Dell EMC VxRail with VMware Horizon

A Reference Architecture document for the design, configuration and implementation of a VxRail Appliance with Horizon.

Dell Engineering
April 2017

A Dell Reference Architecture

Revisions

Date	Description
December 2016	Initial release
February 2017	Updated VxRail description, various diagrams and cache disk information.
April 2017	Test Results Standardization

Table of contents

1 Introduction

1.1 Purpose

This document addresses the architecture design, configuration and implementation considerations for the key components of the architecture required to deliver virtual desktops via VMware Horizon on Dell EMC VxRail with vSphere 6.0 Update 2 on VMware vSAN 6.2.

1.2 Scope

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

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

1.3 What's new

- Introduce Dell EMC VxRail Appliance
- Introduce Hybrid & All-Flash configuration for Dell EMC VxRail
- Introduce VDI optimized Dell EMC VxRail V Series Configurations

2 Solution architecture overview

2.1 Introduction

Dell Wyse Datacenter solutions provide a number of deployment options to meet your desktop virtualization requirements. Our solution is able to provide a compelling desktop experience to a range of employees within your organization from task workers to knowledge workers to power users. The deployment options for Dell Wyse Datacenter include:

- Linked Clones(Non-persistent)
- Full Clone Virtual Desktops (Persistent)
- RDSH with VMware vSAN

2.2 What is Dell EMC VxRail Appliance?

The Dell EMC VxRail appliances are very powerful Hyper Converged Infrastructure Appliances (HCIA) delivered in 1U/2U Rack Building Blocks. The appliances are built on VMware vSAN technology on VMware VSphere and EMC software. VxRail allows the seamless addition of additional nodes to the appliances from the minimum supported three nodes up to 64 nodes.

The Dell EMC VxRail Appliance platforms are equipped with Broadwell processors and you can now start a cluster with 3 nodes at 25% lower entry price to support smaller deployments and this would be ideal for small deployments/POC environments, the recommended starting block is with a four node appliance configuration. VxRail can now support storage-heavy workloads with storage dense nodes, graphics-heavy VDI workloads with GPU hardware, and entry-level nodes for remote and branch office environments finally, you can upgrade from VxRail 3.5 to 4.0 software with a single click via the VxRail manager interface.

VxRail allows customers to start small and scale as there requirements increase. Single-node scaling and low-cost entry-point options give you the freedom to buy just the right amount of storage and compute, whether just beginning a project or adding capacity to support growth. A single node V Series appliance can scale from 16-40 CPU cores, it can have a maximum of 24TB raw with the hyrbid configuration and 46TB Raw with the all-flash. A 64 Node all-flash cluster delivers a maximum of 2,560 cores and 1,840 TB of raw storage.

2.2.1 What is included in Dell EMC VxRail 4.0?

A full suite of capabilities are included with the Dell EMC VxRail 4.0 appliance with no additional cost.

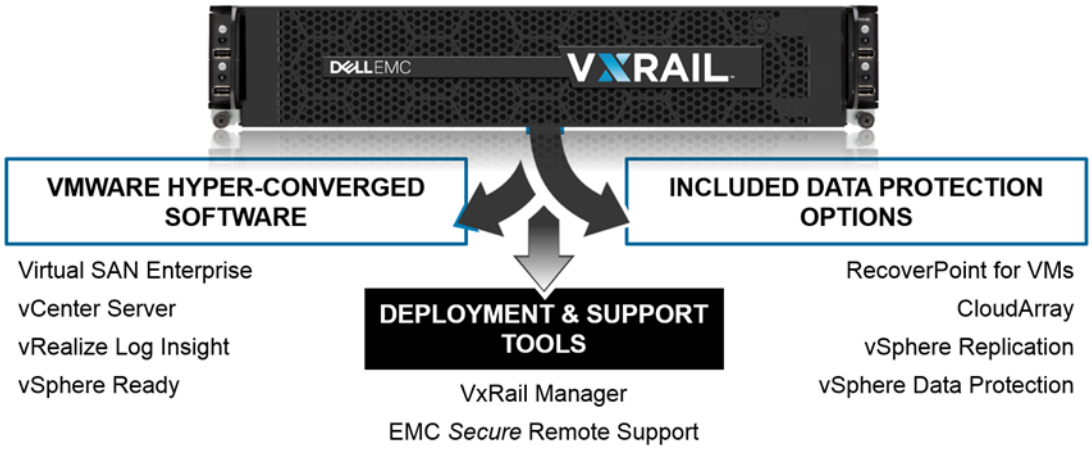

VxRail contains the following software from VMware and EMC.

- vSAN
- vCenter
- ESXi
- vRealize Log Insight
- VxRail Manager

Software License Included with VxRail

- vSAN
- vCenter
- vRealize Log Insight

The customer is prompted during deployment to input an existing vSphere license as although ESXi is installed as part of the factory process a license for ESXi is not included with VxRail.

Optional Software

VxRail also includes optional licensed software that is not pre-installed and configured but the customer is entitled to licenses for this software. They are EMC CloudArray and RecoverPoint.

CloudArray

- A cloud gateway that allows you to expand local storage with using capacity in the cloud.
- A licensed is included with every VxRail appliance purchase - 1TB local / 10TB cloud.
- 1TB acts as hot cache and as it fills, colder data will be moved to the 10TB capacity in the cloud.
- The license does not include the actual cloud storage, but just the ability to manage it. The cloud storage must be purchases separately.
- CloudArray is downloaded and installed from the VxRail Manager Marketplace.

- When CloudArray is used for the first time, customer will be taken to the CloudArray portal and will be prompted to input their PSNT. A license will then be provided to the customer to enable CloudArray.

RecoverPoint

- Data protection for virtual machines.
- A license is included with every VxRail appliance purchase - up to 5 VMs per appliance
- RecoverPoint is downloaded and installed from the VxRail Manager Marketplace.

vSphere Data Protection is also available to be downloaded and installed via the VxRail marketplace. This software is licensed via vSphere and does not come licensed with VxRail.

It is fully integrated with VMware vCenter Server and VMware vSphere Web Client, providing disk-based backup of virtual machines. This software provides full virtual machine restore and file level restore without the need for an agent to be installed in every virtual machine. The patented, variable-length deduplication technology across all backup jobs significantly reduces the amount of backup data disk space needed.

For more information on vSphere Data Protection visit here.

2.3 Physical architecture overview

The core VxRail architecture consists of a Local Tier1 model. This consists of a Cache and a Capacity Tier, the minimum requirements for this configuration is 1 x SSD for the Cache Tier and 1 x HDD/SSD for the Capacity Tier. The management and compute nodes are configured in the same Dell EMC VxRail Cluster and share the VMware vSAN software defined storage. The user data can be hosted on a file server on the VSAN file system.

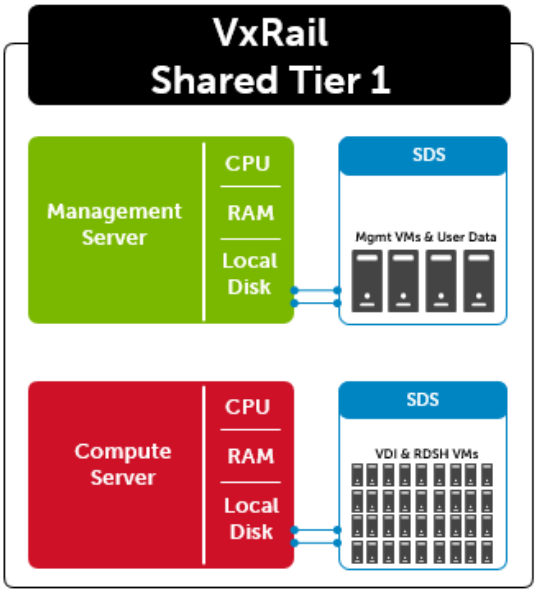

2.4 Solution layers

The Dell EMC VxRail Appliance leverages a core set of hardware and software components consisting of five primary layers:

- Networking Layer
- Compute Server Layer
- Management Server Layer
- Storage Layer (VMware vSAN)
- Thin Client Layer (please refer to section 3.6)

These components have been integrated and tested to provide the optimal balance of high performance and lowest cost per user. The Dell EMC VxRail appliance is designed to be cost effective allowing IT departments to implement high-performance fully virtualized desktop environments.

2.4.1 Networking

Designed for true linear scaling, Dell EMC VxRail series leverages a Leaf-Spine network architecture. A Leaf-Spine architecture consists of two network tiers: an L2 Leaf and an L3 Spine based on 40GbE and non-blocking switches. This architecture maintains consistent performance without any throughput reduction due to a static maximum of three hops from any node in the network.

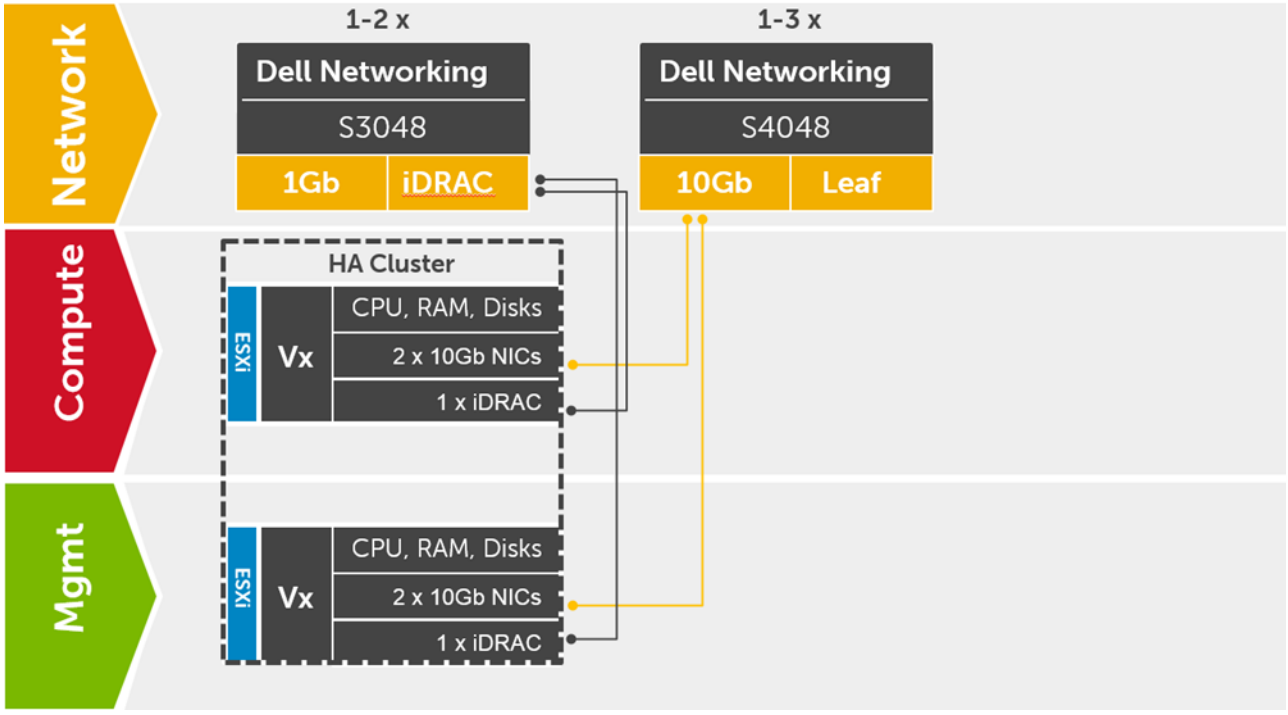

The following figure shows a design of a scale-out Leaf-Spine network architecture that provides 20 GB active throughput from each node to its Leaf and scalable 80 GB active throughput from each Leaf to Spine switch providing scale from 3 VxRail nodes to 64+ without any impact to available bandwidth:

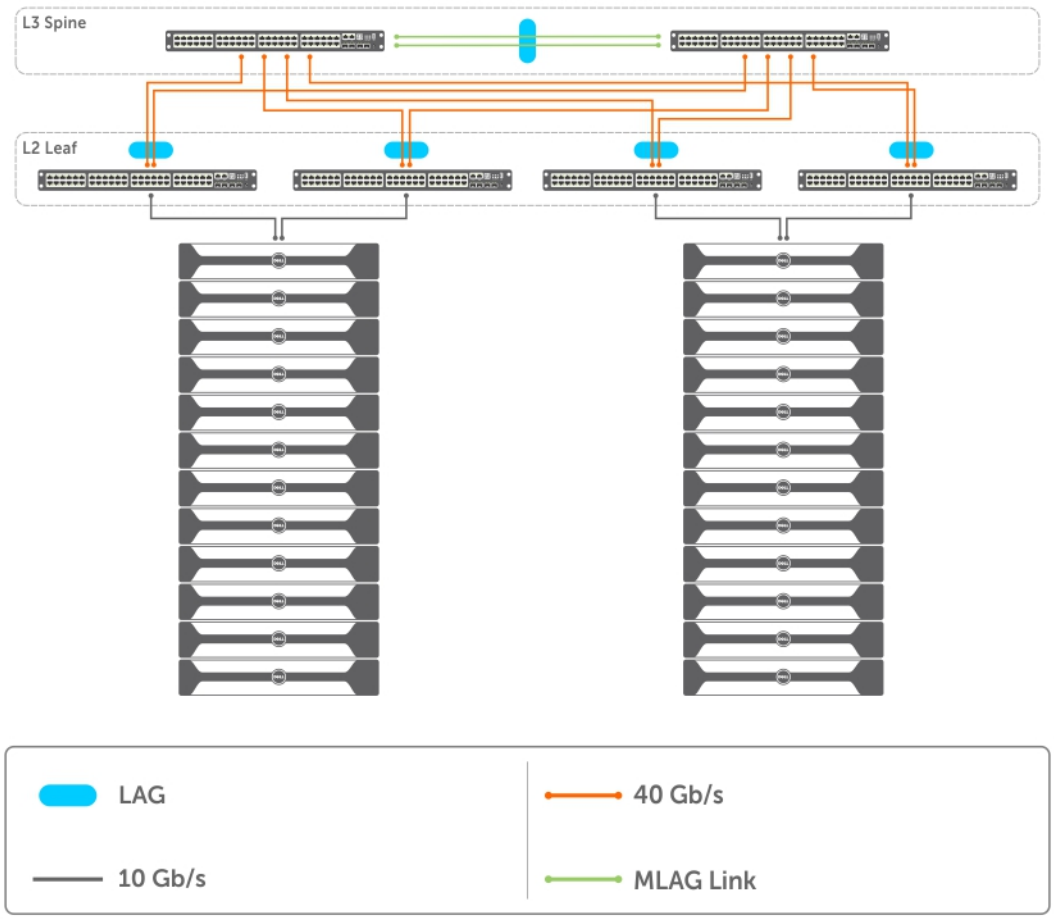

The best practices guide for VxRail 4.0 with S4048-ON is located here.

2.4.2 Dell EMC VxRail Host

The compute, management and storage layers are converged into a single Dell EMC VxRail Series appliance server cluster, hosting VMware vSphere. The recommended boundaries of an individual cluster are based on number of the nodes supported for vSphere 6 which is 64.

Dell recommends that the VDI management infrastructure nodes be separated from the compute resources, in this configuration both management and compute are in the same vSphere HA Cluster. Optionally, the management node can be used for VDI VMs as well with an expected reduction of 30% for these nodes only. The 30% accounts for the amount of resources needed to be reserved for management VMs so this needs to be factored in when sizing. Compute hosts can be used interchangeably for Horizon or RDSH as required.

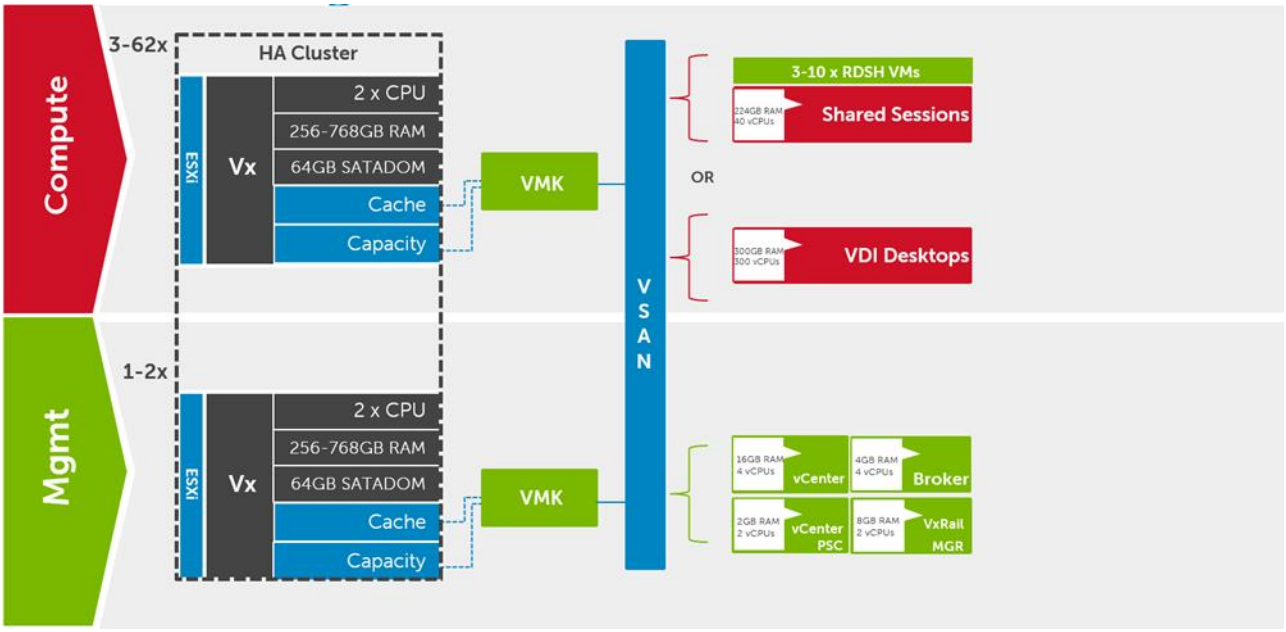

2.4.3 Storage (vSAN)

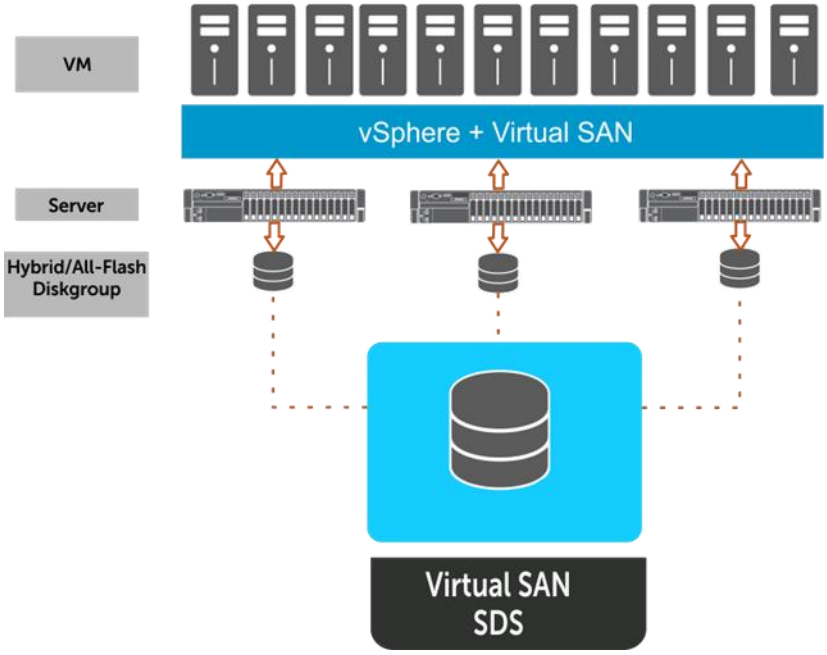

VMware vSAN is software-defined storage solution fully integrated into vSphere. Once enabled on a cluster, all the flash or magnetic hard disks present in the hosts are pooled together to create a shared data store that will be accessible by all hosts in the VMware vSAN cluster. Virtual machines can then be created and a storage policy can be assigned to them. The storage policy will dictate availability / performance and sizing.

From a hardware perspective, at least three ESXi hosts (four recommended) are required for the vSAN cluster. Each host will need at least one SSD and one HDD. In hybrid configurations, the SSD acts as a read cache (70%) and a write buffer (30%). The read cache keeps a list of commonly accessed disk blocks and the write cache behaves as a non-volatile write buffer. It is essential to the performance of the Virtual SAN as all I/O goes to the SSD first. The higher the performance of the disks then the better the performance of your virtual machines. It's important to determine the number of simultaneous write operations that a particular SSD is capable of sustaining in order to achieve adequate performance.

In all-flash configurations, the cache tier is dedicated 100% to writes, allowing all reads to come directly from the capacity tier. This model allows the cache device to protect the endurance of the capacity tier.

All virtual machines deployed to vSAN have an availability policy setting that ensures at least one additional copy of the virtual machine data is available; this includes the write cache contents. When a write is initiated by the VM then it is sent to both the local write cache on the owning host and also to the write cache on the remote hosts. This ensures we have a copy of the in cache data in the event of a host failure and no data will get corrupted. If a block is requested and not found in the read cache, the request is directed to the HDD.

Magnetic hard disk drives (referred to as HDDs from here on) have two roles in vSAN. They make up the capacity of the VMware vSAN data store as well as making up components for a stripe width. SAS and NL-SAS are supported.

VMware recommends configuring 10% of projected consumed capacity of all VMDKs space as SSD storage on the hosts. If a higher ratio is required, then multiple disk groups (up to 4) will have to be created as there is a limit of 1 cache SSD per disk group.

vSAN implements a distributed RAID concept across all hosts in the cluster, so if a host or a component within a host (e.g. an HDD or SSD) fails then virtual machines still have a full complement of data objects available and can continue to run. This availability is defined on a per-VM basis through the use of VM storage policies.

vSAN 6.2 provides two different configuration options, a hybrid configuration that leverages flash-based devices for the cache tier and magnetic disks for the capacity tier, and an all-flash configuration. This delivers enterprise performance and a resilient storage platform. The all-flash configuration uses flash for both the cache tier and capacity tier.

3 Hardware components

3.1 Network

The following sections contain the core network components for the Dell Wyse Datacenter solutions. General uplink cabling guidance to consider in all cases is that TwinAx is very cost effective for short 10 GB runs and for longer runs use fiber with SFPs.

3.1.1 Dell Networking S3048 (1Gb ToR switch)

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

Model	Features	Options	Uses
Dell Networking S3048-ON	48 x 1000BaseT 4 x 10Gb SFP+ Non-blocking, line-rate performance 260Gbps full-duplex bandwidth 131 Mbps forwarding rate	Redundant hot-swap PSUs & fans VRF-lite, Routed VLT, VLT Proxy Gateway User port stacking (up to 6 switches) Open Networking Install Environment (ONIE)	1Gb connectivity (iDRAC)

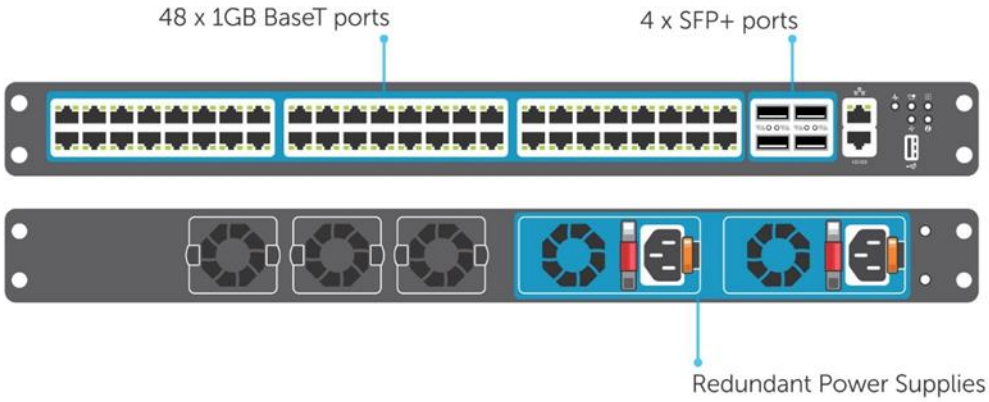

3.1.2 Dell Networking S4048 (10 GB ToR switch)

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

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

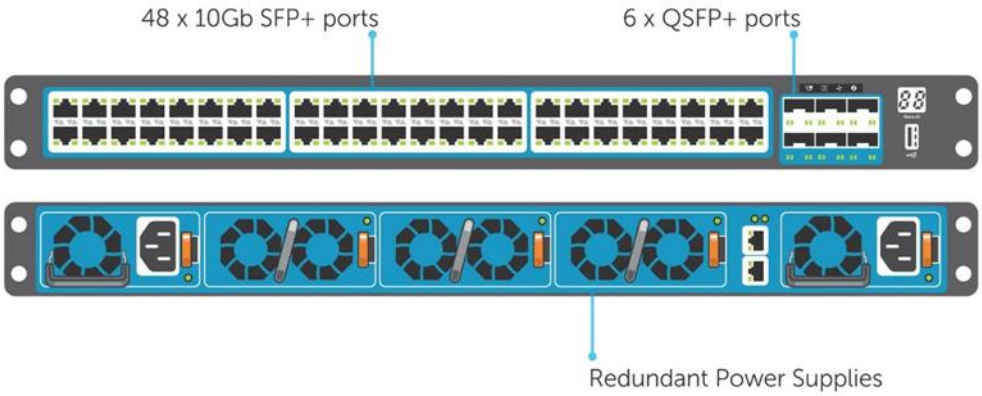

For more information on the S3048, S4048 switches and Dell Networking, please visit this link.

