once all users logged in and dropped as the users logged off.

Memory Resource Utilization Performance Analysis Results

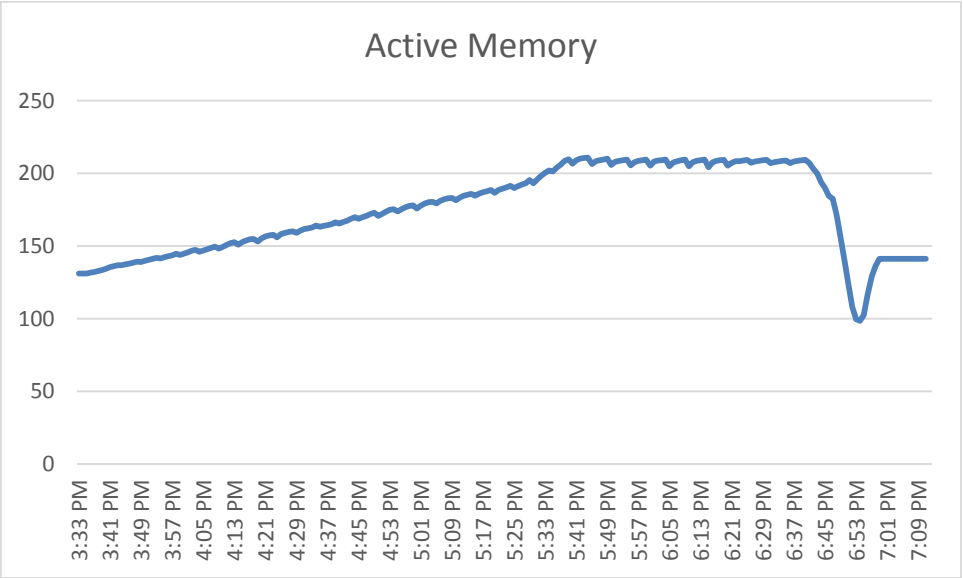

A spike is evident after all VMs start logging in however active memory is about 86.7% of available memory during the peak of validation activity.

Disk IO Resource Utilization Performance Analysis Results

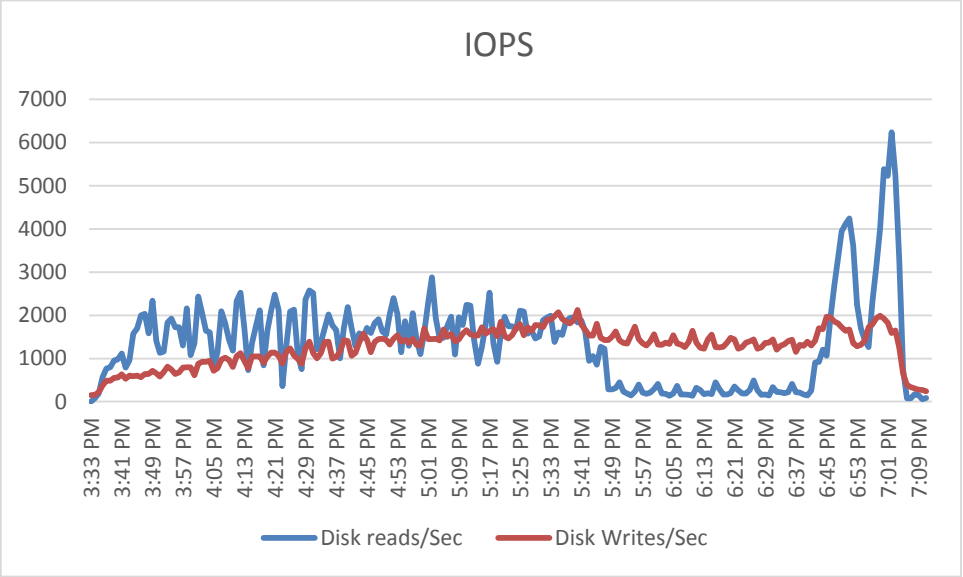

A spike is evident after all VMs start logging in. However, the read/write activity fluctuated until all users finished logging in and then became stable from 05:50PM to 06:40PM. Total Peak IOPS measured during the login state was 4326.92 giving an IOPS value of 17.30 per user and total peak IOPS in steady state (after all users logged in) was 1932.97 giving an IOPS value of 7.73 per user.

Network Resource Utilization Performance Analysis Results

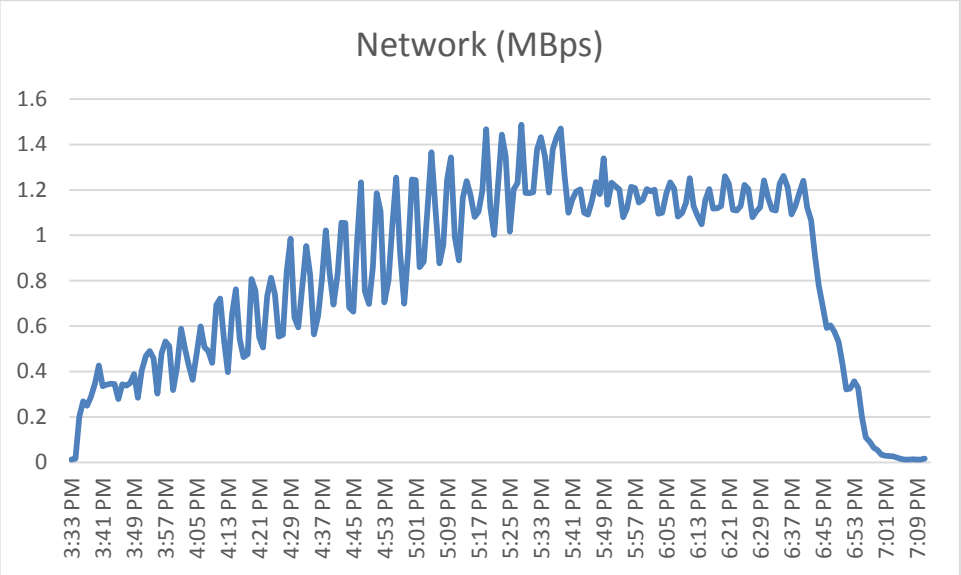

As seen from the above graph, overall network performance was good.

Results Summary Table

Workload	Server Density	CPU	Login StateTier-1 IOPS	Login State IOPS per User Tier-1	Steady State Tier-1 IOPS	Steady State IOPS per User Tier-1
Basic	250	76.12%	4326.92	17.30	1932.97	7.73

D.2 Distributed Server – Dedicated Compute

Configuration Summary

The following validations were performed using hardware as described in the **Local Tier 1 – Base** model. The compute role was dedicated to a single R720 host with dual 2.9 GHz processors, 256GB RAM, and 12x300GB 15k local hard drives configured in a single RAID 10 array. The management roles were installed on VMs running on a separate R720 host.

Microsoft Windows 8 was used as the desktop operating system.

Standard Workload Test

The CPU usage for this test reached 85% thus confirming that the server with this configuration can support up to 155 standard users.

CPU Resource Utilization Performance Analysis Results

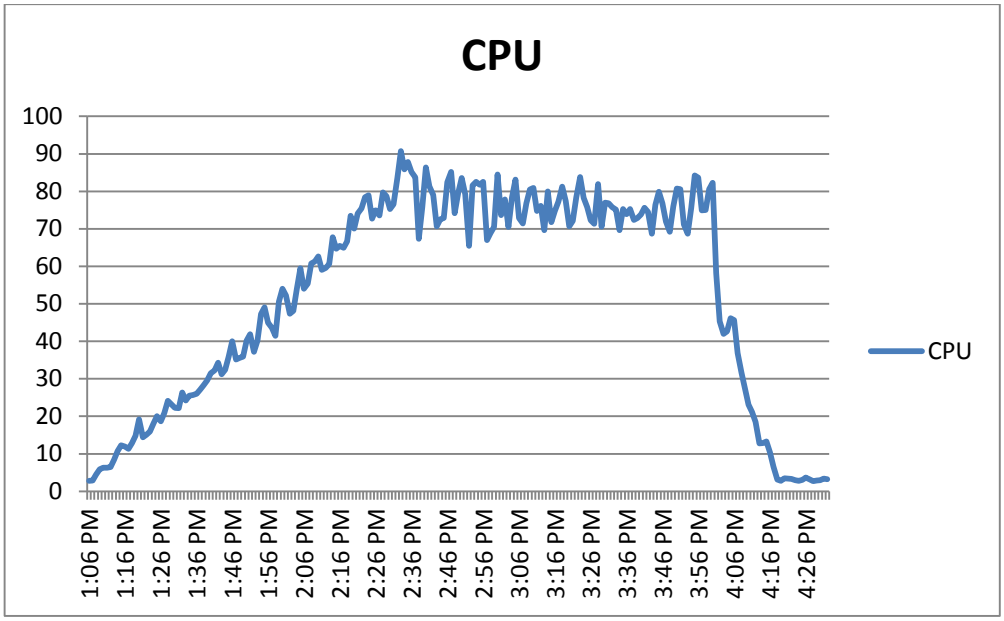

As seen from the above graph, maximum CPU utilization was approximately 85.18%. After all users started logging in, CPU usage was spiked from 10 to 91 and became stable between 65-85.18% once all users logged in and dropped as the users logged off.

Memory Resource Utilization Performance Analysis Results

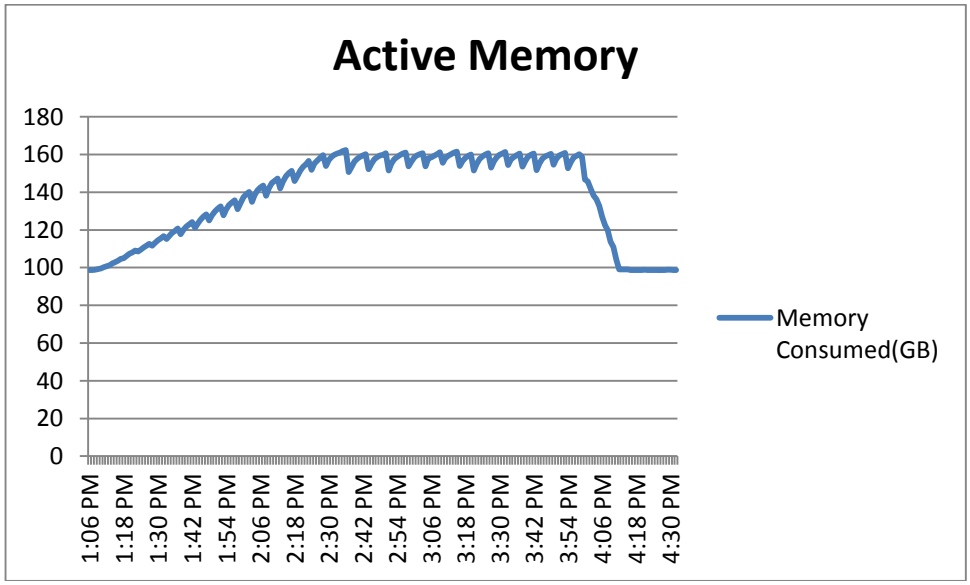

A spike is evident after all VMs start logging in; however, active memory is about 63% of available memory during the peak of validation activity.

Disk IO Resource Utilization Performance Analysis Results

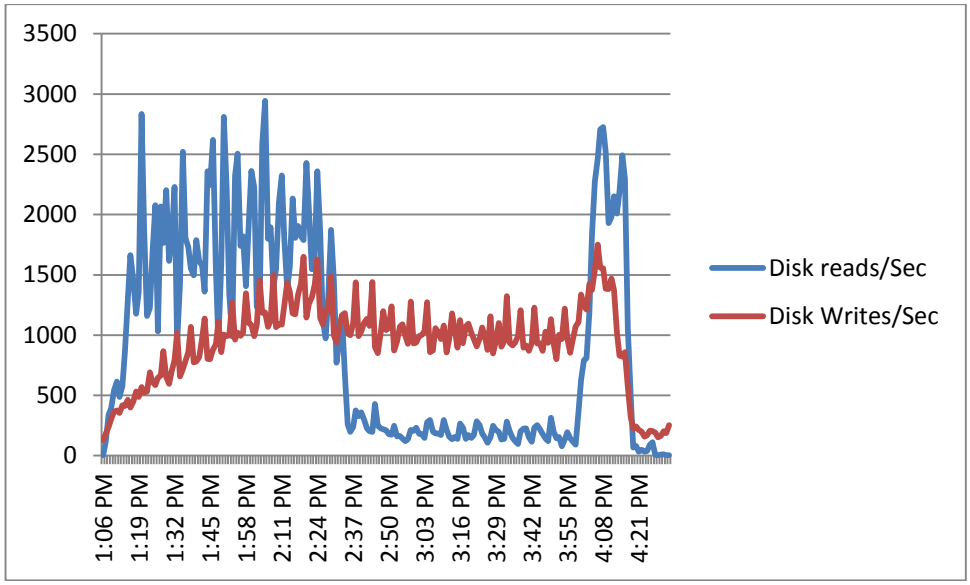

A spike is evident after all VMs start logging in; however, the read/write activity was fluctuating from 01:06PM to 02:35 PM and became stable after all users finished logging in from 02:38 PM to 03:55 PM. Total Peak IOPS measured during the login state was 4131 giving an IOPS value of 26.65 per user and total peak IOPS measured during the steady state (after all users logged in) was 1605 giving an IOPS value of 10.35 per user.

Network Resource Utilization Performance Analysis Results

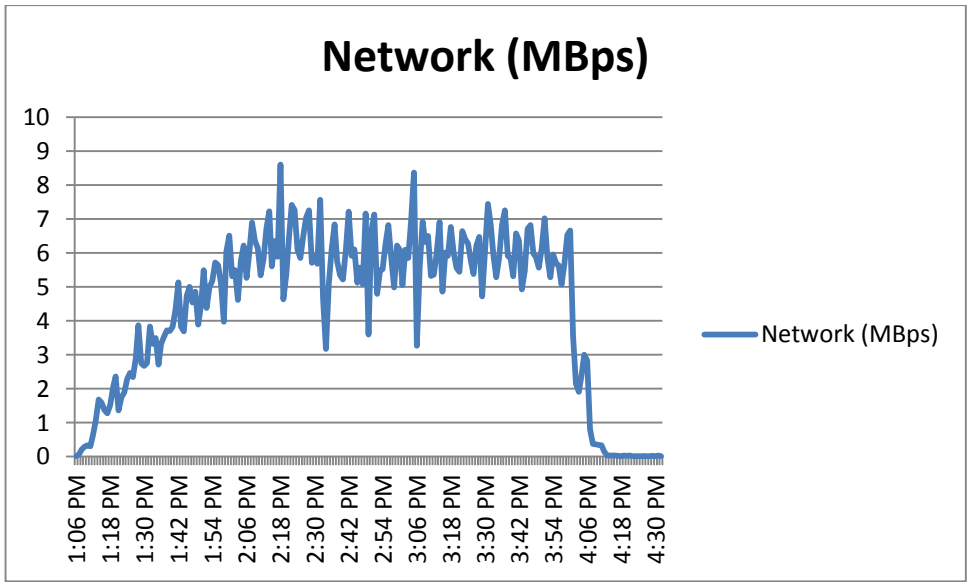

As seen from the above graph, overall network performance was good.

Premium Workload Test

The CPU usage for this test was 63%, thus confirming that the server with this configuration can support up to 135 Premium users.

CPU Resource Utilization Performance Analysis Results

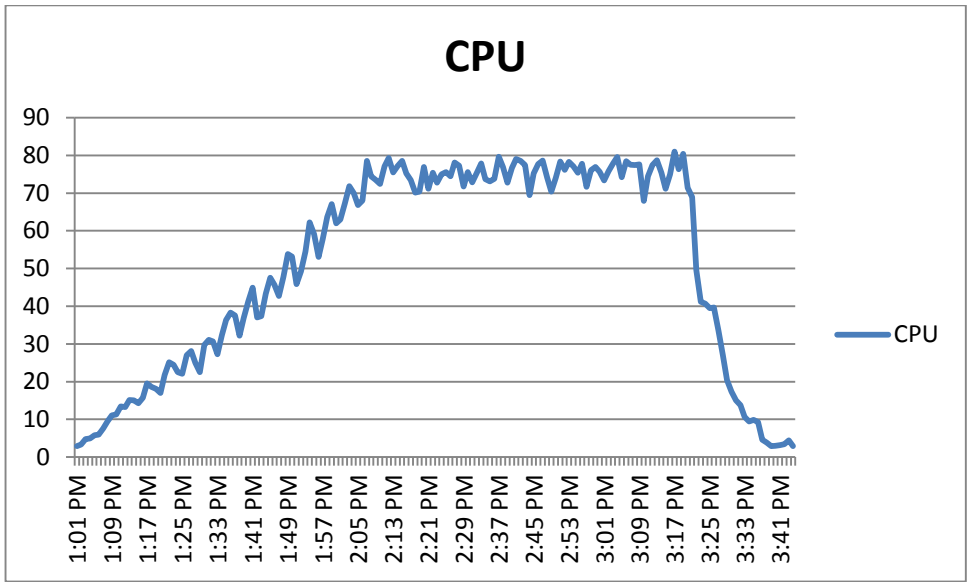

As seen from the above graph, maximum CPU utilization was approximately 80%. After all users started logging in, CPU usage spiked from 10 to 80 and became stable between 70-80% once all users logged in and dropped as the users logged off.

Memory Resource Utilization Performance Analysis Results

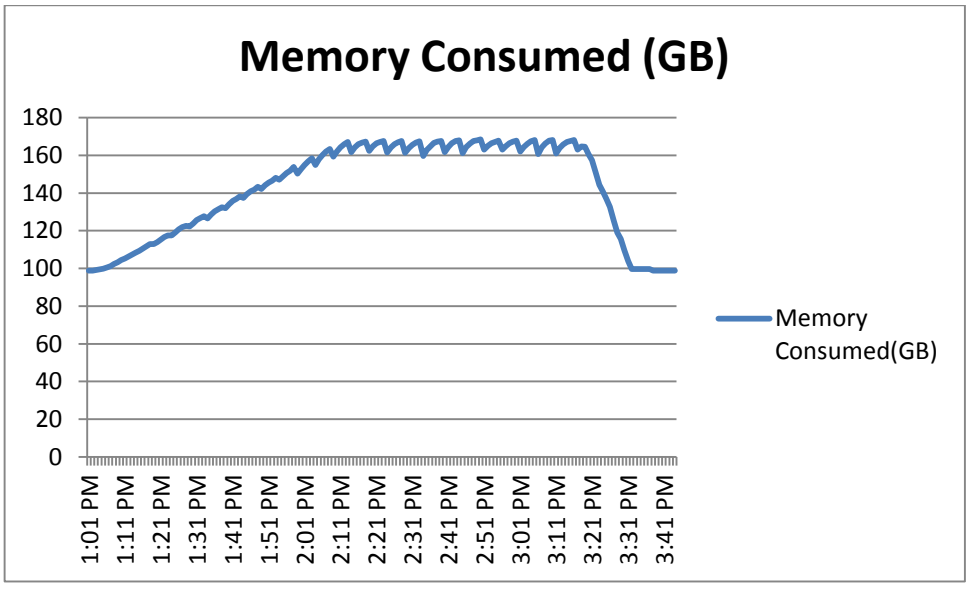

A spike is evident after all VMs start logging in; however, active memory is about 66.01% of available memory during the peak of validation activity.

Disk IO Resource Utilization Performance Analysis Results

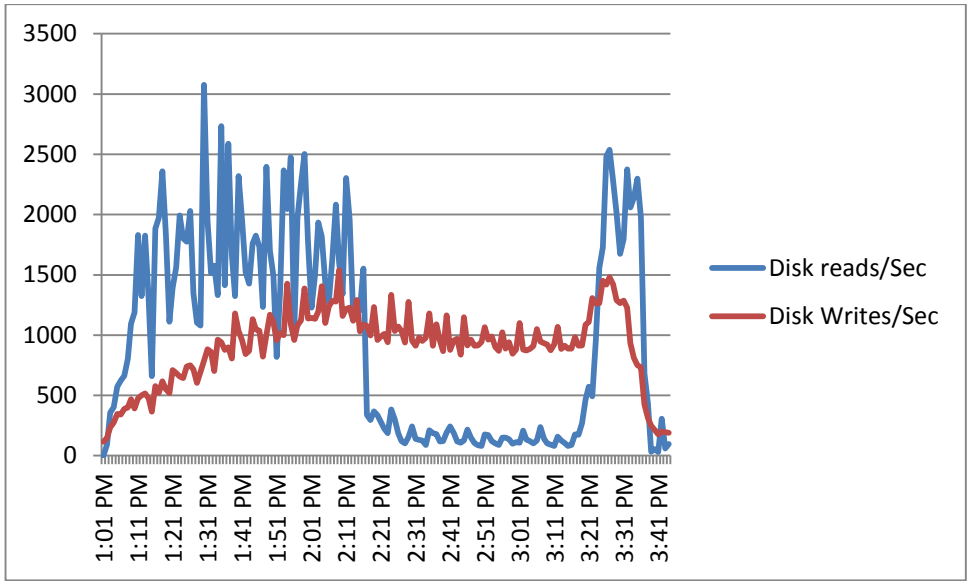

A spike is evident after all VMs start logging in; however, the read/write activity was fluctuating 01:01 PM to 2:13 PM and became stable after all users finished logging in from 02:17 PM to 3:16 PM. Total Peak IOPS measured during the login state was 3890 giving an IOPS value of 28.81 per user and total peak IOPS measured during steady state (after all users logged in) was 1719 giving an IOPS value of 12.73 per user.

Network Resource Utilization Performance Analysis Results

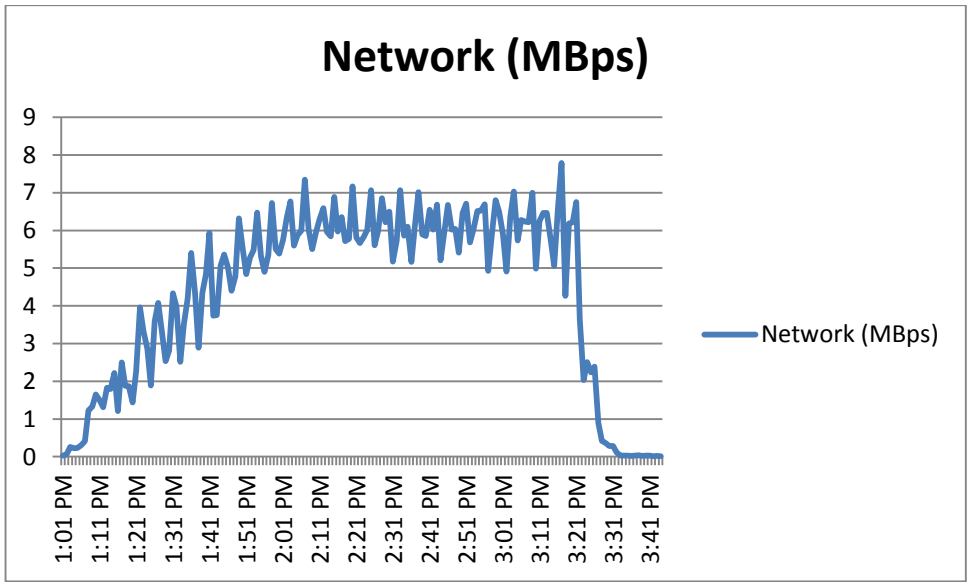

As seen from the above graph, overall network performance was good.

Results Summary Table

Workload	Server Density	CPU	Login StateTier-1 IOPS	Login State IOPS per User Tier-1	Steady State Tier-1 IOPS	Steady State IOPS per User Tier-1
Standard	155	85.18%	4131	26.65	1605	10.35
Premium	135	80%	3890	28.81	1719	12.73

D.3 Distributed Server – Dedicated GPU Compute

Configuration Summary

The following validations were performed using hardware as described in the **Local Tier 1 – Base** model; however, GPU cards were added to the compute host. The compute role was dedicated to a single R720 host with dual 2.9 GHz processors, 256GB RAM, and 12x300GB 15k local hard drives configured in a single RAID 10 array. The management roles were installed on VMs running on a separate R720 host.

Microsoft Windows 8 was used as the desktop operating system.

Premium Plus Workload Test

Three AMD S7000 GPU cards were used for this test. Driver version was 12.102.6.0. The CPU usage for this test reached 87.56% thus confirming that the server with this configuration can support up to 75 Premium Plus users.

CPU Resource Utilization Performance Analysis Results

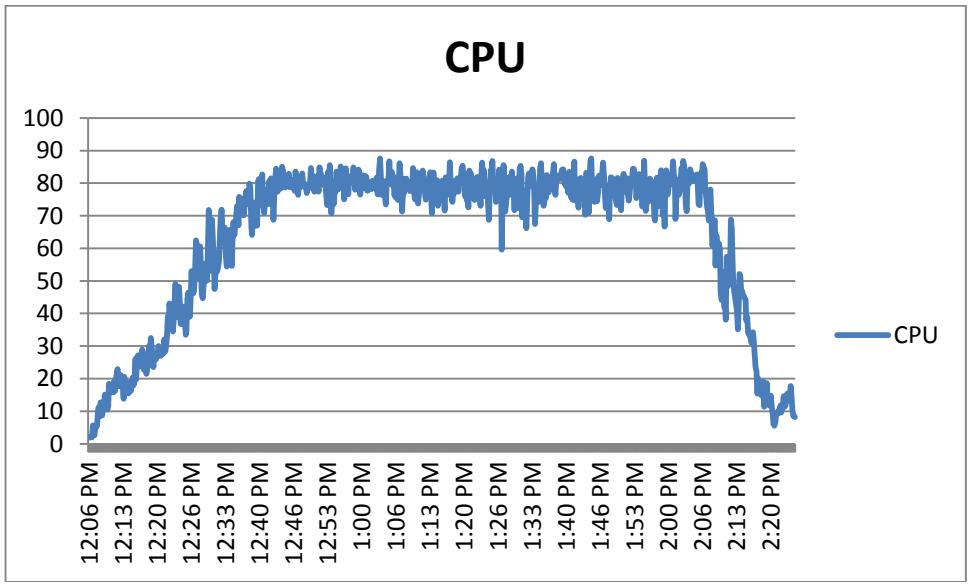

As seen from the above graph, maximum CPU utilization was approximately 87%. After all users started logging in, CPU usage spiked from 0 to 85% and became stable between 59-87% once all users logged in and dropped as the users logged off.

Memory Resource Utilization Performance Analysis Results

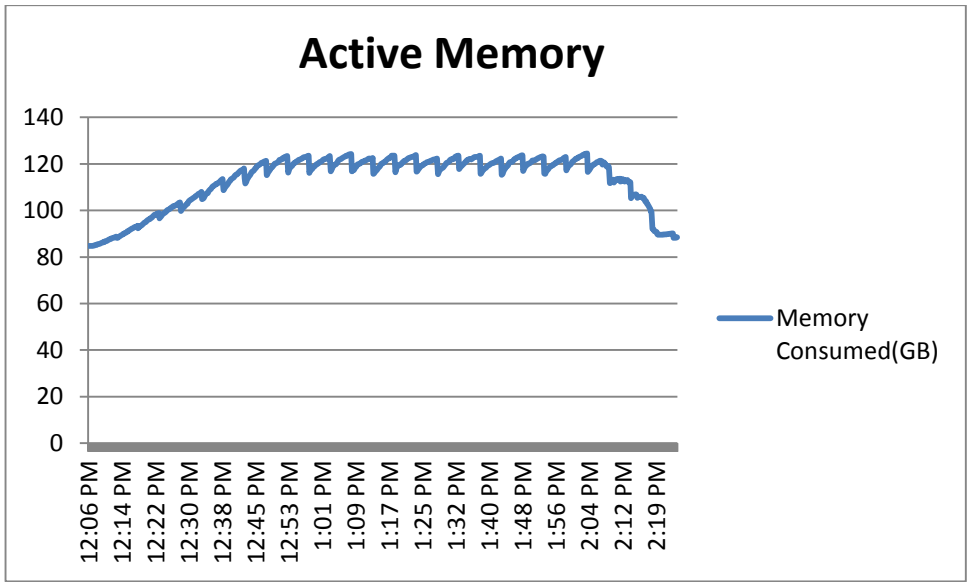

A spike is evident after all VMs start logging in; however, active memory is about 49.21% of available memory during the peak of validation activity.

Disk IO Resource Utilization Performance Analysis Results

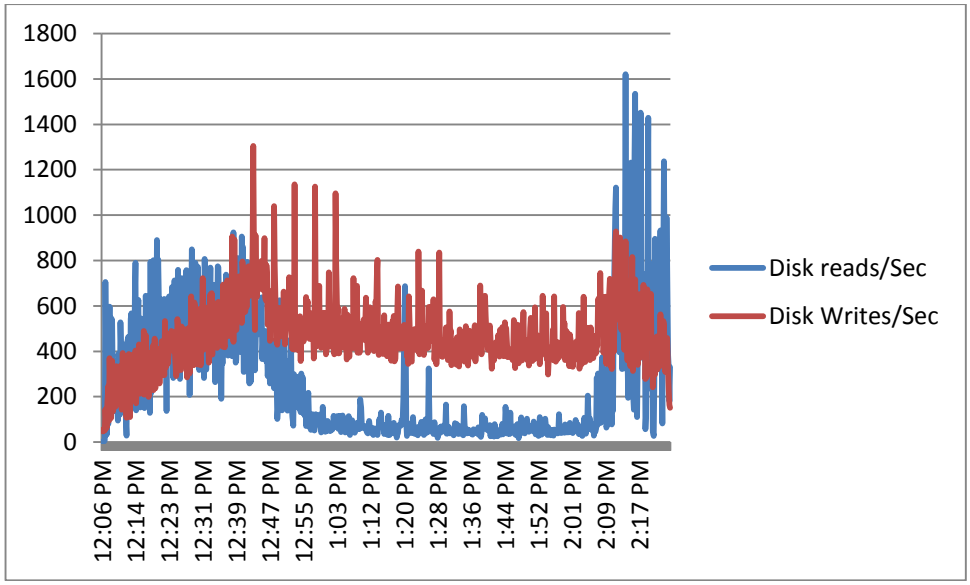

A spike is evident after all VMs start logging in; however, the read/write activity was fluctuating from 12:06 PM to 12:45PM and became stable after all users finished logging in from 12:50 PM to 02:05 PM. Total Peak IOPS measured during the login state was 1594 giving an IOPS value of 21.25 per user and total peak IOPS in steady state (after all users logged in) was 1136 giving an IOPS value of 15.14 per user.

Network Resource Utilization Performance Analysis Results

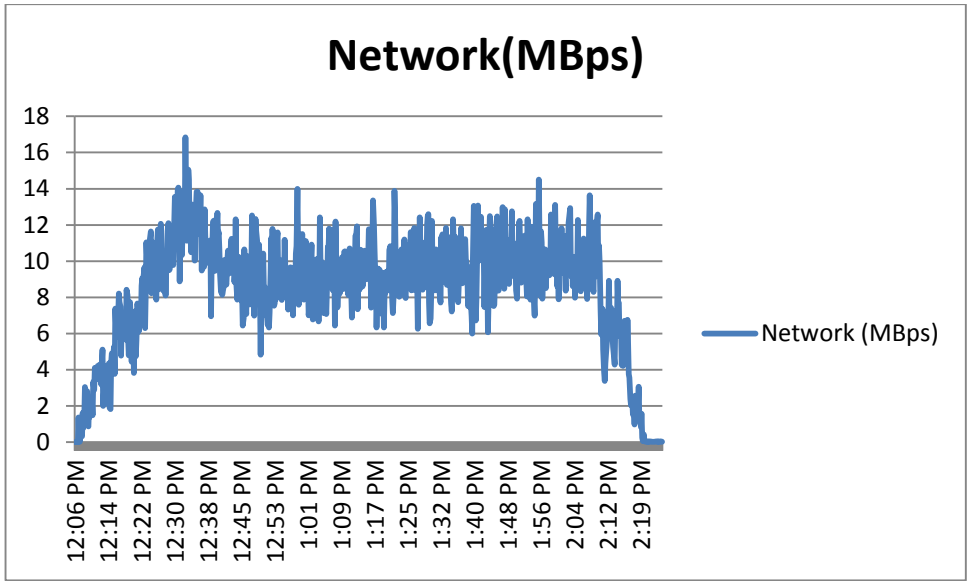

As seen from the above graph, overall network performance was good.