3.2 Dell EMC VxRail Platform Configurations

The Dell EMC VxRail Appliance has multiple platform configuration options. This Reference Architecture focuses primarily on the VDI optimized V Series Platform but this section describes the other optimized platform configurations options that are available also.

Platform	Description	Configurations	Form Factor
G Series	General Purpose	All-Flash & Hybrid	2U4N
E Series	Entry Level	All-Flash & Hybrid	1U1N
V Series	VDI Optimized	All-Flash & Hybrid	2U1N
P Series	Performance Optimized	All-Flash & Hybrid	2U1N
S Series	Storage Dense	Hybrid	2U1N

3.3 Dell EMC VxRail VDI Optimized V Series Configurations

The VDI-optimized 2U/1Node appliance with GPU hardware for graphics-intensive desktop deployments. There is the option to order a V Series configuration without GPU as is details in the A3, B5 & C7 configuration but GPU cards can be added to these configurations at a later date.

In the Local Tier 1 model, VDI sessions execute from local storage on each Compute server. The hypervisor used in this solution is vSphere. In this model, both the Compute and Management server hosts access VMware vSAN storage. The Management, VDI, vMotion and VSAN VLANS are configured across 2 x 10 GB on the NDC.

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

- A5 configuration is perfect for small scale, POC or low density cost-conscience environments. Available in the
- B5 configuration is geared toward larger scale general purpose workloads, balancing performance and cost-effectiveness.
- C7 is the premium configuration offering an abundance of high performance and tiered capacity where user density is maximized.

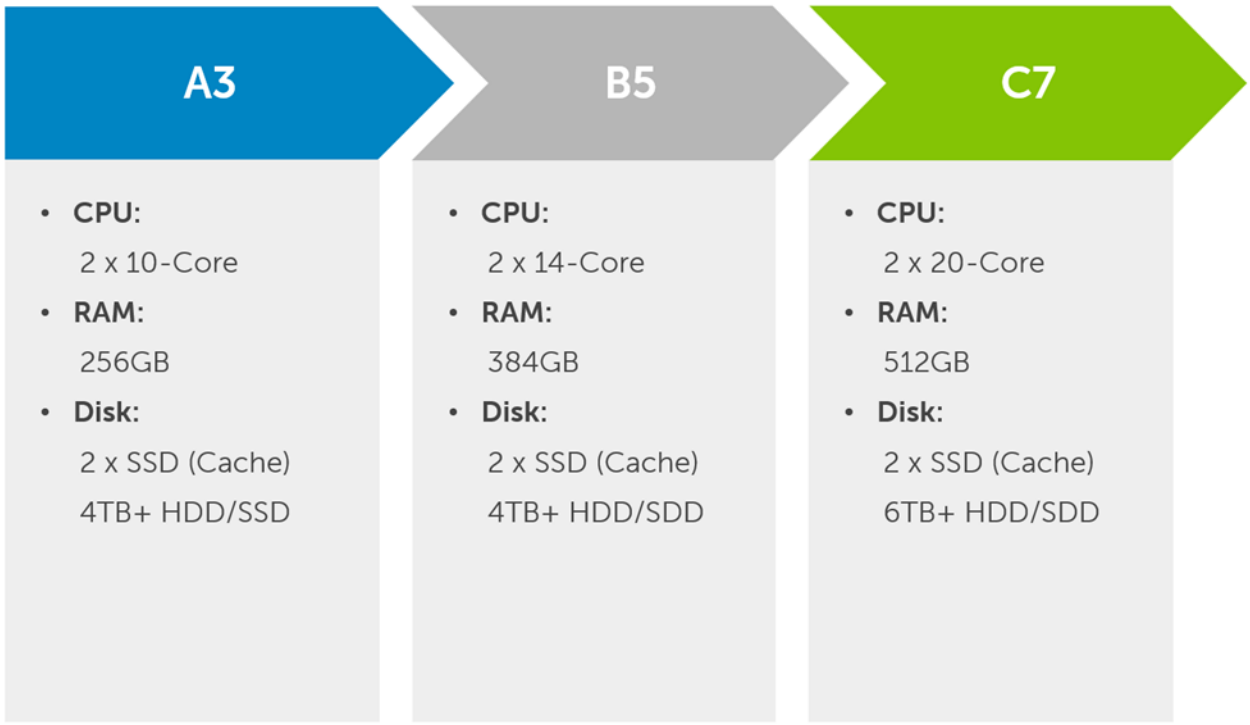

3.3.1 V470/V470F-A3 Configuration

The V470/V470F-A3 configuration is a VDI optimized configuration. The configuration has 256GB of Memory and 2 x E5-2640v4 CPUs with the option of 2 x NVidia M60 GPU cards. The drive configuration consist of two disk groups, 1 cache disk and 2 capacity disks each. The cache disk are populated in slots 0 & 4.

2 x Diskgroups
1 x Cache and 2 x Capacity devices per diskgroup

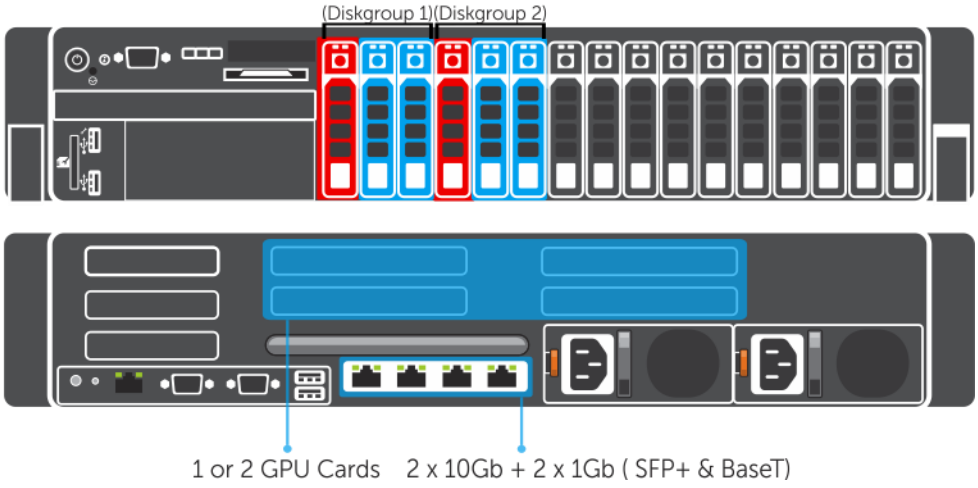

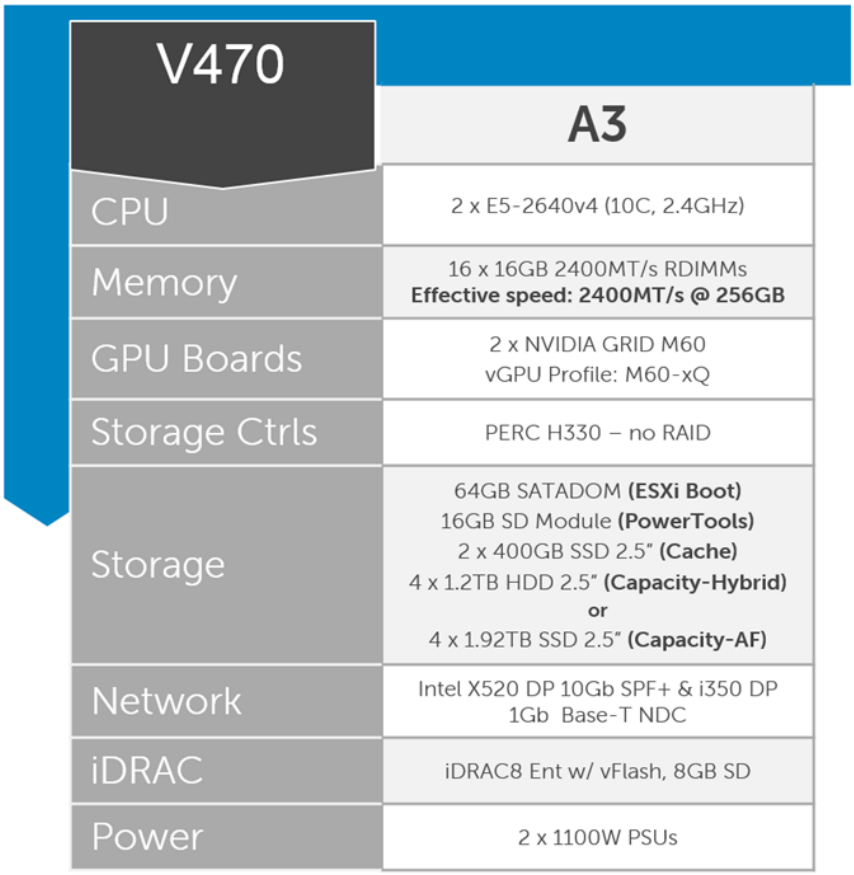

V470	A3
CPU	2 x E5-2640v4 (10C, 2.4GHz)
Memory	16 x 16GB 2400MT/s RDIMMs **Effective speed: 2400MT/s @ 256GB**
GPU Boards	2 x NVIDIA GRID M60 vGPU Profile: M60-xQ
Storage Ctrls	PERC H330 – no RAID
Storage	64GB SATADOM **(ESXi Boot)** 16GB SD Module **(PowerTools)** 2 x 400GB SSD 2.5" **(Cache)** 4 x 1.2TB HDD 2.5" **(Capacity-Hybrid)** **or** 4 x 1.92TB SSD 2.5" **(Capacity-AF)**
Network	Intel X520 DP 10Gb SPF+ & i350 DP 1Gb Base-T NDC
iDRAC	iDRAC8 Ent w/ vFlash, 8GB SD
Power	2 x 1100W PSUs

3.3.2 V470/V470F-B5 Configuration

The V470/V470F-B5 configuration is a VDI optimized configuration. The configuration has 384GB of Memory and 2 x E5-2660v4 CPUs with the option of 2 x NVidia M60 GPU cards. The drive configuration consist of two disk groups, 1 cache disk and 2 capacity disks each.The cache disks are to be populated in slots 0 & 4.

2 x Diskgroups
1 x Cache and 3 x Capacity devices per diskgroup

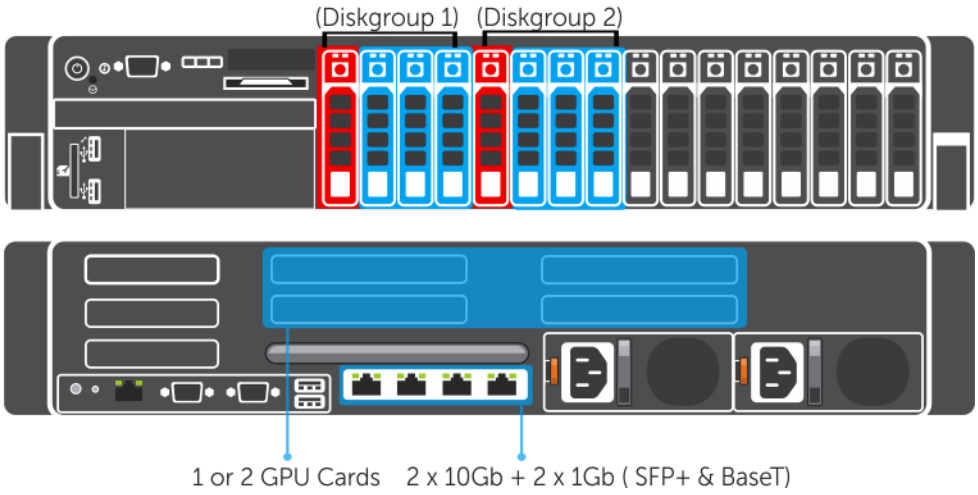

V470	B5
CPU	2 x E5-2660v4 (14C, 2.0GHz)
Memory	24 x 16GB 2400MT/s RDIMMs **Effective speed: 2133MT/s @ 384GB**
GPU Boards	2 x NVIDIA GRID M60 vGPU Profile: M60-xQ
Storage Ctrls	PERC H330 – no RAID
Storage	64GB SATADOM **(ESXi Boot)** 16GB SD Module **(PowerTools)** 2 x 400GB SSD 2.5" **(Cache)** 4 x 1.2TB HDD 2.5" **(Capacity-Hybrid)** **or** 4 x 1.92TB SSD 2.5" **(Capacity-AF)**
Network	Intel X520 DP 10Gb SPF+ & i350 DP 1Gb Base-T NDC
iDRAC	iDRAC8 Ent w/ vFlash, 8GB SD
Power	2 x 1100W PSUs

3.3.3 V470/V470F-C7 Configuration

The V470/V470F-C7 configuration is a VDI optimized configuration. The configuration has 512GB of Memory and 2 x E5-2698v4 CPUs with the option of 2 x NVidia M60 GPU cards. The drive configuration consist of two disk groups, 1 cache disk and 3 capacity disks each. The cache disks are to be populated in slots 0 & 4.