GPU Performance Analysis Results

Below is the GPU utilization for the 3 AMD S7000 cards installed on the server.

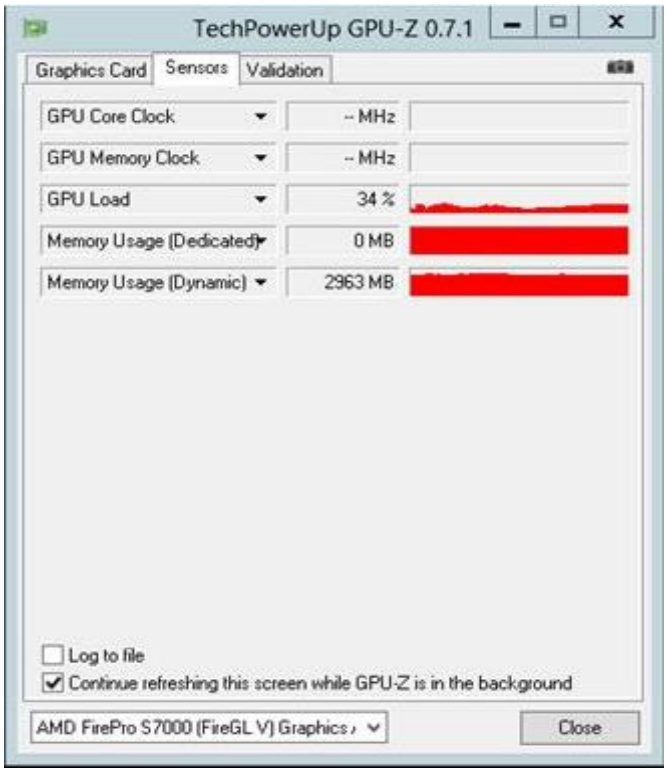

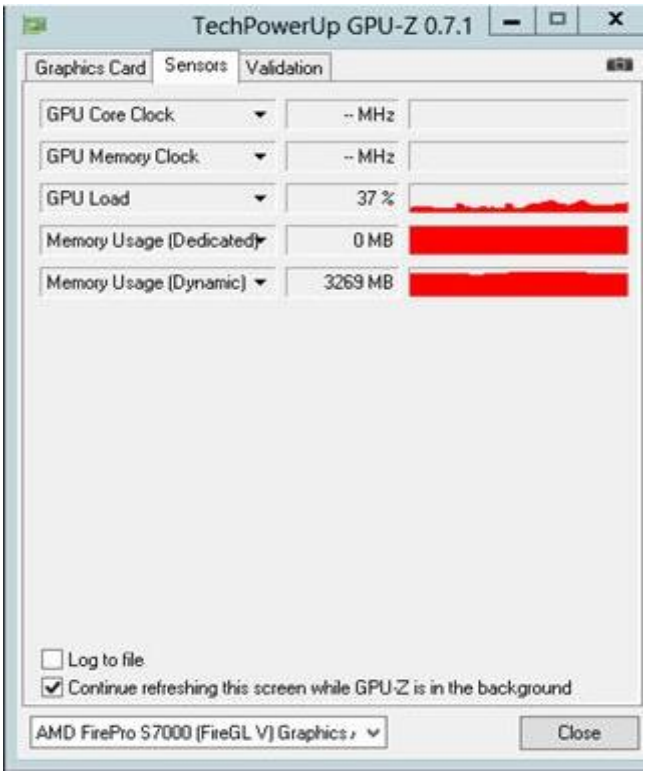

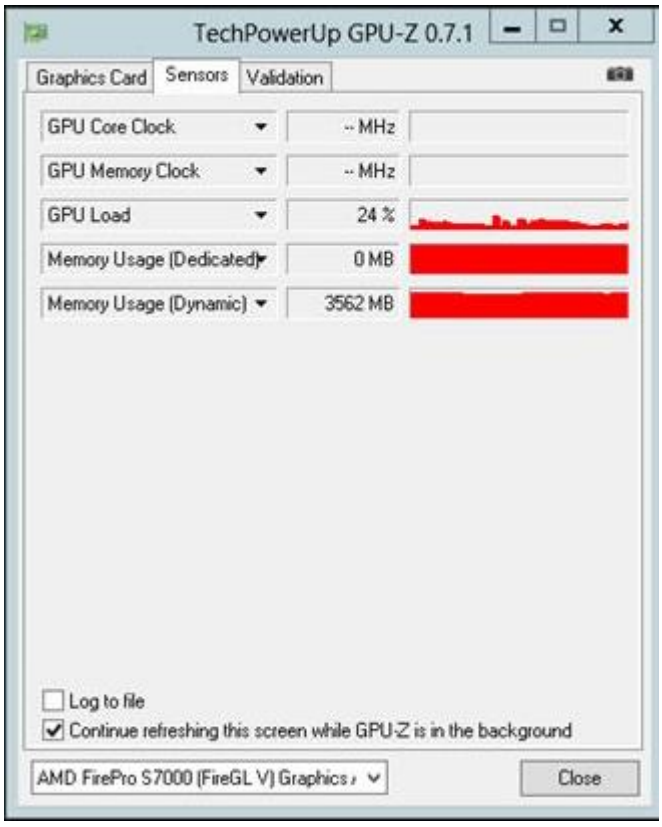

Below is the Process Explorer graph which shows the total GPU utilization during the test.

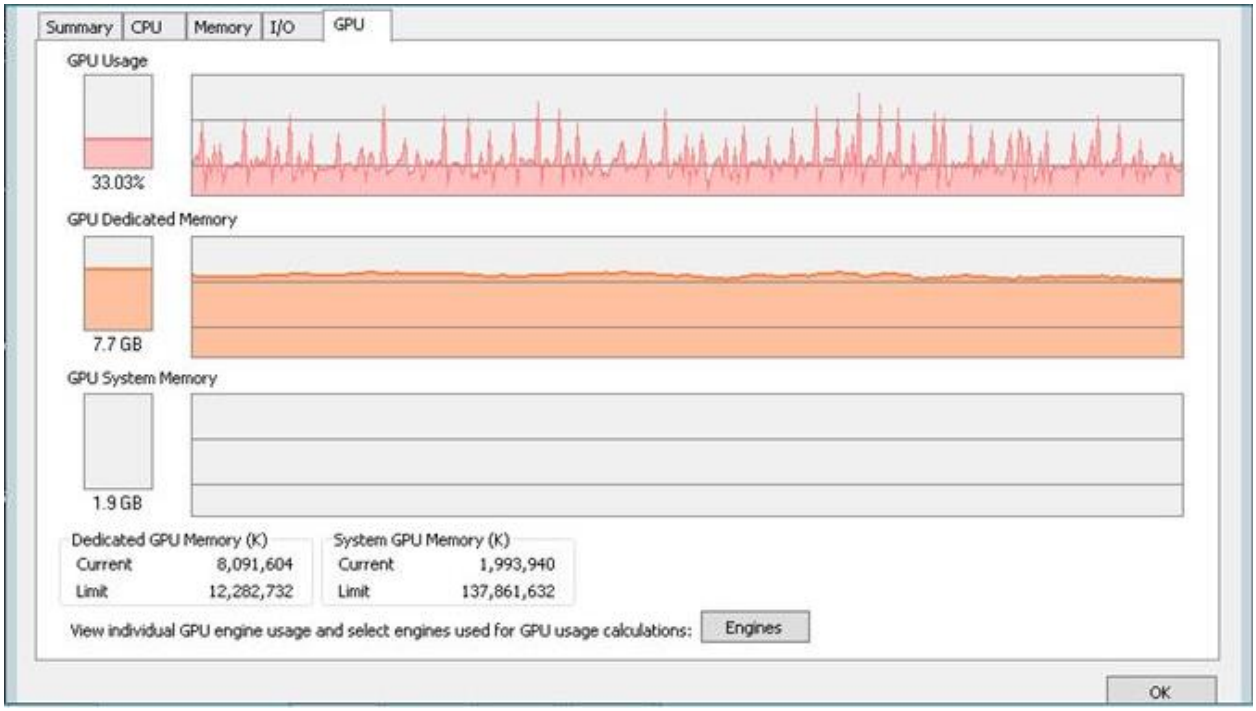

Graph below shows the RemoteFX Output FPS for all sessions.

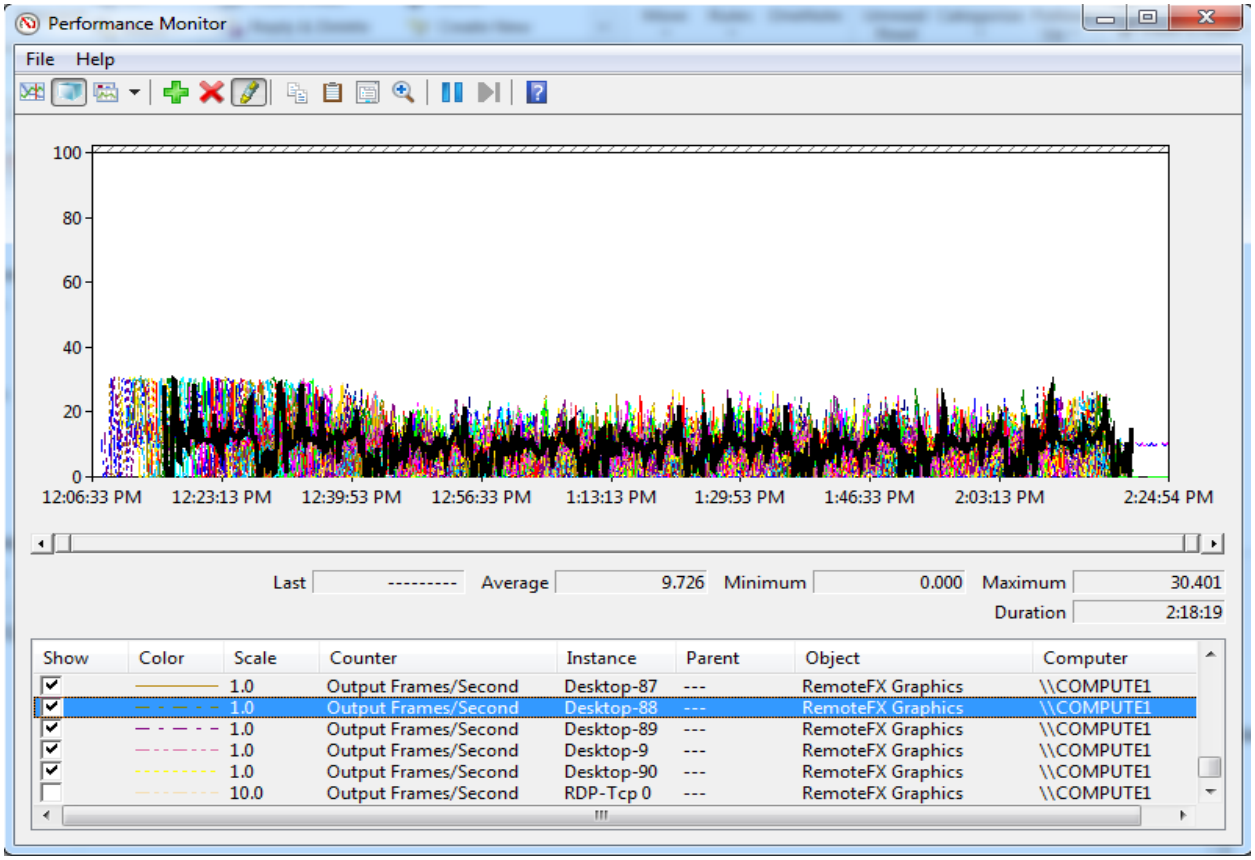

3D Graphics Workload Test

For the high end graphics testing, AMD S9000 GPU cards were used. Two tests were performed showing a comparison of the AMD S9000 against WARP (software based GPU) as well as performance of one AMD S9000 card versus two cards. For more information on Microsoft's Windows Advanced Rasterization Platform (WARP), please refer to: http://msdn.microsoft.com/en-us/library/windows/desktop/gg615082(v=vs.85).aspx

AMD S9000 driver version was 8.97.10.6

Two AMD S9000 GPU cards vs. WARP

The graphs below show that two AMD S9000 GPU cards will outperform WARP in terms of FPS and jitter. The frame rates for the AMD cards remains at an acceptable level above 20 FPS through 85 users launched, but jitter becomes unacceptable after 65 users therefore the threshold is set to 65 users for two AMD S9000 cards. In comparison, WARP maxes at 18 FPS from the start with a higher level of jitter. CPU usage remains at an acceptable level even when the numbers of users launched exceeds the FPS and jitter tolerance levels.

User Experience - Full screen modern app with 3D object

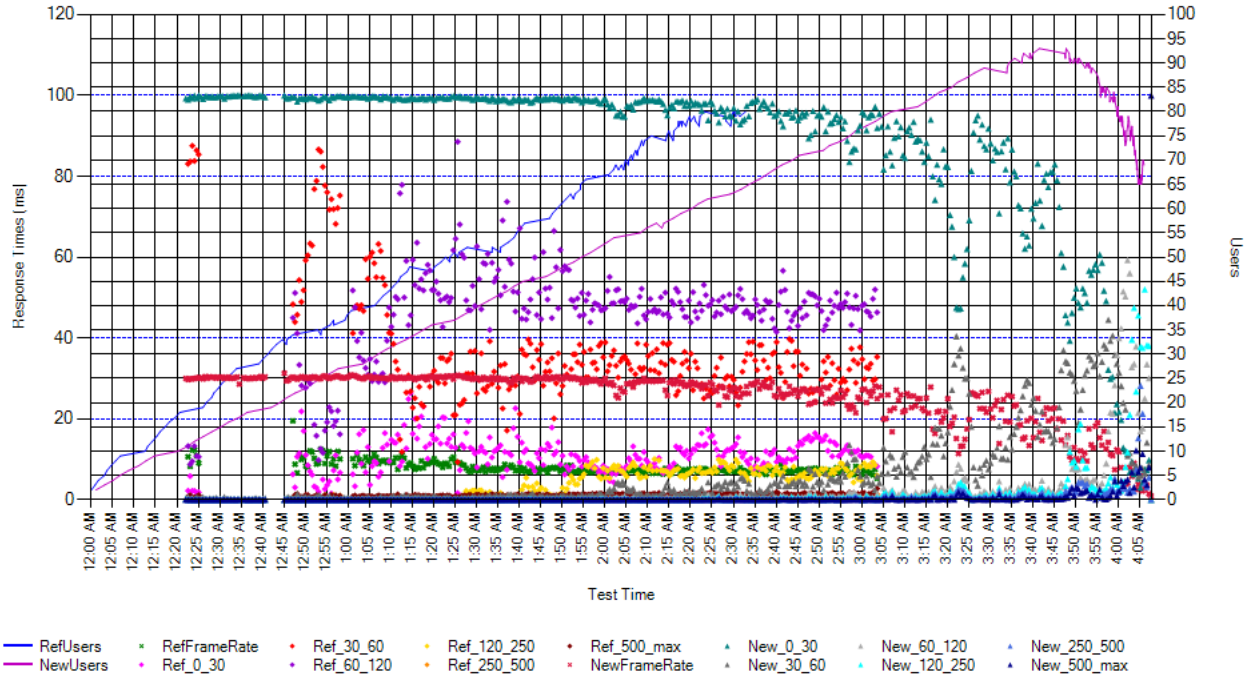

NOTE: Response times on left axis denotes frame rates

———	vGPU (2 x S9000) user progression
———	Warp user progression
*****	vGPU - Frame rates are above 20 FPS all the way to 85 users. Jitter is low and acceptable all the way to 65 users
*****	Warp – Frame rates are max 18 FPS from the start. Jitter is higher

CPU Usage:

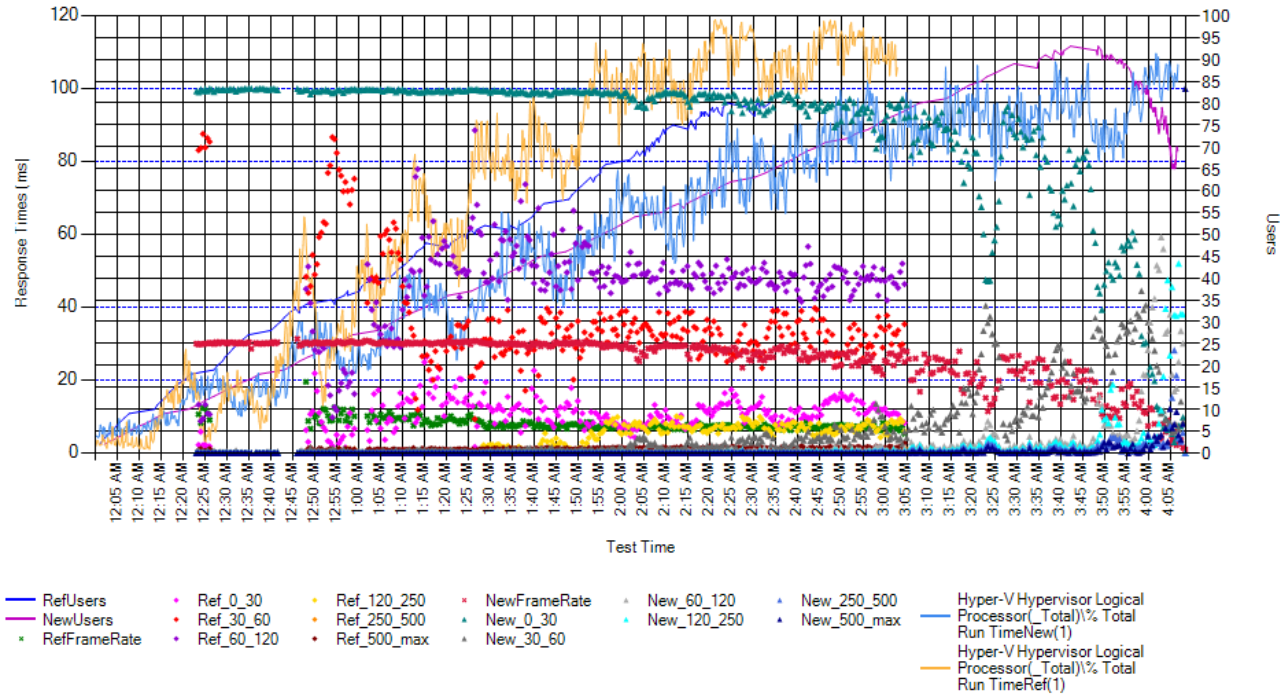

NOTE: Response times on left axis denotes frame rates

vGPU user progression
Warp user progression
At 85 users, vGPU CPU usage is ~80%
At 65 users, Warp CPU usage is ~80%

Results Summary Table

Workload	Server Density	CPU	Login StateTier-1 IOPS	Login State IOPS per User Tier-1	Steady State Tier-1 IOPS	Steady State IOPS per User Tier-1
Premium Plus	75	87%	1594	21.25	1136	15.14

D.4 DVS 10-Seat Trial Kit Test

Results Summary Table

Based on the server hardware configuration as described above in the DVS 10-Seat Trial Kit section, 10 basic Windows 8 users will experience excellent performance with additional resource headroom available in reserve. The consumption numbers below are based on average performance:

Task Worker Users	CPU (%)	RAM (GB Consumed)	Disk (IOPS)	Network (Kbps)
10	40	7.5	50	262

D.5 VRTX – ROBO Model

Configuration Summary

The following validations were performed using hardware as described in the **Appendix B – Remote Office/Branch Office** model. All nodes (servers) were PowerEdge M620 blades with dual 2.7GHz processors, 256GB RAM, and 2x146GB 15k local drives for OS. All blades accessed shared storage as described in the Appendix.

Tests were performed using two nodes and four nodes to determine density for both configurations. In all cases, all nodes were configured with both compute and management roles in a failover cluster. This is the recommended configuration and was done to determine the maximum user density in failover scenarios. For example, in a two node configuration, one node would have to handle the entire workload for the management VMs as well as virtual desktops if the other node experienced a failure.

Microsoft Windows 8 was used as the desktop operating system.

250 User (2 Node) Basic Workload Test

Since the goal was to determine the maximum density for a single node for failover purposes, all virtual desktops and management VMs executed on the same node for this test. The CPU usage for this test averaged between 85-90% during peak thus confirming that the 2 Node configuration can support up to 250 basic users as typical load would be distributed across two nodes under normal operation. Under normal operation with 125 basic users per node, the CPU utilization averages about 70% as seen with the 500 user testing below.

CPU Resource Utilization Performance Analysis Results

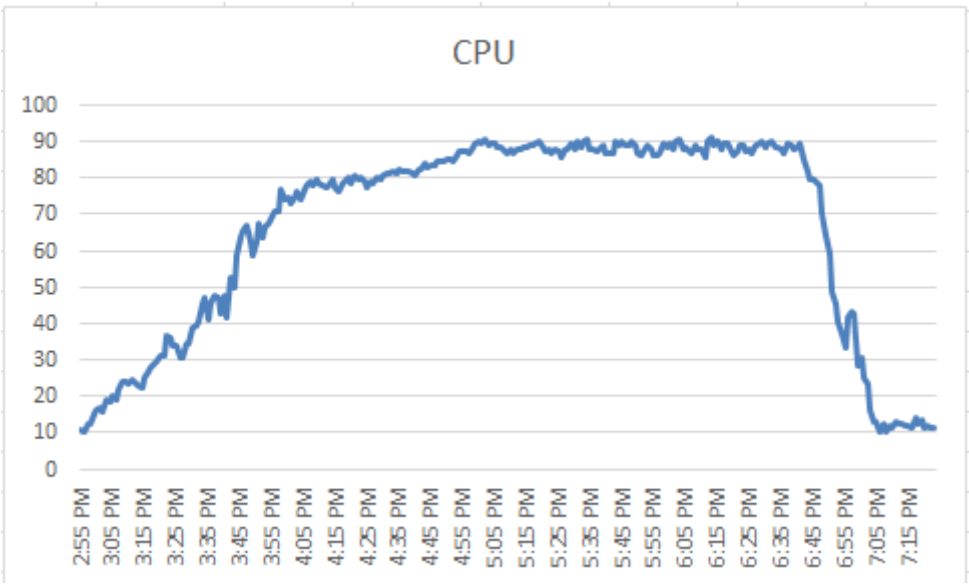

CPU load reached 90% during peak but was considered acceptable as the load would normally be distributed across two nodes during normal operations. The graph above shows the performance of one node to simulate a failover scenario – although the usage is high, it would be tolerable until the other node is repaired/replaced.

Memory Resource Utilization Performance Analysis Results

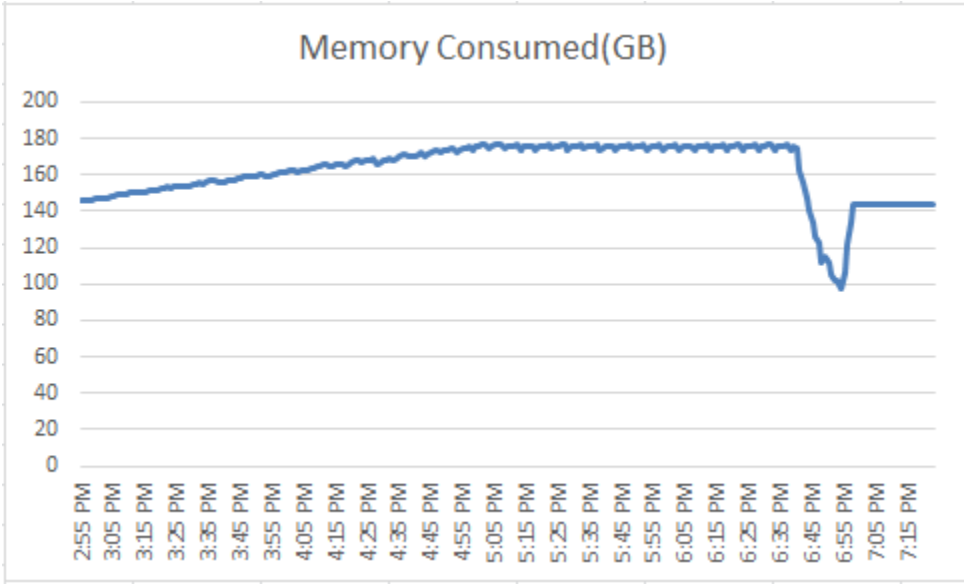

Memory consumed is well within amount of memory available.

Disk IO Resource Utilization Performance Analysis Results

This solution utilizes shared storage built in to the enclosure that is accessible by all of the nodes. Two volumes were created on the shared storage: one to house the virtual desktop disks and another for the management VM disks. The charts below show the disk IOPS for the volumes while the system was under test.

NOTE: CSV Cache was enabled resulting in low Read IO from disk as represented in the graphs below. Refer to the CSV Caching section below for more information.

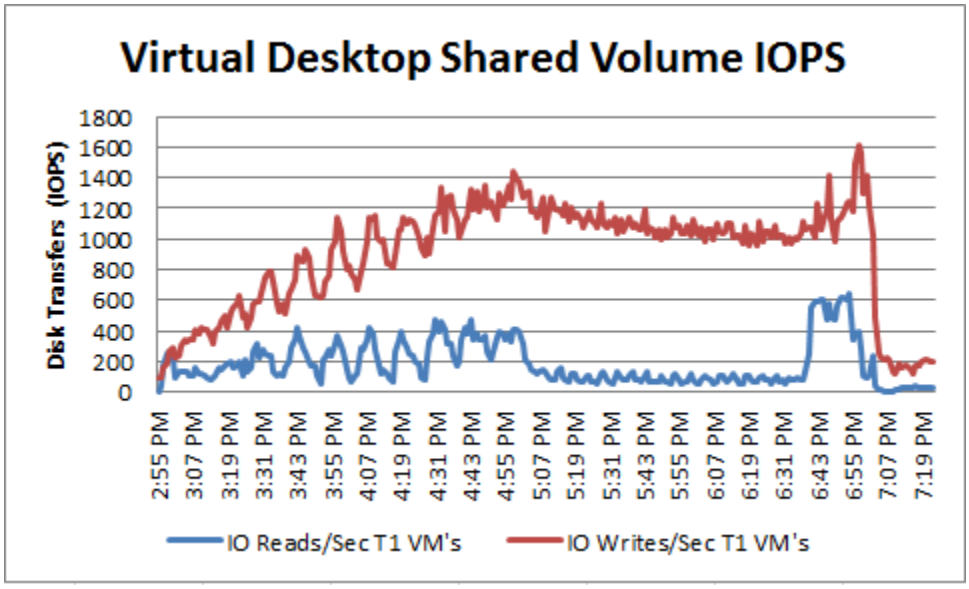
Virtual Desktop Shared Volume IOPS
Disk Transfers (IOPS)
1800
1600
1400
1200
1000
800
600
400
200
0
2:55 PM
3:07 PM
3:19 PM
3:31 PM
3:43 PM
3:55 PM
4:07 PM
4:19 PM
4:31 PM
4:43 PM
4:55 PM
5:07 PM
5:19 PM
5:31 PM
5:43 PM
5:55 PM
6:07 PM
6:19 PM
6:31 PM
6:43 PM
6:55 PM
7:07 PM
7:19 PM
IO Reads/Sec T1 VM's
IO Writes/Sec T1 VM's

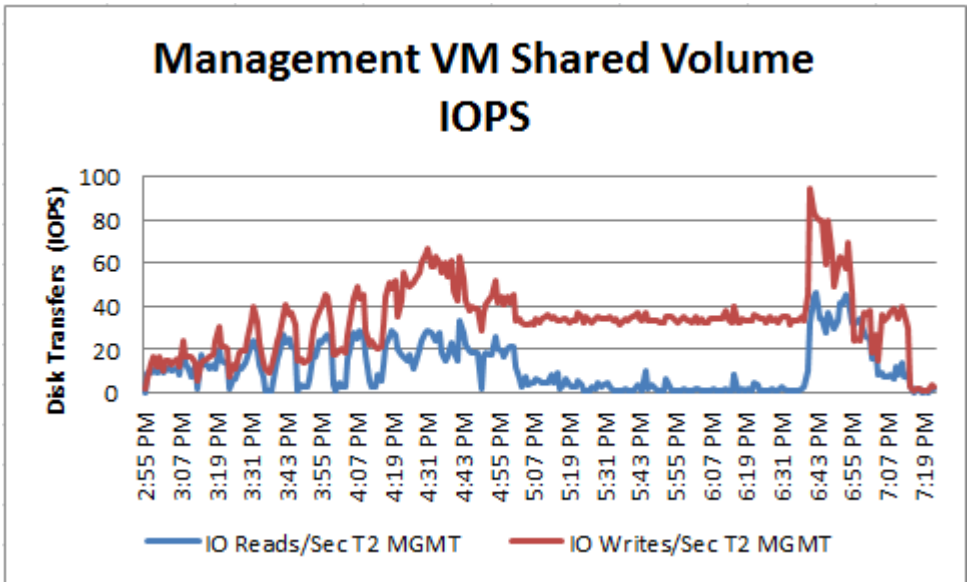
Management VM Shared Volume IOPS
Disk Transfers (IOPS)
100
80
60
40
20
0
2:55 PM
3:07 PM
3:19 PM
3:31 PM
3:43 PM
3:55 PM
4:07 PM
4:19 PM
4:31 PM
4:43 PM
4:55 PM
5:07 PM
5:19 PM
5:31 PM
5:43 PM
5:55 PM
6:07 PM
6:19 PM
6:31 PM
6:43 PM
6:55 PM
7:07 PM
7:19 PM
IO Reads/Sec T2 MGMT
IO Writes/Sec T2 MGMT

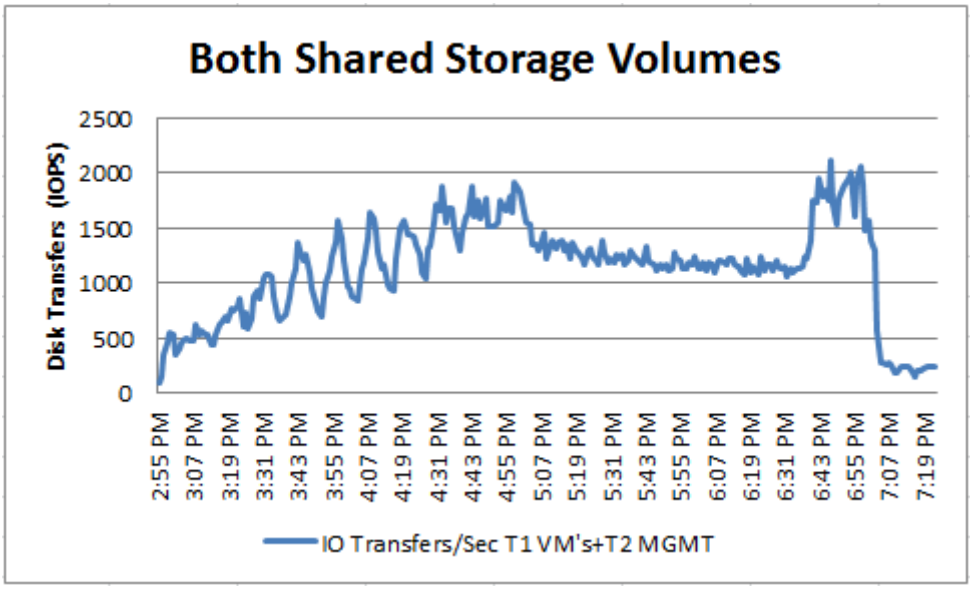

Network Resource Utilization Performance Analysis Results

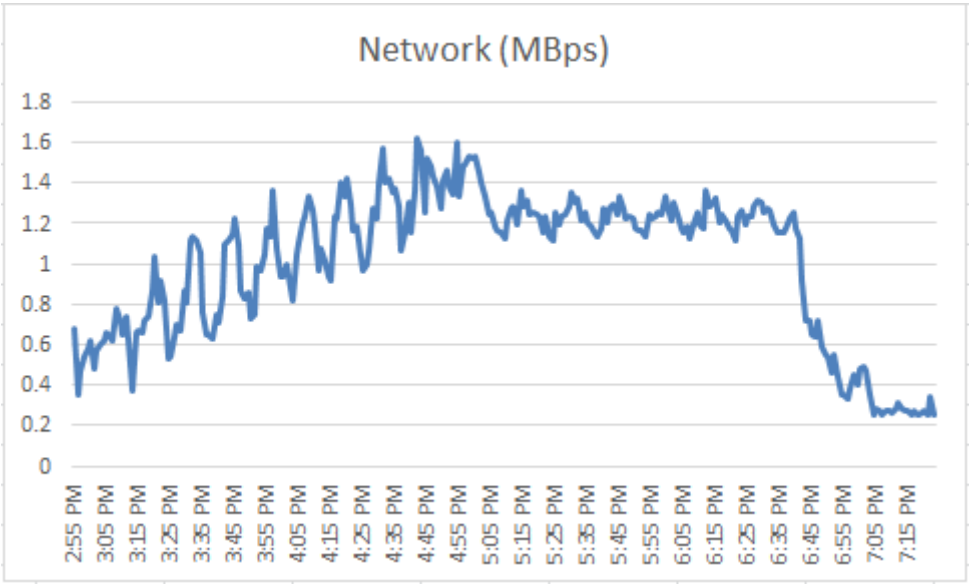

As seen from the above graph, overall network performance was good.