2 x Diskgroups
1 x Cache and 3 x Capacity devices per diskgroup

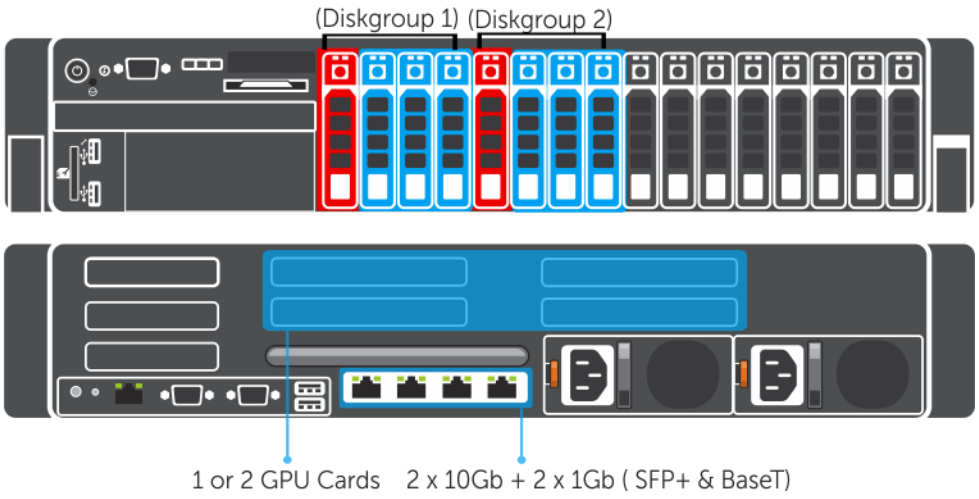

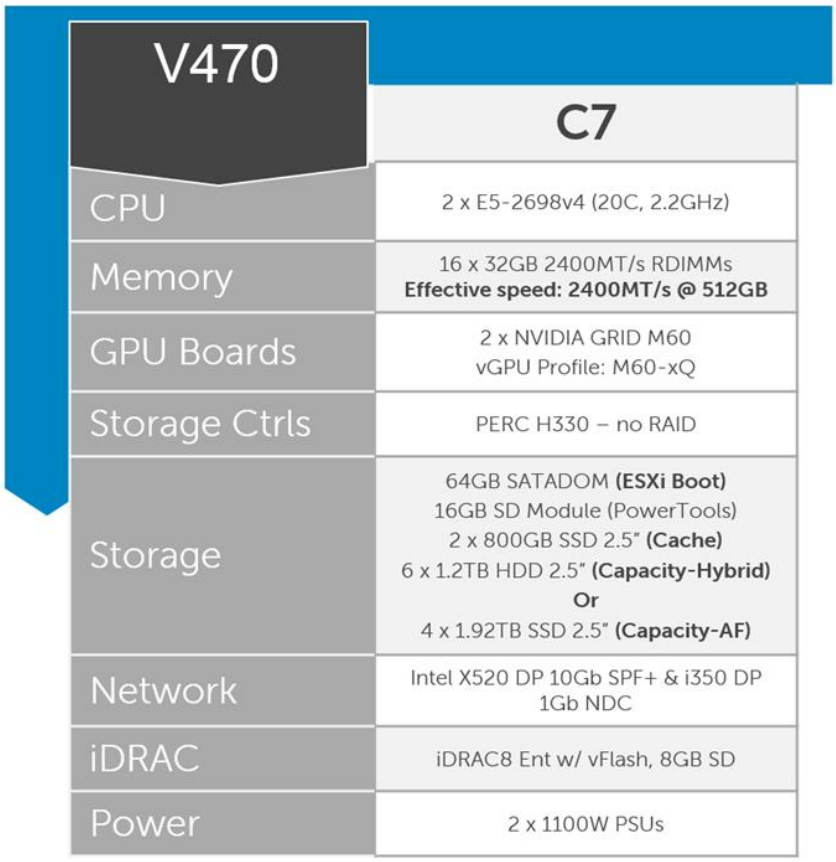

3.4 Dell EMC VxRail Platforms

3.4.1 Dell EMC VxRail E Series Appliance (E460/E460F)

The E Series is the entry level platform, this comes in single or dual socket processor in a 1U configuration per Node. These are aimed for basic workloads, remote office etc. The minimum amount of memory needed for a one CPU configuration is 64GB and the maximum is 768GB. The minimum for a two socket CPU configuration is 128GB and a maximum of 1536GB.The minimum drive configuration is 1 x cache disk and 1 x capacity in a 1 disk group configuration and the maximum for this configuration is 2 x cache disks and 8 capacity in a two disk group configuration. Slot 0 and Slot 5 are to be used for Cache disks only.

2 x Diskgroups
1 x Cache and 4 x Capacity devices per diskgroup

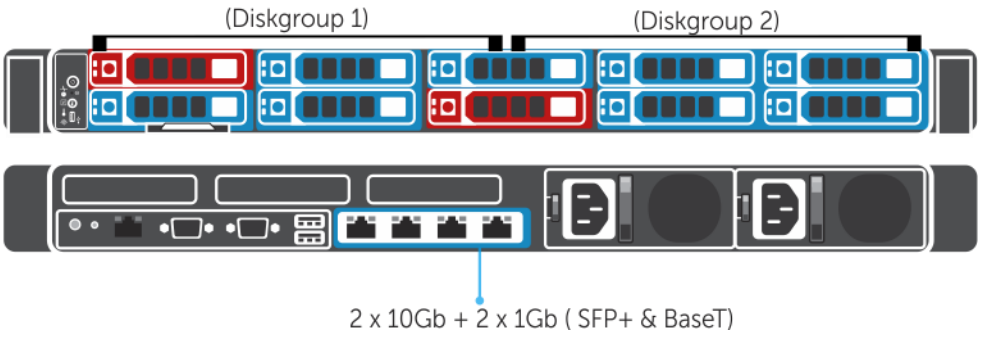

3.4.2 Dell EMC VxRail P Series Appliance (P470/P470F)

The P Series are performance optimized Nodes aimed at high performance scenarios and heavy workloads. There are dual socket processor configuration options with a minimum of 128GB Memory to a maximum of 1536GB. The P470 minimum drive configuration is 1 x cache disk and 1 x capacity in a 1 diskgroup configuration and the maximum for this configuration is 4 x cache disks and 12 capacity in a four diskgroup configuration. The cache disks are be located in Slots 0, 4, 8 and 12 depending on the amount of diskgroups configured.

4 x Diskgroups
1 x Cache and 3 x Capacity devices per diskgroup

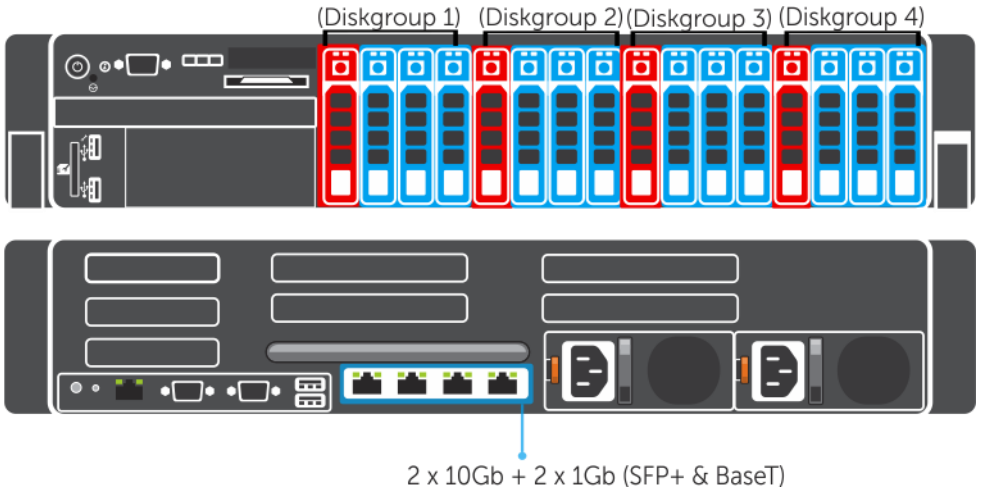

3.4.3 Dell EMC VxRail S Series Appliance (S470)

This is the storage dense platform designed for demanding applications such as virtualized Microsoft SharePoint, Microsoft Exchange, big data, and analytics. This comes in single or dual socket processor configuration the minimum amount of memory needed for a one CPU configuration is 64GB and the maximum is 768GB. The minimum for a two socket CPU configuration is 128GB and a maximum of 1536GB.The minimum drive configuration is 1 x cache disk and 1 x capacity in a 1 disk group configuration and the maximum for this configuration is 2 x cache disks and 12 capacity in a two disk group configuration.

2 x Diskgroups
1 x Cache and 6 x Capacity devices per diskgroup

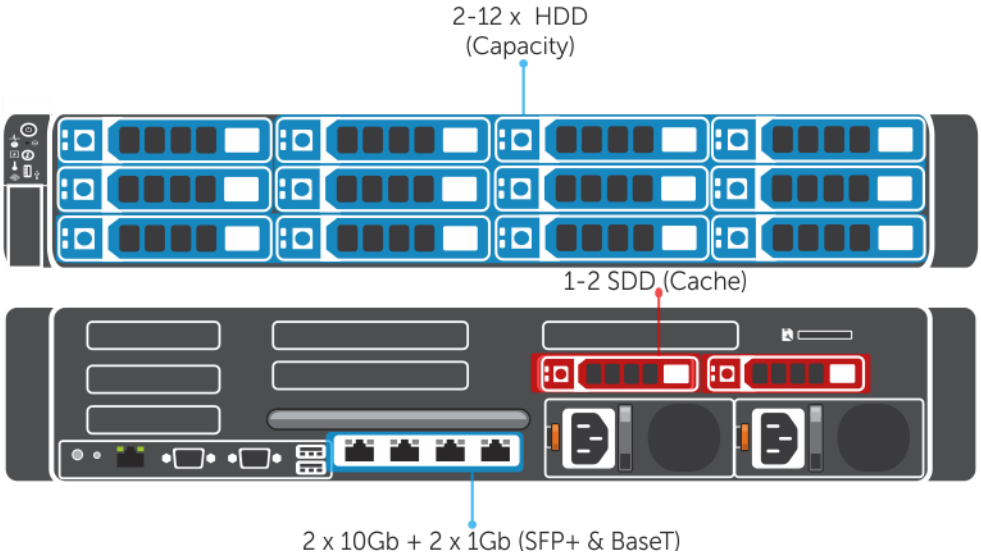

3.5 GPUs

3.5.1 NVIDIA Tesla M60

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

The NVIDIA® Tesla® M60 GPU accelerator works with NVIDIA GRID™ software to provide the industry's highest user performance for virtualized workstations, desktops, and applications. It allows enterprises to virtualize almost any application (including professional graphics applications) and deliver them to any device, anywhere.

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

3.6 Dell Wyse Thin Clients

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

3.6.1 Wyse 3030 LT Thin Client (ThinOS) with PCoIP

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

3.6.2 Wyse 5030 PCoIP Zero Client

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

3.6.3 Wyse 5040 AIO Thin Client with PCoIP

The Dell Wyse 5040 AIO all-in-one (AIO) thin client runs ThinOS with PCoIP, has a 21.5" Full HD display and offers versatile connectivity options for use in a wide range of industries. With four USB 2.0 ports, Gigabit Ethernet and integrated dual band Wi-Fi options, users can link to their peripherals and quickly connect to the network while working with processing-intensive, graphics-rich applications. Built-in speakers, a camera and a microphone make video conferencing and desktop communication simple and easy. It even supports a second attached display for those who need a dual monitor configuration. A simple one-cord design and out-of-box automatic setup makes deployment effortless while remote management from a simple file server, Wyse Device Manager (WDM), or Wyse Thin Client Manager can help

lower your total cost of ownership as you grow from just a few thin clients to tens of thousands. For more information, please visit: link

3.6.4 Wyse 5050 AIO PCoIP Zero Client

The Wyse 5050 All-in-One (AIO) PCoIP zero client has a 23.6" Full HD display and combines the security and performance of the Wyse 5030 PCoIP zero client for VMware with the elegant design of Dell's best-selling P24 LED monitor. The Wyse 5050 AIO provides a best-in-class virtual experience with superior manageability – at a better value than purchasing a zero client and high resolution monitor separately. A dedicated hardware PCoIP engine delivers the highest level of display performance available for advanced applications, including CAD, 3D solids modeling, video editing and more. Elegant in appearance and energy efficient, the Wyse 5050 AIO is a fully functional VMware Horizon endpoint that delivers a true PC-like experience. It offers the full benefits of an efficient and secure centralized computing environment, like rich multimedia, high-resolution 3D graphics, HD media, and full USB peripheral interoperability locally (LAN) or remotely (WAN). For more information, please visit: Link

3.6.5 Wyse 7030 PCoIP Zero Client

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

3.6.6 Wyse 5060 Thin Client (ThinOS) with PCoIP

The Wyse 5060 offers high performance, reliability and flexible OS options, featuring all the security and management benefits of Dell thin clients. Designed for knowledge workers demanding powerful virtual desktop performance, and support for unified communications solutions like Skype for Business, the Wyse 5060 thin client delivers the flexibility, efficiency and security organizations require for their cloud environments. This quad core thin client supports dual 4K (3840x2160) monitors and provides multiple connectivity options with six USB ports, two of which are USB 3.0 for high-speed peripherals, as well as two DisplayPort connectors, wired networking or wireless 802.11 a/b/g/n/ac. The Wyse 5060 can be monitored, maintained, and serviced remotely via Wyse Device Manager (WDM), cloud-based Wyse Cloud Client Manager (CCM) or Microsoft SCCM (5060 with Windows versions). For more information, please visit: Link.

3.6.7 Wyse 7040 Thin Client with Windows Embedded Standard 7P

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

3.6.8 Wyse 7020 Thin Client (Windows 10 IoT)

The versatile Dell Wyse 7020 thin client is a highly efficient and powerful endpoint platform for virtual desktop environments. It is available with Windows Embedded Standard, Windows 10 IoT and Wyse ThinLinux and supports a broad range of fast, flexible connectivity options so that users can connect their favorite peripherals while working with processing-intensive, graphics-rich applications. With a powerful, energy-saving quad core AMD G Series APU in a compact chassis with dual-HD monitor support, the Wyse 7020 thin client delivers stunning performance and display capabilities across 2D, 3D and HD video applications. Its silent diskless and fan less design helps reduce power usage to just a fraction of that used in traditional desktops. Wyse Device Manager (WDM) helps lower the total cost of ownership for large deployments and offers remote enterprise-wide management that scales from just a few to tens of thousands of cloud clients. For more information, please visit Link.

3.6.9 Latitude 3460 mobile thin client

The Latitude 3460 mobile thin client is designed to address a broad range of typical use cases by empowering the mobile workforce to securely access cloud applications and data remotely, while ensuring the security, manageability and centralized control provided by a virtual desktop environment. Optional Advanced Threat Protection in the form of Dell Threat Defense offers proactive malware protection on both virtual desktops and the endpoints. Based on Windows Embedded Standard 7 64-bit for a familiar local Windows experience, this mobile thin client offers high performance with an Intel Celeron 3215U processor, a 14-inch HD (1366 x 768) anti-glare display, a wide range of connectivity options and ports including USB 3.0, HDMI, gigabit Ethernet, and WLAN and Bluetooth options and an extended battery life to enable full productivity in a variety of settings throughout the day. The Latitude 3460 mobile thin client is highly manageable through Wyse Device Manager (WDM), Wyse Cloud Client Manager and Microsoft's System Center.

4 Software components

4.1 VMware

4.1.1 vSphere 6

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

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

Throughout the Dell Wyse Datacenter solution, all VMware and Microsoft best practices and prerequisites for core services are adhered to (NTP, DNS, Active Directory, etc.). The vCenter 6 VM used in the solution is a single Windows Server 2012 R2 VM (Check for current Windows Server OS compatibility at http://www.vmware.com/resources/compatibility) or vCenter 6 virtual appliance, residing on a host in the management Tier. SQL server is a core component of the Windows version of vCenter and is hosted on another VM also residing in the management Tier. It is recommended that Composer is installed on a standalone Windows Server 2012 R2 VM when using the vCenter Server Appliance.

For more information on VMware vSphere, visit http://www.vmware.com/products/vsphere

4.1.2 vSAN

This release of VMware vSAN delivers following important new features and enhancements:

Deduplication and compression: VMware vSAN now supports deduplication and compression to eliminate duplicate data. This technique reduces the total storage space required to meet your needs. When you enable deduplication and compression on a VMware vSAN cluster, redundant copies of data in a particular disk group are reduced to single copy. Deduplication and compression are available as a cluster-wide setting only available as a feature on all-flash clusters.

Enabling deduplication and compression can reduce the amount of storage consumed by as much as 7x. Actual reduction numbers will vary as this depends primarily on the types of data present, number of duplicate blocks, how much these data types can be compressed, and distribution of these unique blocks.

RAID 5 and RAID 6 erasure coding: VMware vSAN now supports both RAID 5 and RAID 6 erasure coding to reduce the storage space required to protect your data. RAID 5 and RAID 6 are available as a policy attribute for VMs in all-flash clusters.

Quality of Service: With the Quality of Service addition to VMware vSAN IOPS limits are now available. Quality of service for VMware vSAN is a Storage Policy Based Management (SPBM) rule. Because quality of

service is applied to VMware vSAN objects through a Storage Policy, it can be applied to individual components or the entire virtual machine without interrupting the operation of the virtual machine.

The term “noisy neighbor” is often used to describe when a workload monopolizes available I/O or other resources, which negatively affect other workloads on the same platform.

For more information on what’s new in VMware vSAN, please visit this link.

VMware vSAN is licensed via the Horizon Advanced or Enterprise license. The Advanced and Enterprise Horizon licenses will cover both the Hybrid and All-Flash configurations of VMware vSAN.

4.1.2.1 vSAN best practices

When determining the amount of capacity required for a VMware vSAN Design we need to pay close attention to the NumberOfFailuresToTolerate(FTT) policy setting. The storage policies that are deployed with Horizon have FTT=1 and that is the recommended default FTT policy setting. When we have FTT=1 set in our policy it will mirror each VMDK in the virtual machine configuration, so if you have two VMDKs that are 40Gb & 20Gb respectively the amount of virtual machine space needed for that virtual machine is 120Gb (40GB x 2 + 20GB x 2).

RAID-5 uses x1.33 the capacity with FTT=1 and requires a minimum of four hosts in the VSAN Cluster. RAID-6 with FTT=2 uses x1.5 the capacity and requires a minimum of six hosts in the VMware vSAN Cluster.

The general recommendation for sizing flash capacity for VMware vSAN is to use 10% of the anticipated storage capacity before the number for FTT is considered.

We also need to factor in how much free capacity or “Slack Space” needs to be preserved when designing the capacity requirement for the VMware vSAN Cluster. The recommendation by VMware is that this should be 30%. The reasoning for this slack space size that the VMware vSAN will begin automatically rebalancing when a disk reaches the 80% full threshold and the additional 10% has been added as a buffer. This is not a hard limit or set via a security policy so the customer can actually use this space but should be made aware of the performance implications of going over the 80% full threshold. More information can be found on the design and sizing of VMware vSAN6.2 Cluster here

4.1.2.2 All-Flash versus Hybrid

The most signification new features in this latest version of VMware vSAN are Deduplication & Compression and erasure coding. These features are only supported in an All-Flash VMware vSAN configuration. The hesitance of a customer going the all flash route is cost but if you factor in the capacity savings achieved by these new features is bridges the gap between the Hybrid & All Flash configurations.

The scenario below is using a VM which consumes 50 GB of space. The hybrid configuration has a default FTT value of 1 and Failure Tolerance Method (FTM) of RAID-1 which has 2x overhead and with FTT=2 that has 3x overhead. The FTM of RAID5/6 is only available with the all-flash configuration and with FTT=1 the overhead is 1.33x, for FTT=2 is 1.5x.

Comparing both FTT=1 scenarios below for both the hybrid and all-flash we can see the capacity savings of over 33GBs per VM so if we had 200VMs per Host that’s a capacity saving of over 660GB of usable VM space per Host.

VM Size	FTM	FTT	Overhead	Configuration	Capacity Required	Hosts Required
50GB	RAID-1	1	2x	Hybrid	100GB	3
50GB	RAID-5	1	1.33x	All-Flash	66.5GB	4
50GB	RAID-1	2	3x	All-Flash	150GB	4
50GB	RAID-6	2	1.5x	All-Flash	75GB	6

Prior to VMware vSAN 6.2, RAID-1 (Mirroring) was used as the failure tolerance method. VMware vSAN 6.2 adds RAID-5/6 (Erasure Coding) to all-flash configurations. While RAID 1(Mirroring) may be favored where performance is the most important factor it is costly with regards to the amount of storage needed.

RAID-5/6 (Erasure Coding) data layout can be configured to help ensure the same levels of availability, while consuming less capacity than RAID-1 (Mirroring). Use of erasure coding reduces capacity consumption by as much as 50% versus mirroring at the same fault tolerance level. This method of fault tolerance does require additional write overhead in comparison to mirroring as a result of data placement and parity.

Deduplication and Compression are two new features that are only available with the all-flash configuration. These features cannot be enabled separately and are implemented at the cluster level. When enabled, VMware vSAN will aim to deduplicate each block and compress the results before destaging the block to the capacity layer. Deduplication and compression work at a disk group level and only objects that are deployed on the same disk group can contribute towards space savings, if components from identical VMs are deployed to different disk groups there will not be any deduplication of identical blocks of data.

The VMware vSAN Read/Write process for both hybrid and all-flash are not the same.

VMware vSAN Hybrid Read: For an object placed on a VMware vSAN datastore, when using RAID-1 configuration it is possible that there are multiple replicas when the number of failure to tolerate are set to greater than 0.Reads may now be spread across the replicas, different reads may be sent to different replicas according to the logical block address and this is to ensure that VMware vSAN does not consume more read cache than is necessary, this avoids caching the data in multiple locations.

VMware vSAN All-Flash Read: Since there is no read cache in an All Flash configuration the process is much different to the Hybrid read operation. The write buffer is first checked to see if the block is present when a read is issued on an all-flash VMware vSAN. This is also the case on hybrid but the difference being with hybrid is that if the block is located in the write buffer it will not be fetched from here. If the requested block is not in the write buffer it will be fetched from the capacity tier but since the capacity tier is also SSD the latency overhead in the first checking the cache and then the capacity tier is minimal. This is main reason why there isn't a read cache with all-flash, the cache tier is a dedicated write buffer which in turns frees up the cache tier for more writes boosting overall IOPS performance.

VMware vSAN Hybrid Write: When a VM is deployed on a hybrid cluster the components of the VM are spread across multiple hosts so when an application within that VM issues a write operation, the owner of the object clones the write operation. This means that the write is sent to the write cache on Host 1 and Host 2 in parallel.

VMware vSAN All-Flash Write: The write process on all-flash is similar to the write process on hybrid, the major difference between both is that with all-flash 100% of the cache tier is assigned to the write buffer whereas with hybrid only 30% is assigned to the write buffer, and the other 70% is assigned to the read cache.

4.1.2.3 VM storage policies for VMware vSAN

Storage policy plays a major role for VMware vSAN strategy and performances. After data store creation you can create VM storage policies to meet VM availability, sizing and performance requirements. The policies are applied down to the VMware vSAN layer when a VM is created. The VM virtual disk is distributed across the VMware vSAN datastore per policy definition to meet the requirements.

VMware Horizon 7 has a built in storage policy for VMware vSAN. When creating a desktop pool with Horizon, select Use VMware vSAN option for Storage Policy Management.