500 User (4 Node) Basic Workload Test

The CPU usage for this test averaged around 70% for each node during peak thus confirming that the 4 Node configuration can support up to 500 basic users.

CPU Resource Utilization Performance Analysis Results

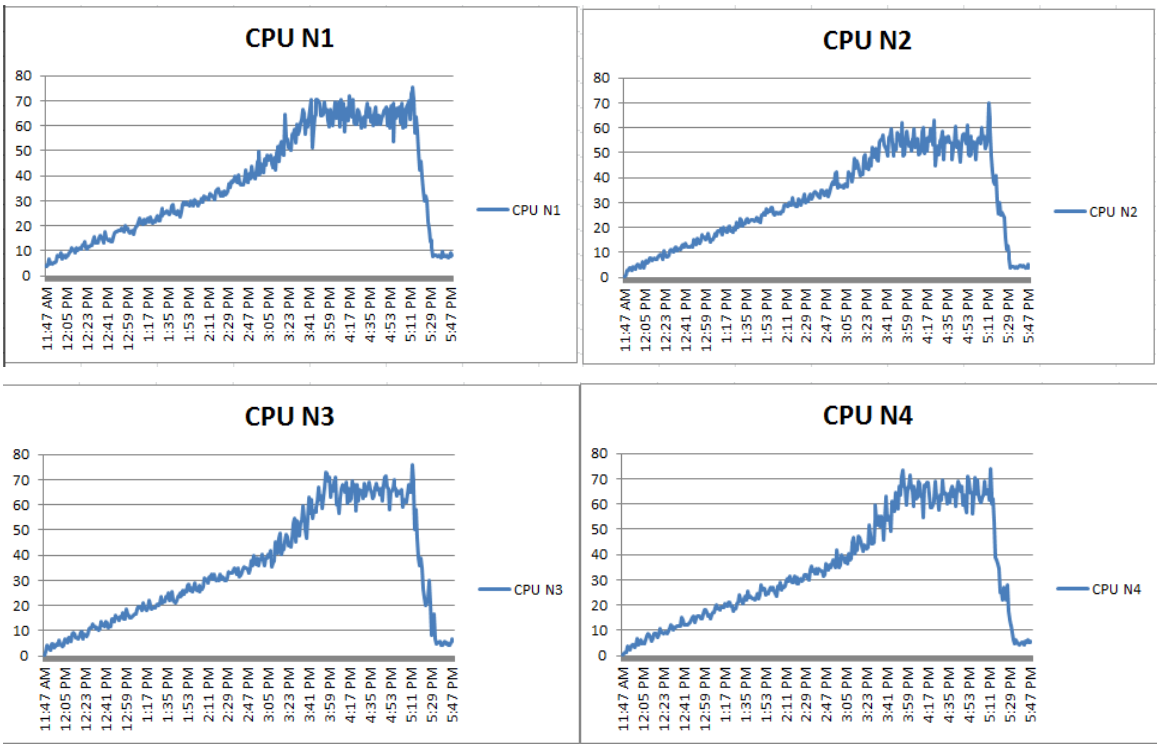

CPU usage for averaged around 70% for each node during peak load with maximums under 80%.

Memory Resource Utilization Performance Analysis Results

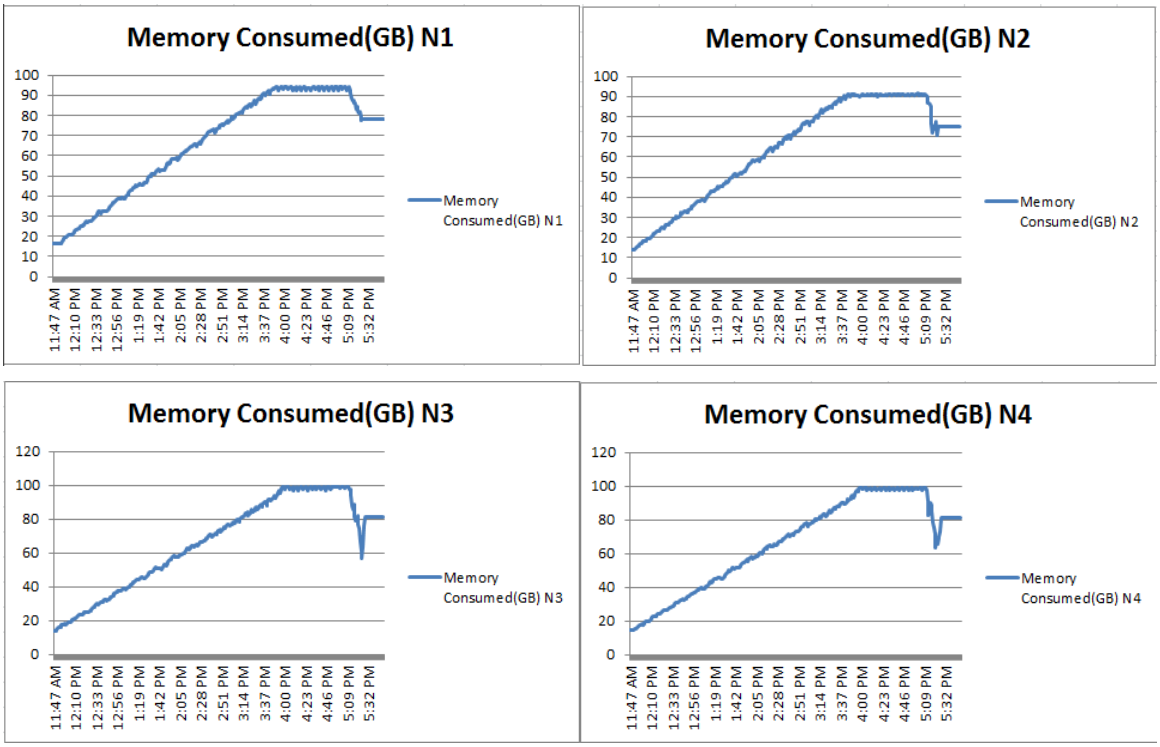

Memory consumed averaged around 95GB per node during peak load.

Disk IO Resource Utilization Performance Analysis Results

This solution utilizes shared storage built in to the enclosure that is accessible by all of the nodes. Two volumes were created on the shared storage: one to house the virtual desktop disks and another for the management VM disks. The charts below show the disk IOPS for the volumes while the system was under test.

NOTE: CSV Cache was enabled resulting in low Read IO from disk as represented in the graphs below. Refer to the CSV Caching section below for more information.

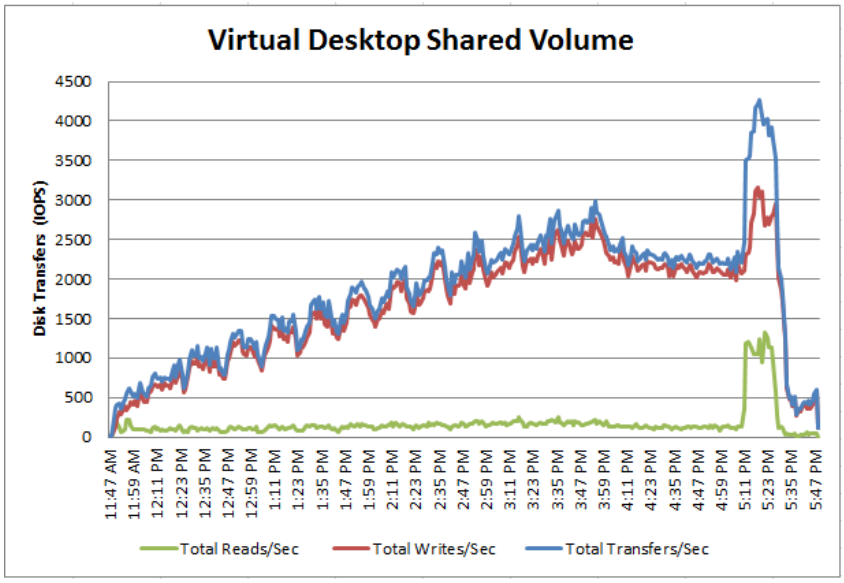

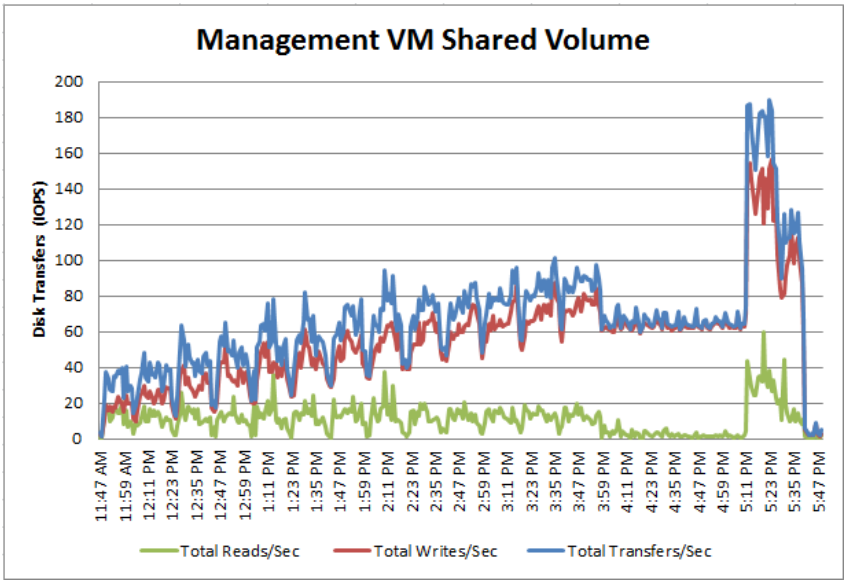

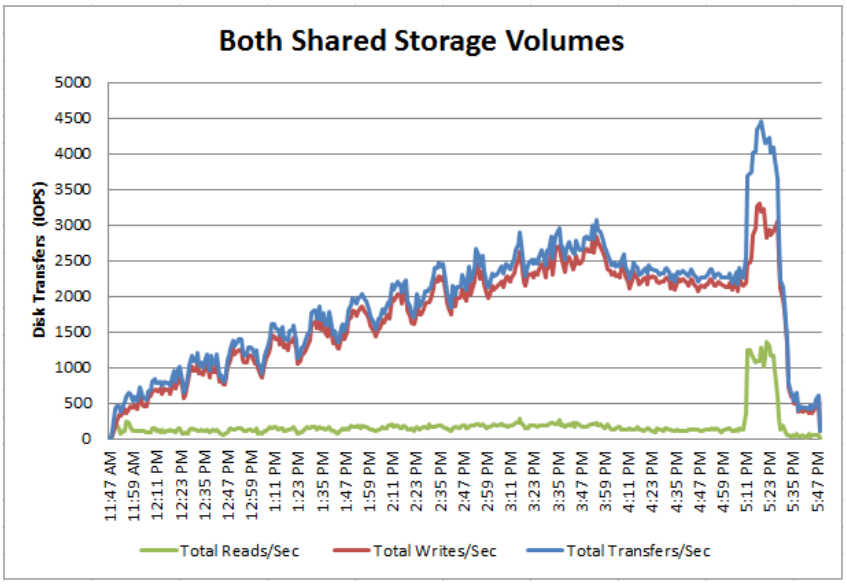

Network Resource Utilization Performance Analysis Results

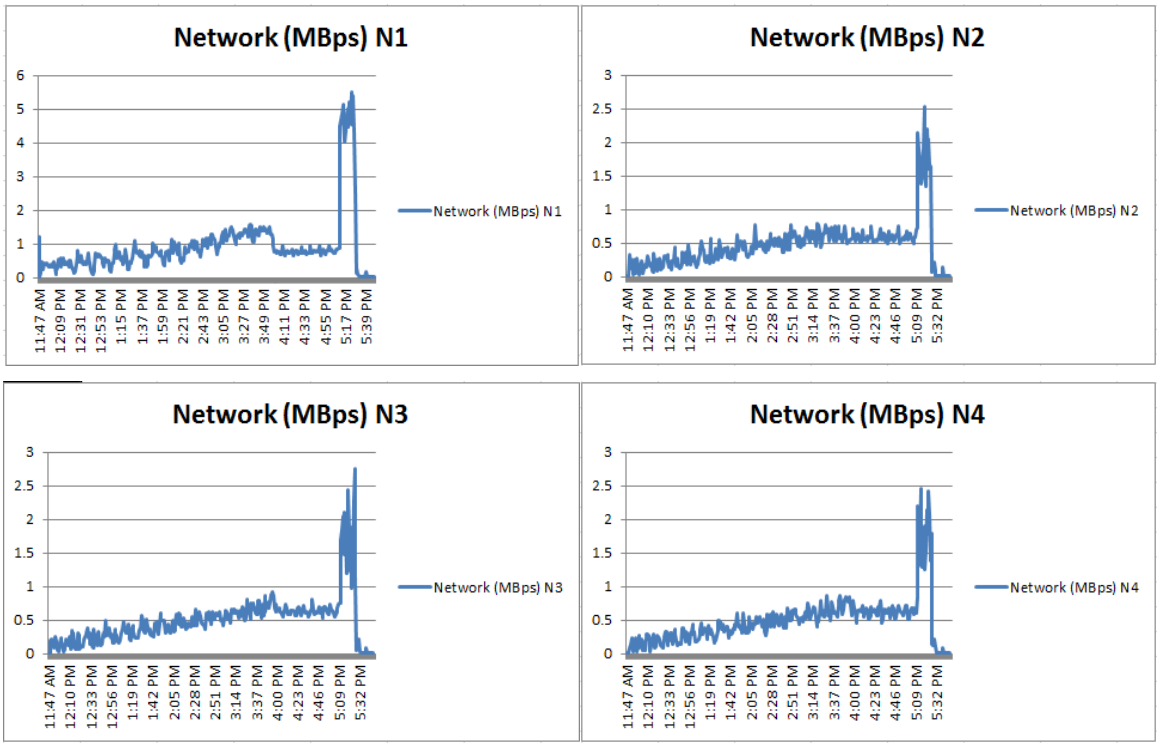

As seen from the graphs above, there was an increase in network traffic on all nodes during logoff, but overall, network performance was good.

CSV Caching

CSV cache was enabled for the shared storage volumes for both user tests. For the 2 Node test, the cache size was set to 512MB (default) per node and for the 4 Node test, it was set to 2048MB

per node. Additional testing will need to be performed to fully understand the impact of CSV caching and to determine proper guidance when setting the cache size. However, as seen in the graphs below, the read IO benefit achieved from CSV caching can be extensive depending on the read intensive workload.

For more information on CSV cache, refer to:
http://blogs.msdn.com/b/clustering/archive/2012/03/22/10286676.aspx

Results below are from using 512MB cache size per node during the 2 Node test. One graph shows the total Read IOPS, the read IOPS satisfied from cache, and read IOPS satisfied from disk activity. The other graph shows the read IO satisfied from cache as a percentage of total read IOPS.

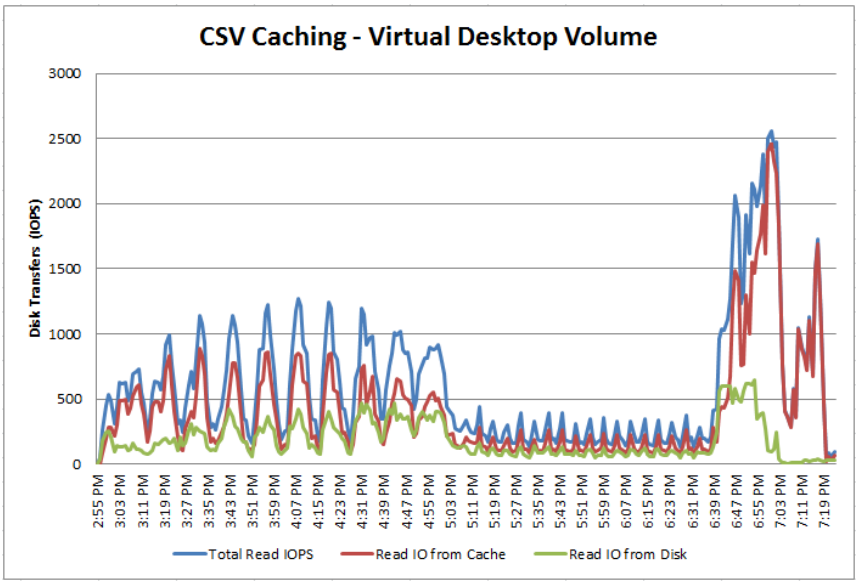

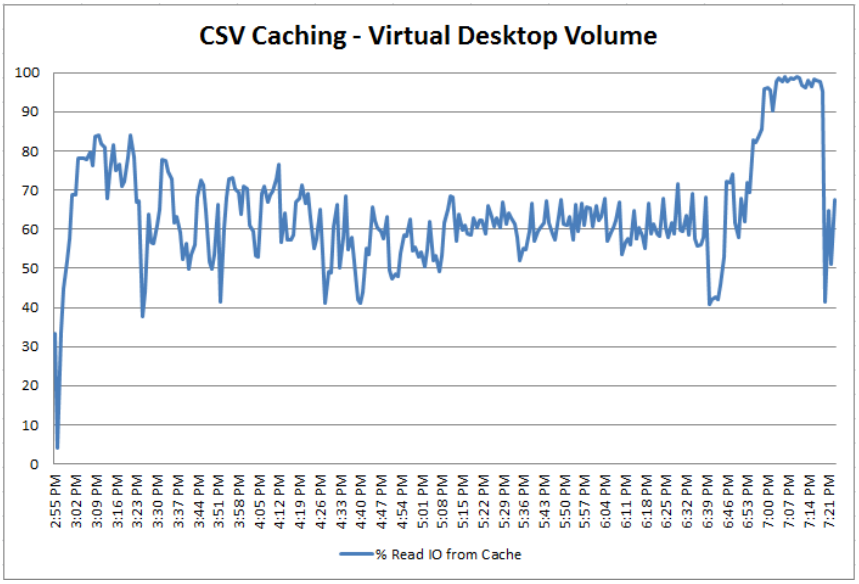

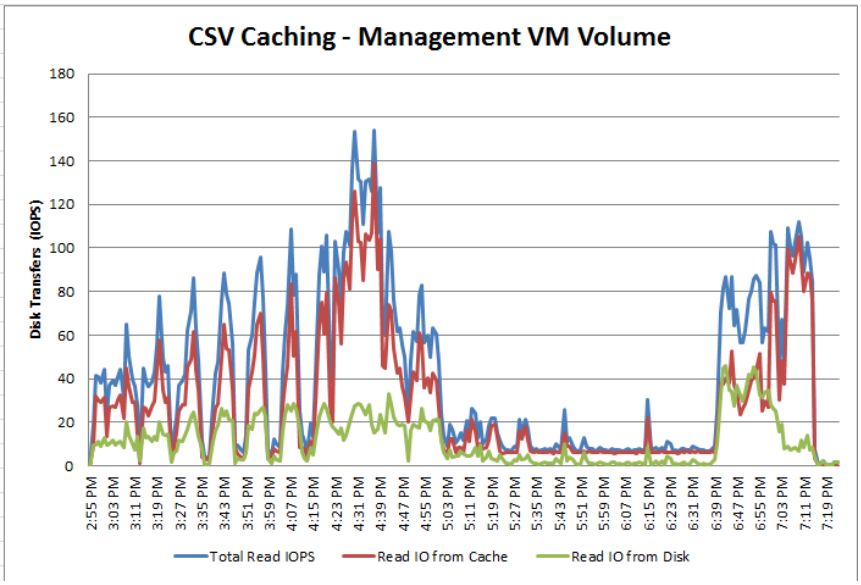

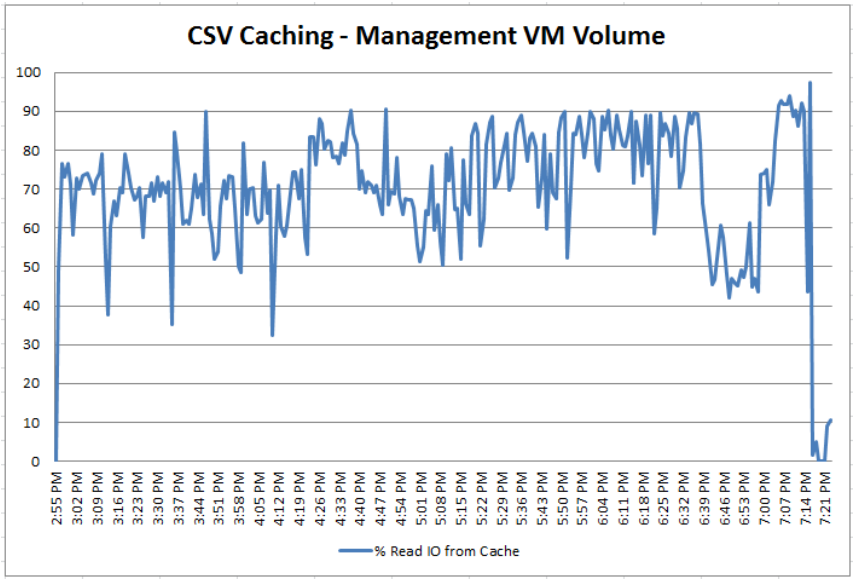

Failover Testing

To observe the effects of a single node failure, a test was performed using a 4 node configuration. A 60 user basic workload was started with desktops distributed evenly across all nodes (15 per node). All management VMs were running on the node that was designated as the "fail test" node. To begin, the "fail test" node was ejected from the enclosure without proper shutdown effectively simulating a hard failure. The 15 virtual desktop VMs running on the failed node were distributed and restarted across the 3 remaining nodes within 2 minutes. Since they are restarted, connectivity to these VMs as well as their running state is lost. The management VMs were also restarted on a remaining node. Within 4 minutes, the management VMs were started and all services on those VMs were up and ready to respond to client requests.

As can be expected and seen in the graphs below, resource utilization will increase on the active nodes to handle the load of the failed node.

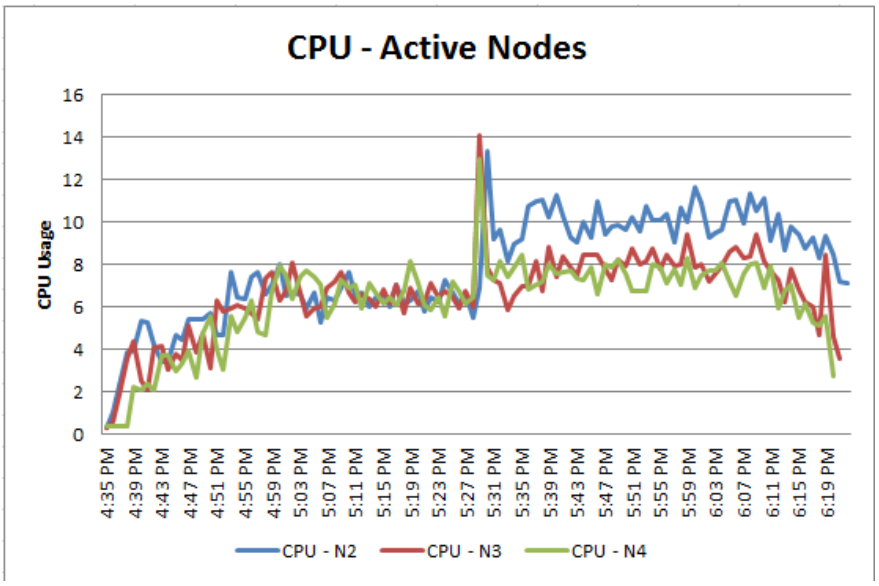
CPU - Active Nodes
CPU Usage
16
14
12
10
8
6
4
2
0
4:35 PM
4:39 PM
4:43 PM
4:47 PM
4:51 PM
4:55 PM
4:59 PM
5:03 PM
5:07 PM
5:11 PM
5:15 PM
5:19 PM
5:23 PM
5:27 PM
5:31 PM
5:35 PM
5:39 PM
5:43 PM
5:47 PM
5:51 PM
5:55 PM
5:59 PM
6:03 PM
6:07 PM
6:11 PM
6:15 PM
6:19 PM
CPU - N2
CPU - N3
CPU - N4

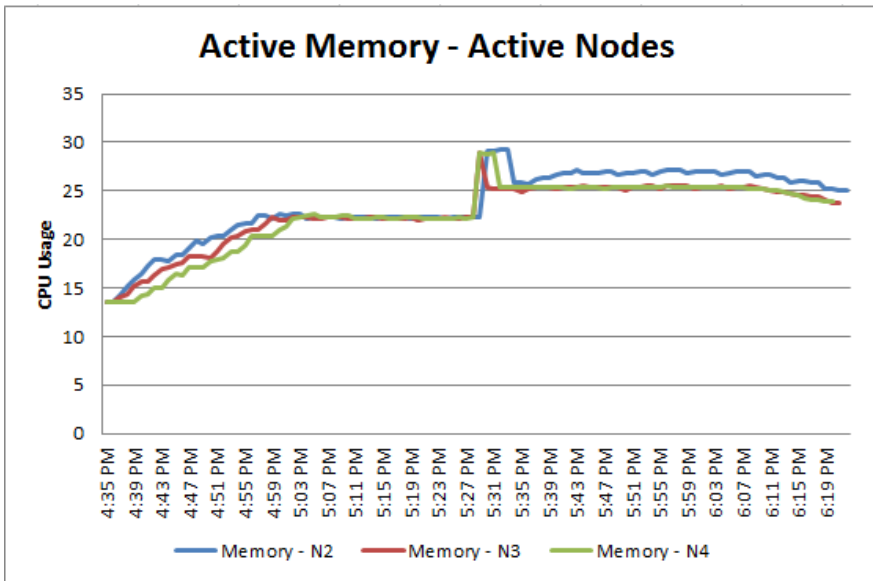
Active Memory - Active Nodes
CPU Usage
35
30
25
20
15
10
5
0
4:35 PM
4:39 PM
4:43 PM
4:47 PM
4:51 PM
4:55 PM
4:59 PM
5:03 PM
5:07 PM
5:11 PM
5:15 PM
5:19 PM
5:23 PM
5:27 PM
5:31 PM
5:35 PM
5:39 PM
5:43 PM
5:47 PM
5:51 PM
5:55 PM
5:59 PM
6:03 PM
6:07 PM
6:11 PM
6:15 PM
6:19 PM
Memory - N2
Memory - N3
Memory - N4

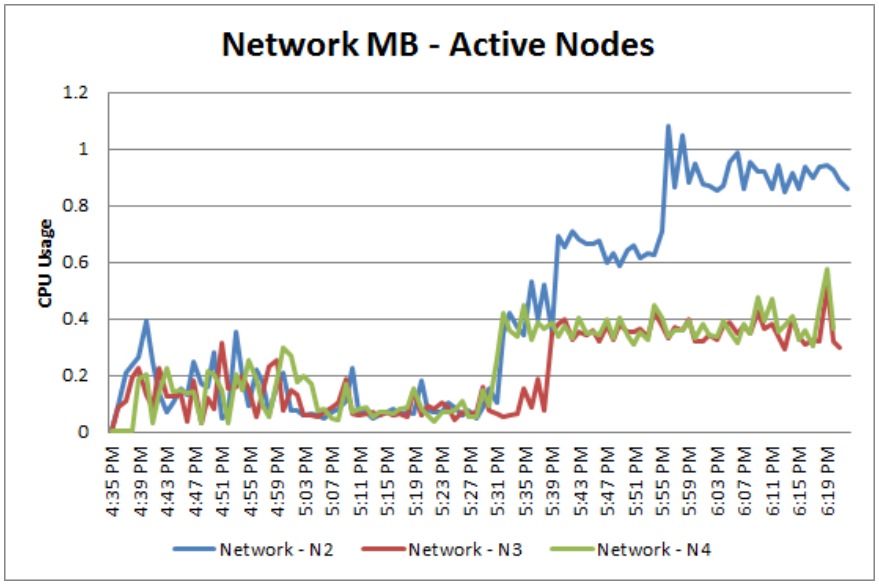

The values above are just to illustrate increases at the time one of the nodes fails. Resource utilization will vary depending on the number of virtual desktops and workload being executed.

Results Summary Table

Workload	Server Density	CPU	Login State Shared Storage IOPS	Login State IOPS per User	Steady State Shared Storage IOPS	Steady State IOPS per User
2 Node – Basic Users	250	70%	1914	7.65	1381	5.52
4 Node – Basic Users	500	70%	3442	6.88	2561	5.12

Appendix E – Shared (RDSH) Performance Analysis Results

E.1 Distributed Server – Dedicated Compute

Configuration Summary

The following validations were performed using hardware as described in the **Local Tier 1 – Base** model. The compute role was dedicated to a single R720 host with dual 2.9 GHz processors, 128GB RAM, and 12x300GB 15k local hard drives configured in a single RAID 10 array. The management roles were installed on VMs running on a separate R720 host.

Microsoft Windows 8 was used as the desktop operating system.

The table shown below summarizes the important configuration information in relation to the RDSH environment on the compute host.

	Hyper-V Compute Host	**RDSH VM1**	**RDSH VM2**	**RDSH VM3**	**RDSH VM4**
CPU Resource	32 logical cores (16 physical cores)	8 vCPUs	8 vCPUs	8 vCPUs	8 vCPUs
Memory (GB)	128	Dynamic 16GB (max 31GB)	Dynamic 16GB (max 31GB)	Dynamic 16GB (max 31GB)	Dynamic 16GB (max 31GB)
RDSH Session Capacity	300	75	75	75	75

Basic Workload Test

The CPU usage for this test reached 75% thus confirming that the server with this configuration can support up to 300 basic users.

CPU Resource Utilization Performance Analysis Results

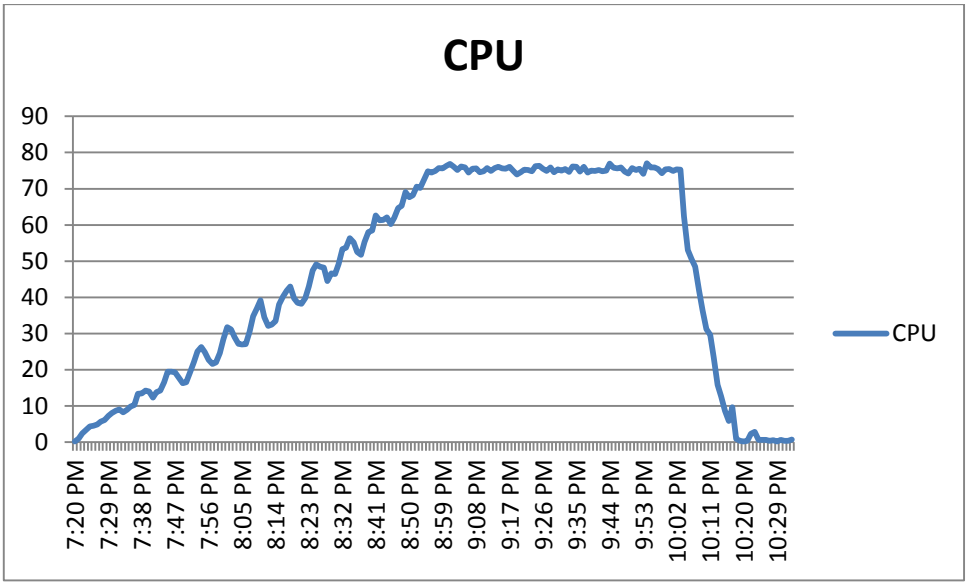

As seen from the above graph, maximum CPU utilization was approximately 75%. After all users started logging in, CPU usage spiked from 0 to 58% and became stable between 74-76% once all users logged in and dropped as the users logged off.