When this is selected a set of storage policies are deployed and visible from with the vSphere Web Console (monitoring/VM Storage Policies).

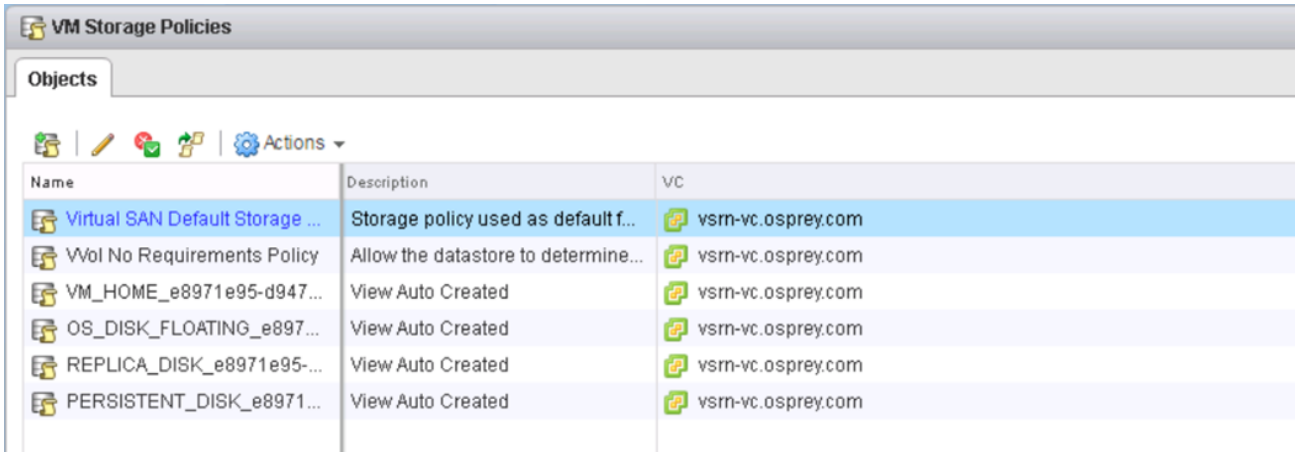

Each policy can be edited but it is recommended to refer to design and sizing guide for VMware vSAN 6.2 located here before making any change to the policy.

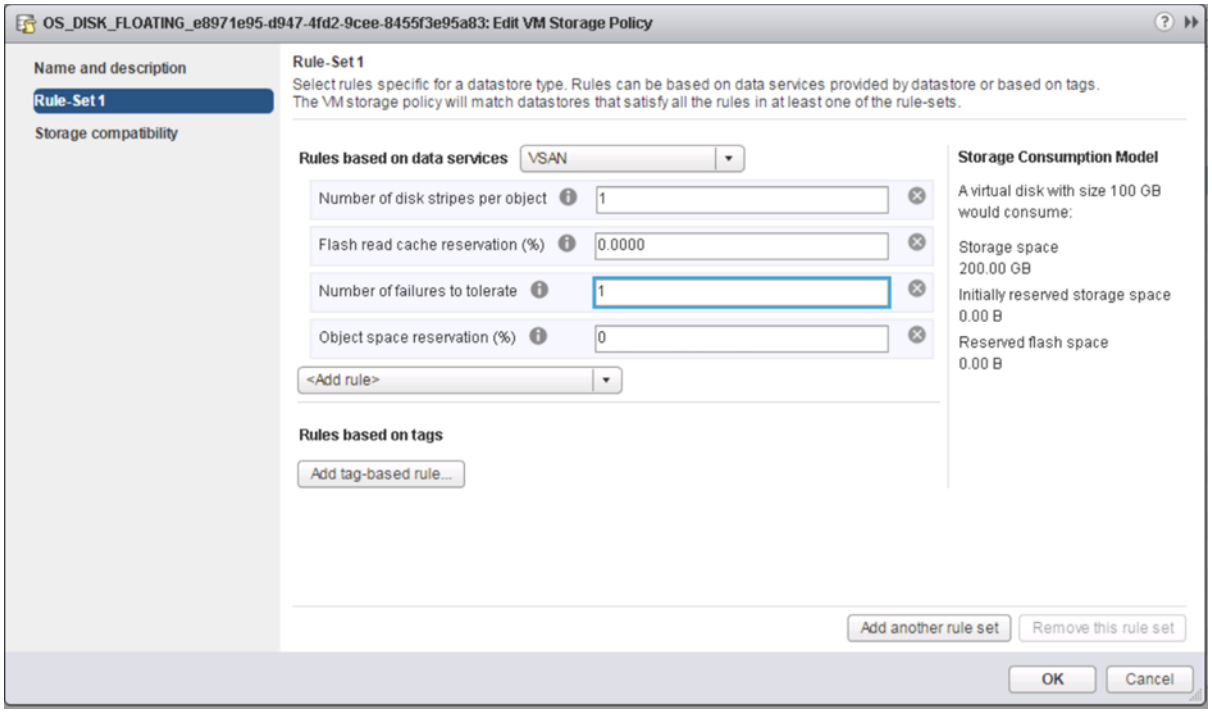

4.1.3 Horizon

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

VMware Horizon provides a complete virtual desktop delivery system by integrating several distributed components with advanced configuration tools that simplify the creation and real-time management of the virtual desktop infrastructure.

The Horizon License matrix can be found here . The Horizon Enterprise license will cover Just in time desktops and App Volumes whereas these new features are not covered under the Standard and Advanced Horizon licenses.

The core Horizon components include:

Horizon Connection Server (HCS) – Installed on servers in the data center and brokers client connections, The VCS authenticates users, entitles users by mapping them to desktops and/or pools, establishes secure connections from clients to desktops, support single sign-on, sets and applies policies, acts as a DMZ security server for outside corporate firewall connections and more.

Horizon Client – Installed on endpoints. Is software for creating connections to Horizon desktops that can be run from tablets, Windows, Linux, or Mac PCs or laptops, thin clients and other devices.

Horizon Portal – A web portal to access links for downloading full Horizon clients. With HTML Access Feature enabled enablement for running a Horizon desktop inside a supported browser is enabled.

Horizon Agent – Installed on all VMs, physical machines and Terminal Service servers that are used as a source for Horizon desktops. On VMs the agent is used to communicate with the Horizon client to provide services such as USB redirection, printer support and more.

Horizon Administrator – A web portal that provides admin functions such as deploy and management of Horizon desktops and pools, set and control user authentication and more.

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

vCenter Server – This is a server that provides centralized management and configuration to entire virtual desktop and host infrastructure. It facilitates configuration, provision, management services. It is installed on a Windows Server 2008 host (can be a VM).

Horizon Transfer Server – Manages data transfers between the data center and the Horizon desktops that are checked out on the end users' desktops in offline mode. This Server is required to support desktops that run the Horizon client with Local Mode options. Replications and syncing are the functions it will perform with offline images.

4.1.3.1 Horizon Key Features

This current release of VMware Horizon delivers following important new features and enhancements:

4.1.3.2 Just in time delivery with Instant Clone Technology

Reduce infrastructure requirements while enhancing security with Instant Clone technology and App Volumes. Instantly deliver brand new personalized desktop and application services to end users every time they log in. Just in Time Delivery with Instant Clone Technology is turning the traditional VDI provisioning model on its head.

The booted-up parent VM can be “hot-cloned” to produce derivative desktop VMs rapidly, leveraging the same disk and memory of the parent, with the clone starting in an already “booted-up” state. This process bypasses the cycle time incurred with traditional cloning where several power cycle and reconfiguration calls are usually made.

When Instant Clone technology is used in conjunction with VMware App Volumes and User Environment Manager, administrators can use Instant Clone Technology to rapidly spin up desktops for users that retain user customization and persona from session to session, even though the desktop itself is destroyed when the user logs out. Virtual desktops benefit from the latest O/S and application patches automatically applied between user logins, without any disruptive recompose.

4.1.3.3 Transformational user experience with Blast Extreme

A new VMware controlled protocol for a richer app & desktop experience Protocol optimized for mobile and overall lower client TCO. All existing Horizon remote experience features work with Blast Extreme and updated Horizon clients. Deliver rich multimedia experience in lower bandwidth Rapid client proliferation from strong Horizon Client ecosystem.

Blast Extreme is network-friendly, leverages both TCP and UDP transports, powered by H.264 to get the best performance across more devices, and reduces CPU consumption resulting in less device power consumed for longer battery life.

4.1.3.4 Modernize application lifecycle management with App Volumes

Transform application management from a slow, cumbersome process into a highly scalable, nimble delivery mechanism that provides faster application delivery and application management while reducing IT costs by up to 70%.

VMware App Volumes is a transformative solution that delivers applications to Horizon virtual desktops. Applications installed on multi-user AppStacks or user-specific writable volumes attach instantly to a desktop at user login. The App Volumes user experience closely resembles that of applications natively installed on the desktop with App Volumes, applications become VM-independent objects that can be moved easily across data centers or to the cloud and shared with thousands of virtual machines.

4.1.3.5 Smart policies with streamlined access

Improve end user satisfaction by simplifying authentication across all desktop and app services while improving security with smarter, contextual, role-based policies tied to a user, device or location.

Policy-Managed Client Features, which enables IT to use policy to define which specific security-impacting features, are accessible upon login. These include clipboard redirection, USB, printing, and client-drives. All of these can be enforced contextually, based on role, evaluated at logon/logoff, disconnect/reconnect and at pre-determined refresh intervals for consistent application of policy across the entirety of the user experience. For example, a user logging in from a network location consider unsecured, can be denied access to USB and printing. Additionally, PCoIP bandwidth profile settings allow IT to customize the user experience based on user context and location.

True SSO streamlines secure access to a Horizon desktop when users authenticate via VMware Identity Manager. A short lived VMware Horizon virtual certificate is generated, enabling a password-free Windows login, bypassing the usual secondary login prompt users would encounter before getting to their desktop.

4.2 Microsoft RDSH

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

- One or more Windows Server OS instances added to the Horizon site

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

4.2.1 NUMA Architecture Considerations

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

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

4.2.2 V470/V470F A3 NUMA Alignment

10 physical cores per CPU in the V470/V470F A3 configuration, 20 logical with Hyper-threading active, gives us a total of 20 consumable cores per appliance.

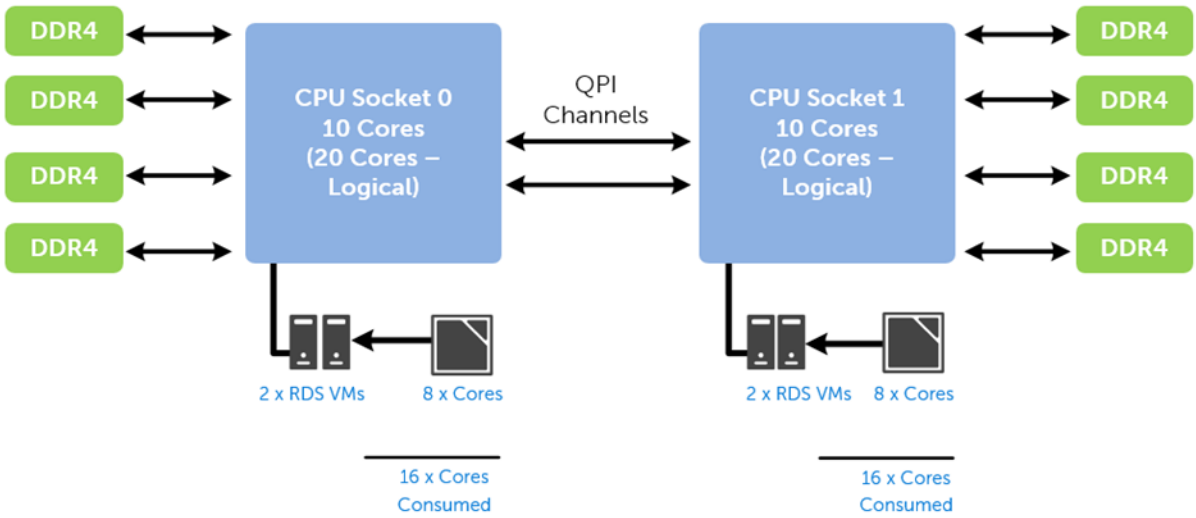

4.2.3 V470/V470F B5 NUMA Alignment

14 physical cores per CPU in the V470/V470F B5 configuration, 28 logical with Hyper-threading active, gives us a total of 56 consumable cores per appliance.

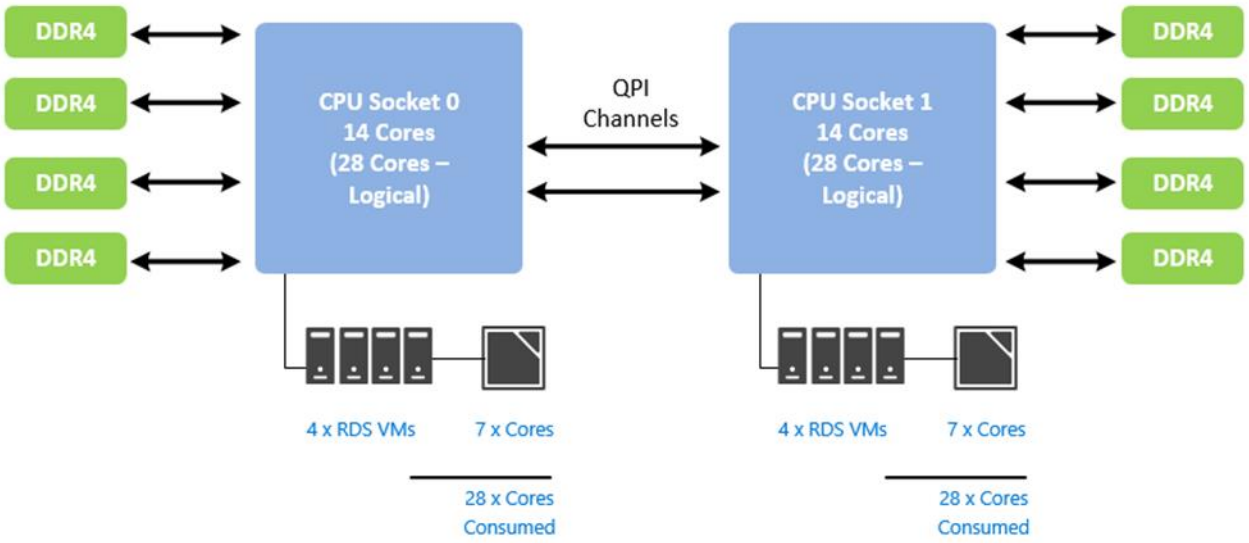

4.2.4 V470/V470F C7 NUMA Alignment

20 physical cores per CPU in the V470/V470F C7 configuration, 40 logical with Hyper-threading active, gives us a total of 80 consumable cores per appliance.

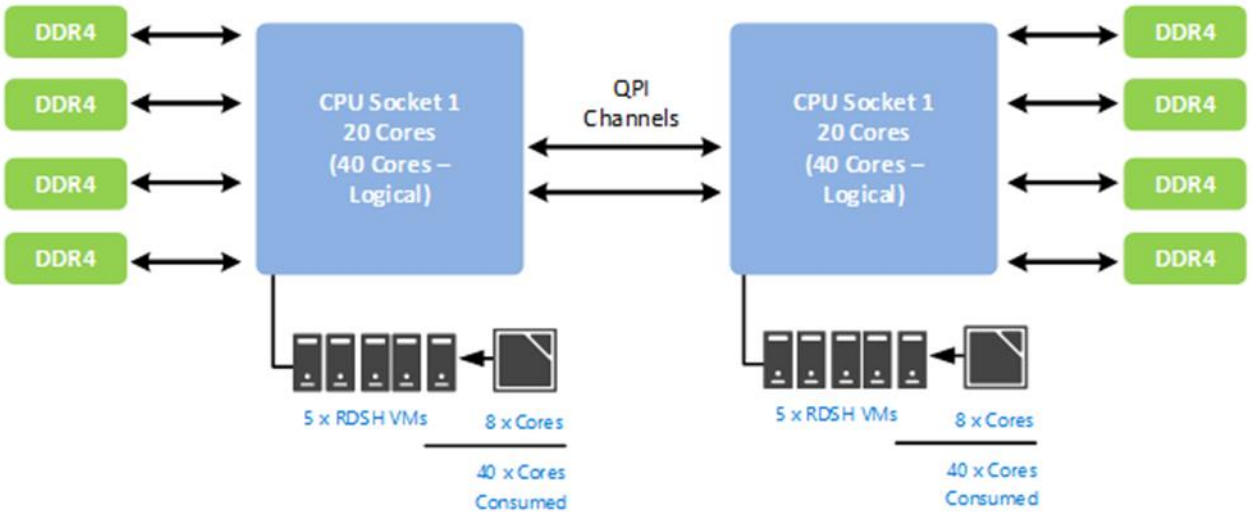

4.3 NVIDIA GRID vGPU

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

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

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

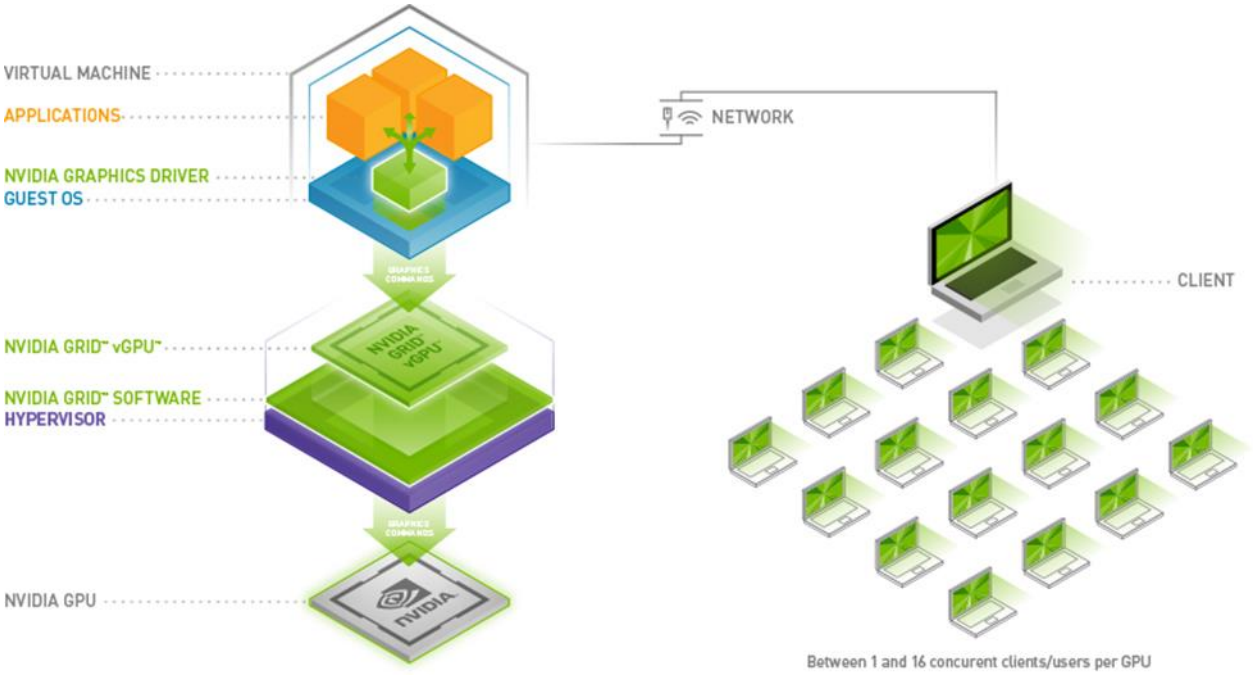

Image provided courtesy of NVIDIA Corporation, Copyright NVIDIA Corporation

4.4 vGPU Profiles

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

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

For more information about NVIDIA GRID vGPU, please visit this link.

NVIDIA Tesla M60 GRID vGPU Profiles:

Card	vGPU Profile	Graphics Memory (Frame Buffer)	Virtual Display Heads	Maximum Resolution	Maximum Graphics-Enabled		
					Per GPU	Per Card	Per Server (2 cards)
Tesla M60	M60-8Q	8GB	4	4096x2160	1	2	4
	M60-4Q	4GB	4	4096x2160	2	4	8
	M60-2Q	2GB	4	4096x2160	4	8	16
	M60-1Q	1GB	2	4096x2160	8	16	32
	M60-0Q	512MB	2	2560x1600	16	32	64
	M60-1B	1GB	4	2560x1600	8	16	32
	M60-0B	512MB	2	2560x1600	16	32	64
	M60-8A	8GB	1	1280x1024	1	2	4
	M60-4A	4GB			2	4	8
	M60-2A	2GB			4	8	16
	M60-1A	1GB			8	16	32

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

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

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

4.4.1 GRID vGPU Licensing and Architecture

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

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

*0Q profiles only support up to 2560x1600 resolution

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

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

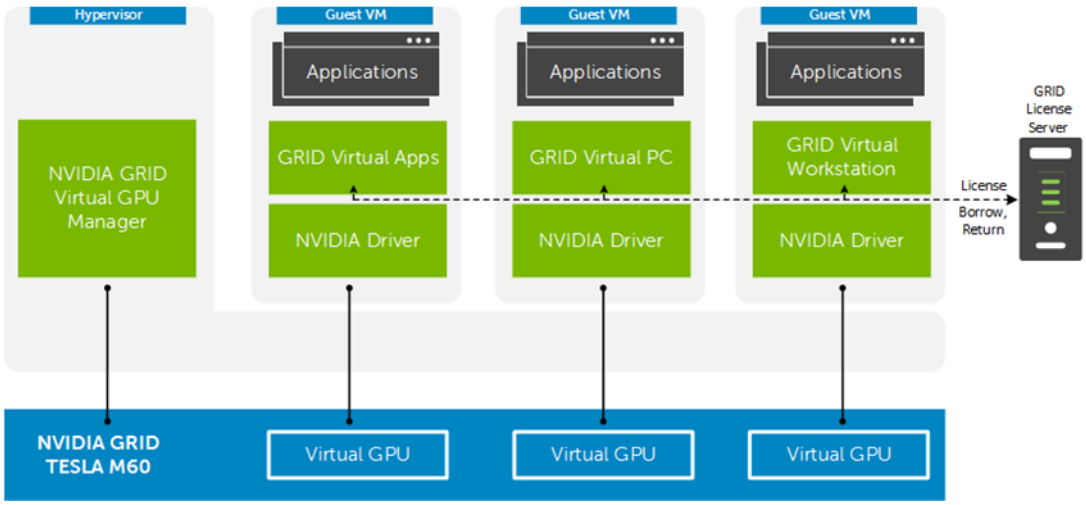

5 Solution architecture for Dell EMC VxRail with Horizon

5.1 Management server infrastructure

There is the option to use and existing Virtual Center during the VxRail deployment but the sizing information below shows the details of the VC appliance and PSC that will be deployed during the factory install.

Role	vCPU	RAM (GB)	NIC	OS + Data vDisk (GB)	Tier 2 Volume (GB)
VMware vCenter Appliance	2	8	1	150	
Horizon Connection Server	2	8	1	40	-
Platform Services Controller	2	2	1	30	
SQL Server	5	8	1	40	210 (VMDK)
File Server	1	4	1	40	2048 (VMDK)
VxRail Manager	2	8	1	32	
Log Insight	4	8	1	530	
Total	**16 vCPU**	**46GB**	**7 vNICs**	**862GB**	**2258GB**

5.1.1 SQL databases

The VMware databases will be hosted by a single dedicated SQL 2012 SP1 Server VM (check DB compatibility at Link in the Management layer. Use caution during database setup to ensure that SQL data, logs and TempDB are properly separated onto their respective volumes. Create all Databases that will be required for:

- Events
- Composer

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

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

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

5.1.2 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 VMware software components. All hosts, VMs and consumable software components need to have a presence in DNS, preferably via a dynamic and AD-integrated namespace. Microsoft best practices and organizational requirements are to be adhered to.

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

5.1.2.1 DNS for SQL

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

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

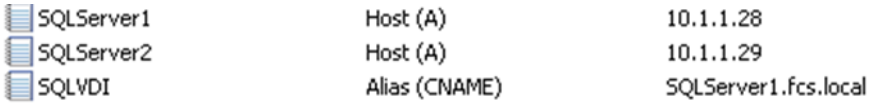

5.2 Storage architecture overview

All Dell EMC VxRail appliances come with two tiers of local storage by default, SSD for performance and SSD or HDD for capacity depending on if it’s a Hybrid or All-Flash configuration. These disk groups need a minimum of 1 x cache device and 1 x capacity device per disk group. These local storage disk groups are configured into one Software Defined Storage pool via VSAN which are shared across all hosts in the VSAN Cluster.

5.2.1 VMware vSAN local storage

VMware vSAN is enabled and configured during the VxRail deployment so there is no manual configuration of vSAN needed with VxRail.