Memory Resource Utilization Performance Analysis Results

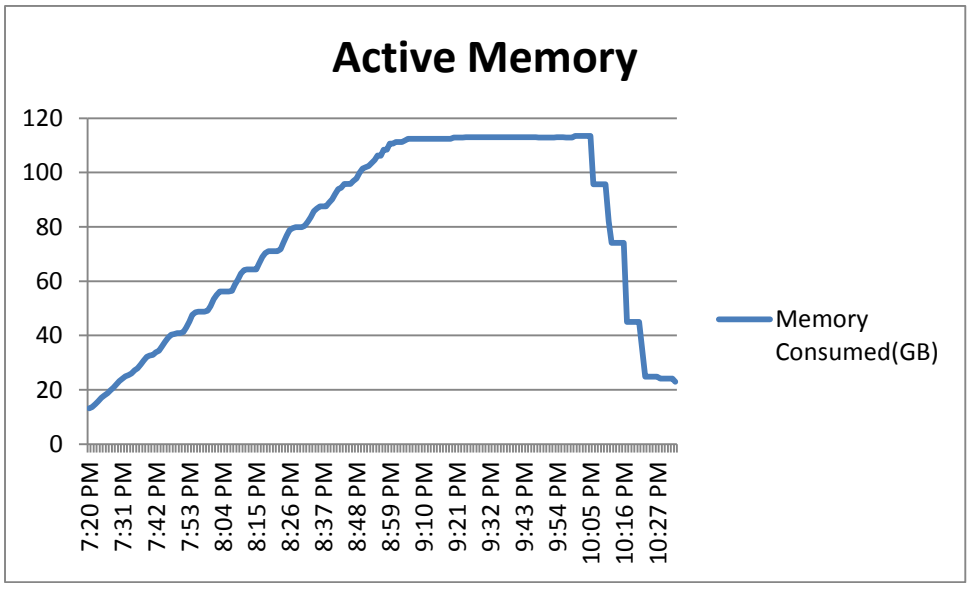

A spike is evident after all VMs start logging in; however, active memory is about 88.28% of available memory during the peak of validation activity.

Disk IO Resource Utilization Performance Analysis Results

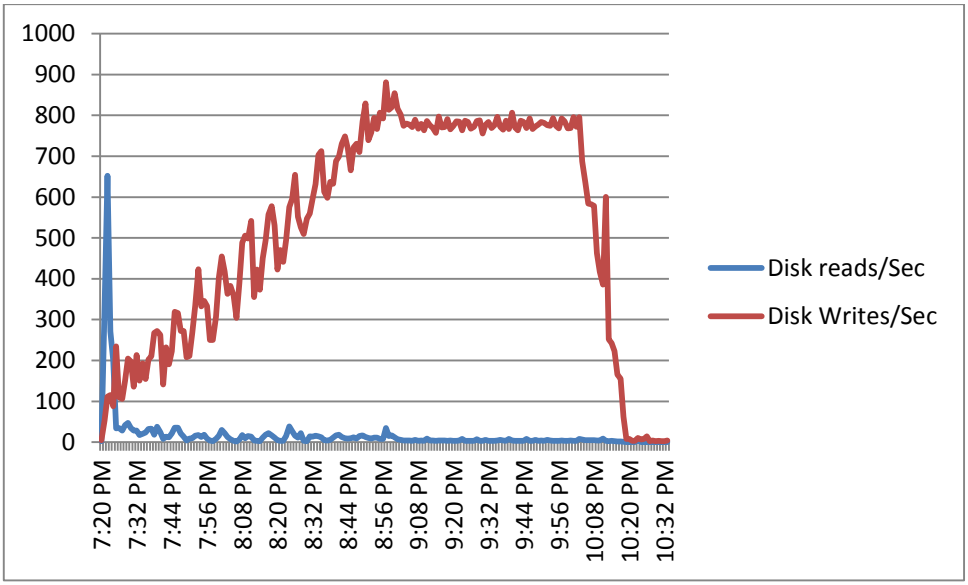

A spike is evident after all VMs start logging in; however, the read/write activity was fluctuating from 7:20 PM to 09:00 PM and became stable after all users finished logging in from 09:04 PM to 10:02 PM. Total Peak IOPS measured during the login state was 916 giving an IOPS value of 3.05 per user and total peak IOPS in steady state (after all users logged in) was 799 giving an IOPS value of 2.66 per user.

Network Resource Utilization Performance Analysis Results

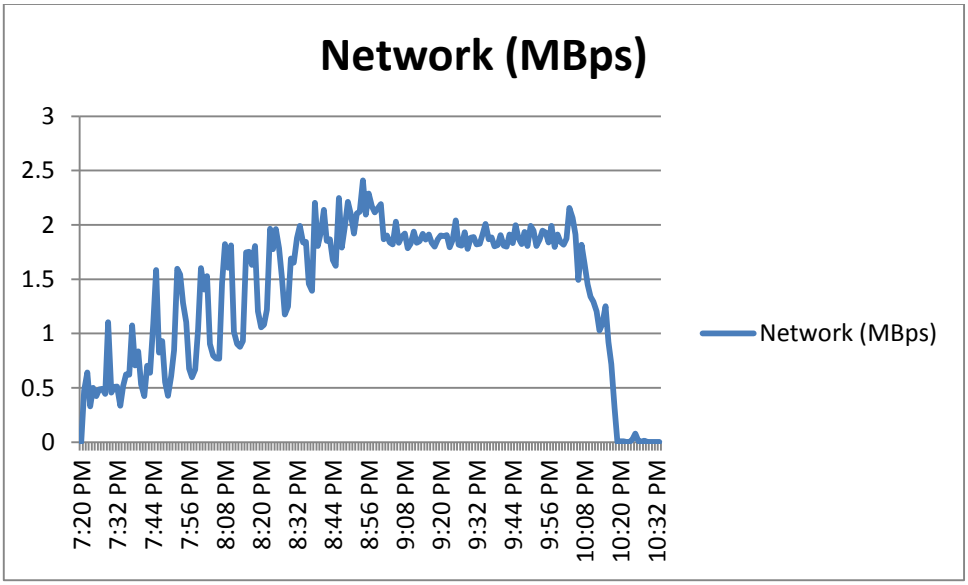

As seen from the above graph, overall network performance was good.

Standard Workload Test

Configuration Summary

The table shown below summarizes the important configuration information in relation to the knowledge worker RDSH environment.

	Hyper-V Host	RDSH VM1	RDSH VM2	RDSH VM3	RDSH VM4
CPU Resource	32 logical cores (16 physical cores)	8 vCPUs	8 vCPUs	8 vCPUs	8 vCPUs
Memory (GB)	96	16	16	16	16
RDSH Session Capacity	200	50	50	50	50

CPU Resource Utilization Performance Analysis Results

The Hyper-V CPU utilization graph shown below shows sustained CPU utilization peaking at approximately 72%, illustrating the capacity of the solution to host 200 users with headroom for additional spikes in user activity. It should be noted that only the "Hyper-V Hypervisor Logical Processor % Total Run Time" is shown but the "Hyper-V Hypervisor Virtual Processor % Total Run Time" graph is very similar, peaking at 71%; a close match between these parameters is indicative of a well-tuned environment.

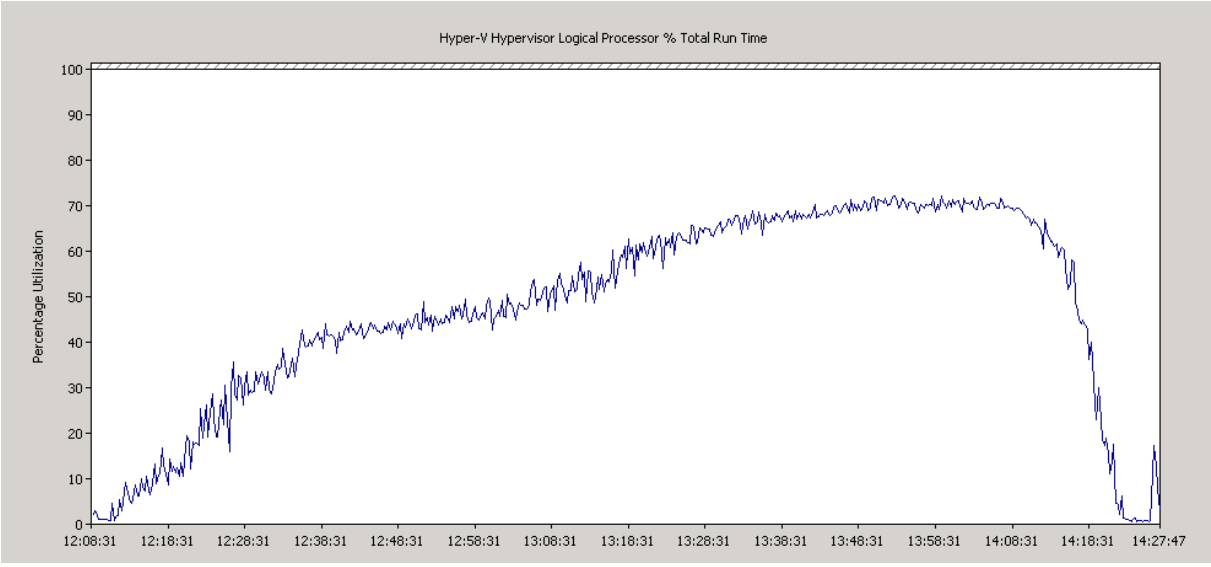

The table shown below shows sustained CPU utilization level at peak load for each of the RDSH VMs. It can be seen that sustained CPU utilization level lies between 85 and 88% for all RDSH VMs; this is a level that represents an ideal combination of maximizing resource utilization while also allowing headroom for spikes in user activity.

	Sustained CPU at Peak Load (%)

RDSH VM1	88
RDSH VM2	88
RDSH VM3	86
RDSH VM4	85

Memory Resource Utilization Performance Analysis Results

The memory graph shown below shows available memory during performance analysis at the Hyper-V hypervisor level. It can be seen that available memory never drops below approximately 16GB at the hypervisor level, which illustrates that sufficient memory headroom exists in the environment: maximum memory utilization is approximately 83%, since the total memory on the compute host is 96GB.

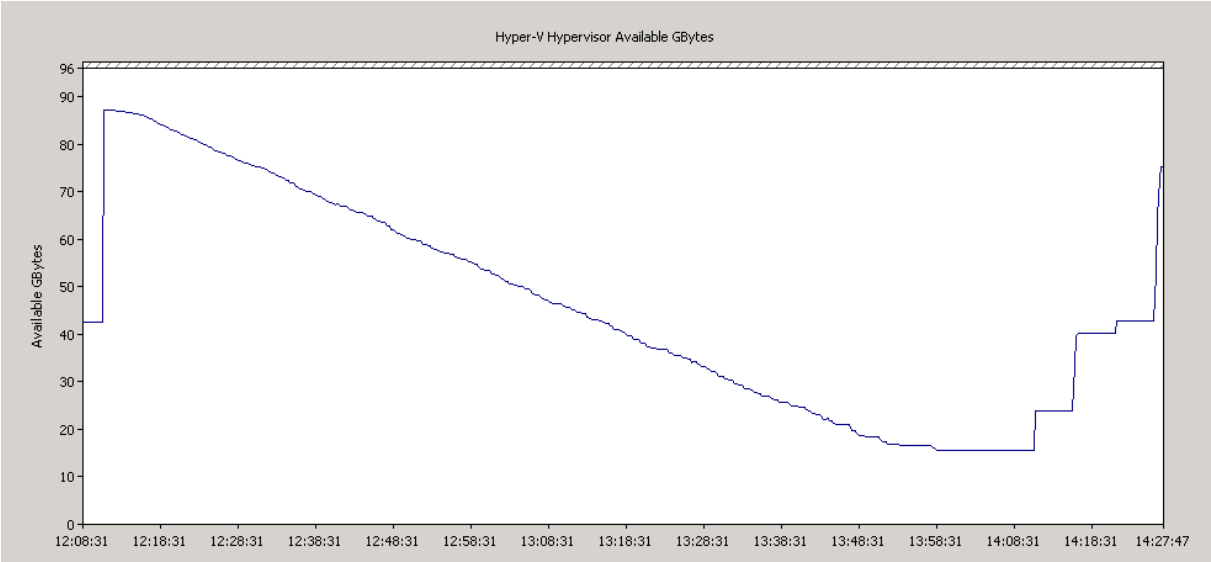

The table shown below shows available memory on the RDSH VMs during performance analysis. It can be seen that available memory at peak load across the 4 VMs doesn't drop below 3.4GB, which illustrates that sufficient memory headroom exists in the environment: maximum memory utilization is approximately 78%, since the total memory on each VM is 16GB.

	Available Memory at Peak Load GB)
RDSH VM1	41
RDSH VM2	23
RDSH VM3	21
RDSH VM4	28

Disk IO Resource Utilization Performance Analysis Results

The graph shown below shows disk IO on the D: partition of the Windows Server 2012 system that hosts the Hyper-V hypervisor (this is the partition where the RDSH VMs reside). Sustained IOPS at peak load on the hypervisor partition reached a level of less than 100.

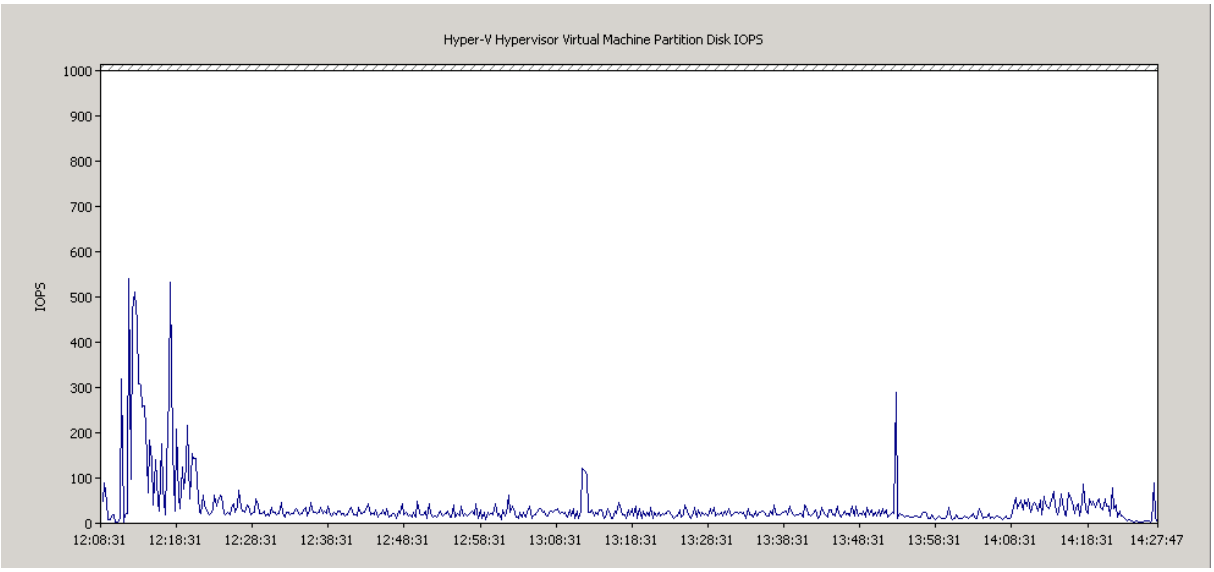

The table shown below shows sustained IOPS at peak load on the RDSH VMs during performance analysis. Peak utilization of 130 IOPS (on RDSH VM1) corresponds with approximately 2.6 IOPS per user.

	Sustained IOPS at Peak Load
RDSH VM1	130
RDSH VM2	115
RDSH VM3	125
RDSH VM4	100

The relevant Hyper-V partition latency information is shown below, it can be seen that latency remains significantly below the 15ms threshold specified by Microsoft in the TechNet best practices article described above for the entire duration of the test.

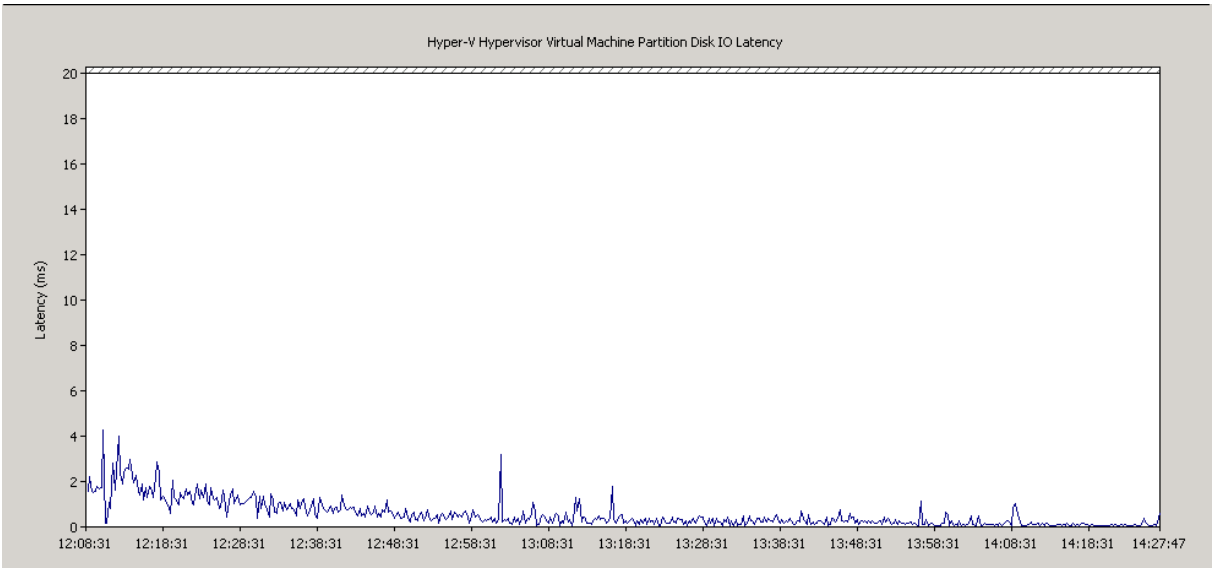

The table shown below shows sustained latency during performance analysis for each of the RDSH VMs; it can be seen that this doesn't exceed 5ms at any stage during performance analysis, indicating that disk IO latency is not a performance bottleneck.

	Sustained Latency at Peak Load (ms)
RDSH VM1	4
RDSH VM2	4
RDSH VM3	5
RDSH VM4	4

Network Resource Utilization Performance Analysis Results

Network resource utilization in the RDSH environment was seen to be low. At the hypervisor level, network resource utilization reaches a maximum of approximately 31 MB/s, representing network bandwidth utilization of approximately 25%, based on a 1Gb/s switching infrastructure. The hypervisor shows zero output queue length throughout the duration of the performance analysis activity. Individual RDSH VM network utilization is consistent with this overall hypervisor network utilization (i.e. total VM network utilization equals hypervisor network utilization) and all of the VMs also show zero output queue length throughout the duration of performance analysis. These hypervisor and RDSH VM network utilization results illustrate that network resource utilization is very low and does not approach a level at which it could become a bottleneck in the environment.

Results Summary Table

Workload	Server Density	CPU	Login StateTier-1 IOPS	Login State IOPS per User Tier-1	Steady State Tier-1 IOPS	Steady State IOPS per User Tier-1
RDSH Basic	300	75%	916	3.05	799	2.66
RDSH Standard	200	72%	470	2.35		

Appendix F – OS and MS Office Comparisons

Office 2010 vs. Office 2013

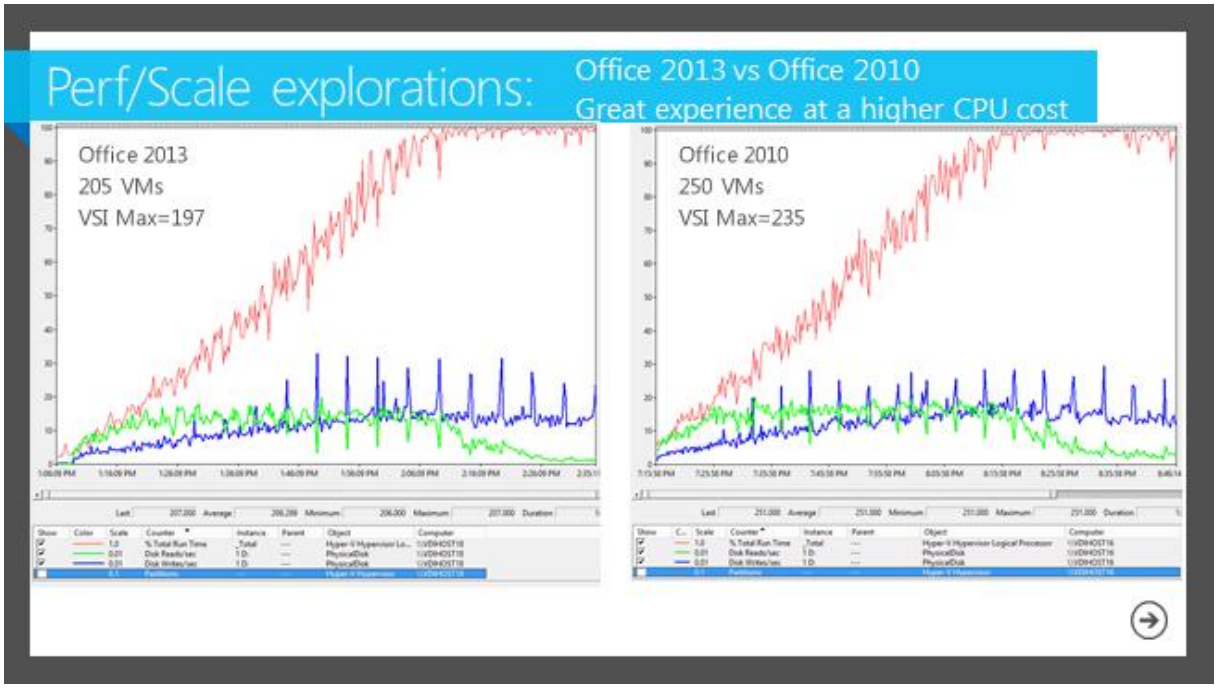

A closer look:

Office 2013:

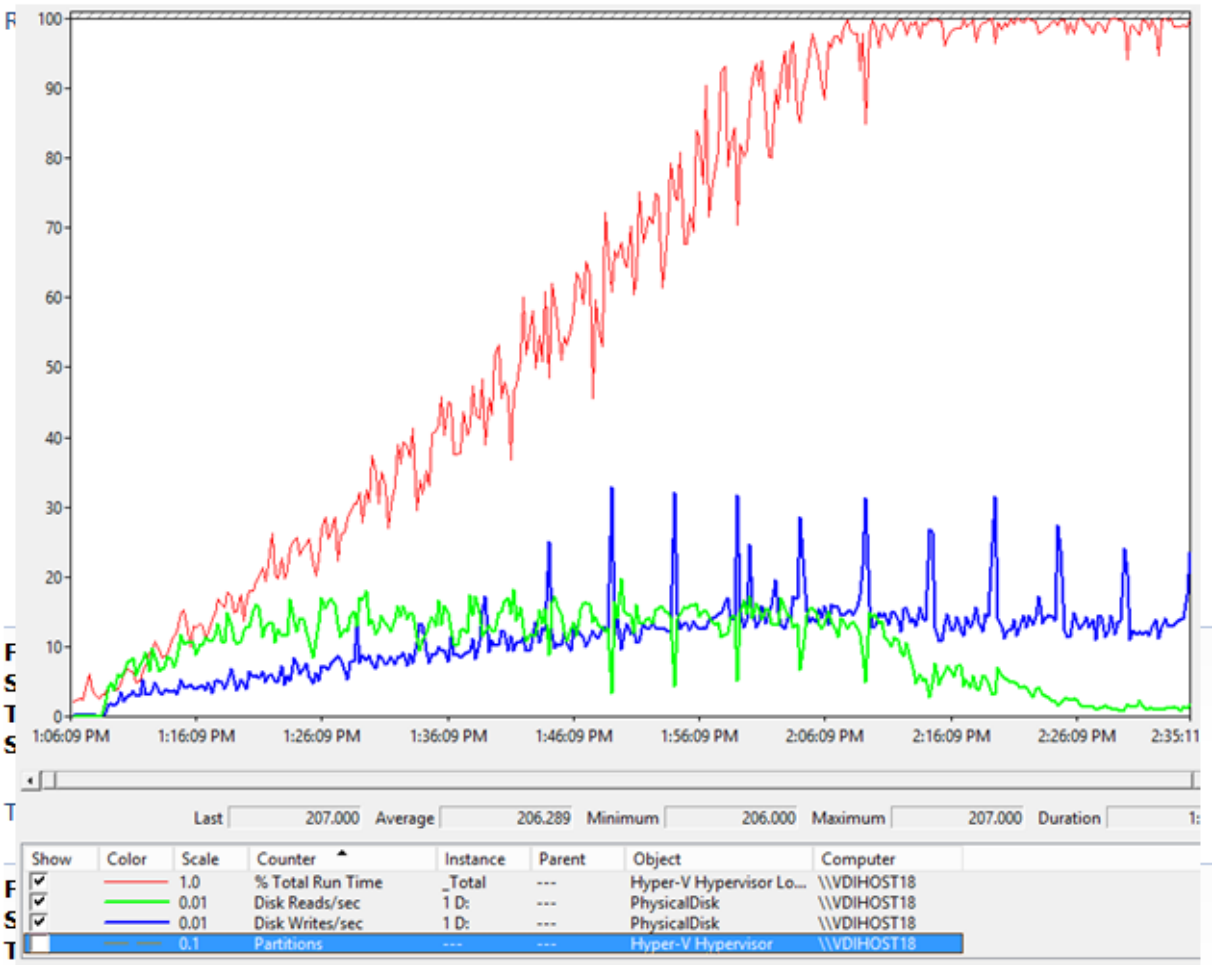
100
90
80
70
60
50
40
30
20
10
0
1:06:09 PM
1:16:09 PM
1:26:09 PM
1:36:09 PM
1:46:09 PM
1:56:09 PM
2:06:09 PM
2:16:09 PM
2:26:09 PM
Last 207.000 Average 206.289 Minimum 206.000 Maximum 207.000 Duration
Show Color Scale Counter Instance Parent Object Computer
1.0 % Total Run Time _Total --- Hyper-V Hypervisor Lo... \\VDIHOST18
0.01 Disk Reads/sec 1 D: --- PhysicalDisk \\VDIHOST18
0.01 Disk Writes/sec 1 D: --- PhysicalDisk \\VDIHOST18
0.1 Partitions --- --- Hyper-V Hypervisor \\VDIHOST18

Office 2010:

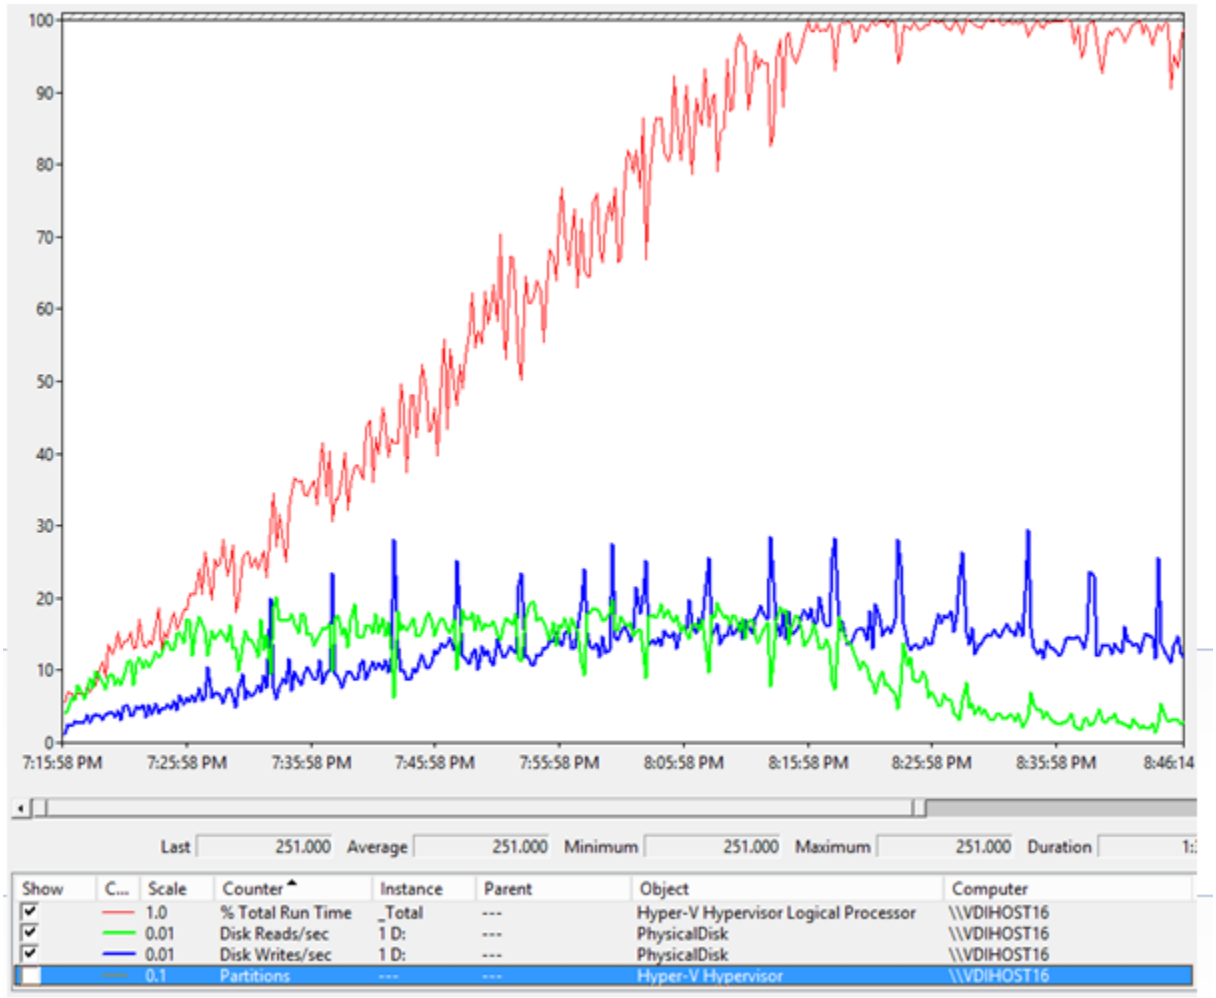
100
90
80
70
60
50
40
30
20
10
0
7:15:58 PM
7:25:58 PM
7:35:58 PM
7:45:58 PM
7:55:58 PM
8:05:58 PM
8:15:58 PM
8:25:58 PM
8:35:58 PM
8:46:14
Last 251.000 Average 251.000 Minimum 251.000 Maximum 251.000 Duration
Show C... Scale Counter Instance Parent Object Computer
1.0 % Total Run Time _Total --- Hyper-V Hypervisor Logical Processor \\VDIHOST16
0.01 Disk Reads/sec 1 D: --- PhysicalDisk \\VDIHOST16
0.01 Disk Writes/sec 1 D: --- PhysicalDisk \\VDIHOST16
0.1 Partitions --- --- Hyper-V Hypervisor \\VDIHOST16

About the Authors

Peter Fine is the Sr. Principal Engineering Architect in the Desktop Virtualization Solutions Group at Dell. Peter has extensive experience and expertise on the broader Microsoft solutions software stack as well as in enterprise virtualization, storage, networking and enterprise datacenter design.

Senthil Baladhandayutham is the Solutions Development Manager in the Desktop Virtualization Solutions Group at Dell, managing the development and delivery of Enterprise class Desktop virtualization solutions based on Dell Datacenter components and core virtualization platforms.

Farzad Khosrowpour is a Systems Sr. Principal Engineering Architect in the Desktop Virtualization Solutions Group at Dell. Farzad is recognized for enterprise level technology development and has led various R&D teams across a broad set of solutions and products including storage, servers, data management, and system architectures.

Jerry Van Blaricom is a Sr. Systems 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.

Shruthin Reddy is a Solutions Engineer in the Desktop Virtualization Solutions Group at Dell with extensive experience validating VDI solutions by Microsoft (Hyper-V), VMware (View), and Citrix (XenDesktop).