5.3 Virtual Networking

5.3.1 Dell EMC VxRail network configuration

The network configuration for the Dell EMC VxRail appliances utilizes a 10 GB converged infrastructure model. All required VLANs will traverse 2 x 10Gb NICs configured in an active/ active team. For larger scaling it is recommended to separate the infrastructure management VMs from the compute VMs to aid in

predictable compute host scaling. The following outlines the VLAN requirements for the Compute and Management hosts in this solution model:

- VxRail VLAN configuration
  - Management VLAN: Configured for hypervisor infrastructure traffic – L3 routed via core switch
  - VDI VLAN: Configured for VDI session traffic – L3 routed via core switch
  - VMware vSAN VLAN: Configured for VMware vSAN traffic – L2 switched only via ToR switch
  - vMotion VLAN: Configured for Live Migration traffic – L2 switched only, trunked from Core (HA only)
  - VDI Management VLAN: Configured for VDI infrastructure traffic – L3 routed via core switch
- A VLAN for iDRAC is configured for all hardware management traffic – L3 routed via core switch

The following screenshot shows the VMkernel adapter for the management network (vmk0), vMotion and VMware vSAN Network (vmk2) on a distributed switch.

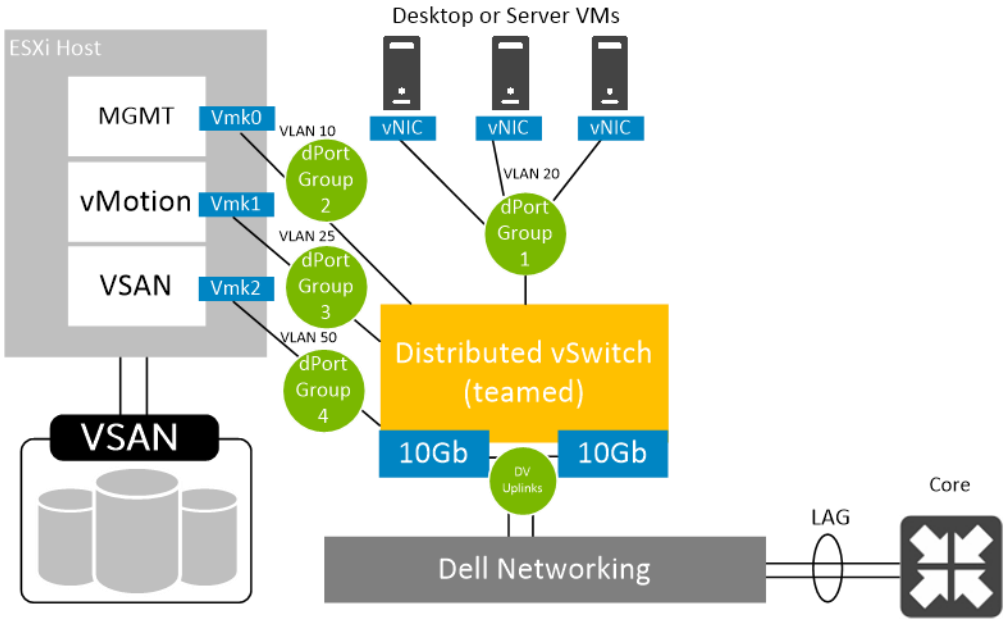

5.3.1.1 vSphere Distributed Switches

The benefit of using a VMware Distributed Switch (vDS) is that it brings a consistent configuration across all hosts. The vDS is configured at the vCenter level and provides central management and monitoring to all hosts configured on the vDS.

dvSwitches should be used as desired for VM traffic especially in larger deployments to ease the management burden across numerous hosts. In the VxRail rack model both the mgmt. hosts connect to shared storage so require additional VMK ports. Network share values should be configured equally among the VMKernel port groups that share a physical set of network adapters.

VMware vSAN cluster networking includes at least two VMkernel ports, one for management traffic and one for VMware vSAN traffic. If vMotion, Storage vMotion or High Availability functionality is required in addition, a third VMkernel port is to be configured for this.

VMware vSAN traffic can be used on 1 GB networks as well as 10 GB networks for Hybrid configuration but 10 GB recommended and is required for All Flash configuration. Standard switch configuration can be used for Proof of Concept, while VMware distributed virtual switch configuration is highly recommended for production versions.

Network VMkernel adapter configuration for the host management traffic using a 10 GB network with standard switch. It is recommended that the network configuration for the VMware vSAN storage is a 10 GB network with distributed switch configuration.

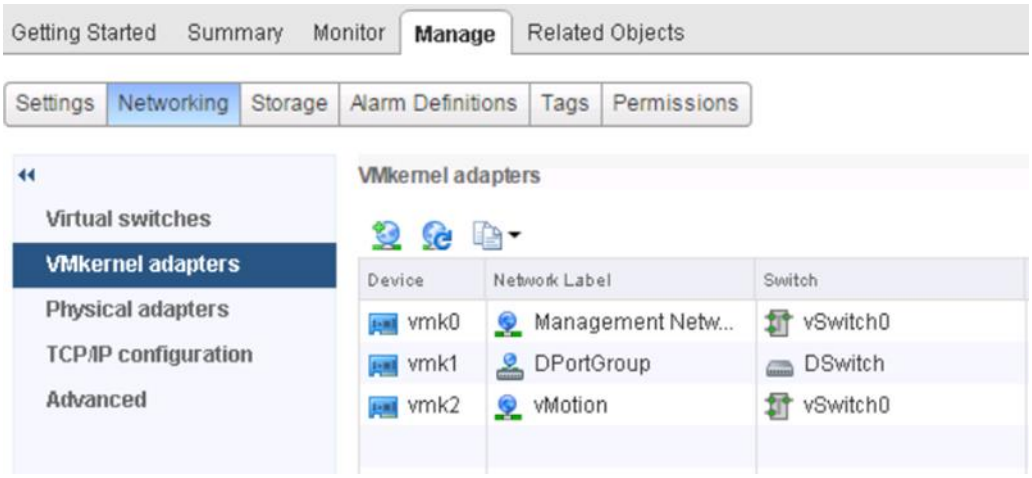

The distributed switch configuration is the same on all VxRail storage hosts. It is recommended to have at least two uplinks for each host to provide load balancing and fail back redundancy. The below image shows an example of a distributed switch configuration for VxRail.

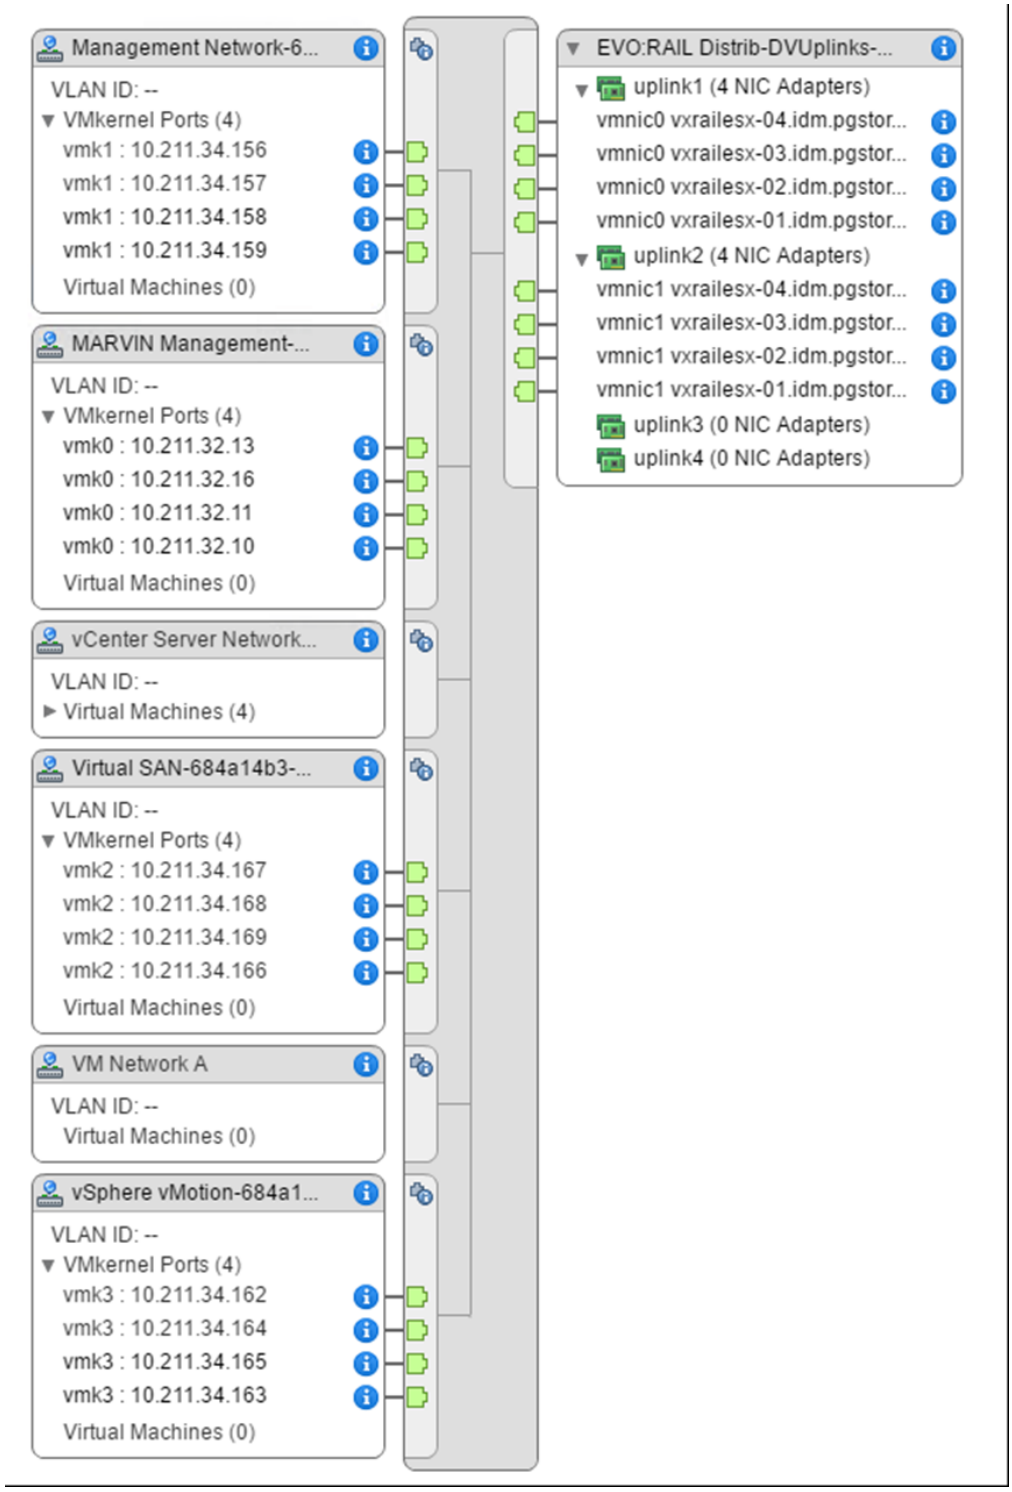

5.3.2 VMware NSX

Dell and VMware's Software Defined Datacenter (SDDC) architecture goes beyond simply virtualizing servers and storage but also extends into the network. VMware NSX is a network virtualization platform deployable on any IP network that is integrated with vSphere Virtual Distributed Switching and provides the same features and benefits to networking as the ESXi hypervisor does to virtual machines. NSX provides a complete set of logical networking elements and services—including logical switching, routing, firewalling, load balancing, VPN, quality of service (QoS), and monitoring. These services are provisioned in virtual networks through any cloud management platform leveraging the NSX APIs. Through Dell's open networking, companies are best able to take advantage of this disaggregation of a virtual network overlay and an open physical underlay. Building a zero-trust security model is easy with NSX as each virtualized workload can be protected with a stateful firewall engine providing extreme policy granularity. Any VM in the datacenter can be rigorously secured or isolated if compromised, especially useful for virtual desktops to prevent malicious code from attacking and spreading through the network.

VMware NSX is implemented via a layered architecture consisting of data, control and management planes. The NSX vSwitch exists within and requires the vSphere Distributed Switch to abstract the physical network while proving access-level switching in the hypervisor. NSX enables the use of virtual load balancers, firewalls, logical switches and routers that can be implemented and scaled seamlessly to suit any deployed architecture. VMware NSX compliments Dell Networking components deployed ToR, leaf/spine or at the core.

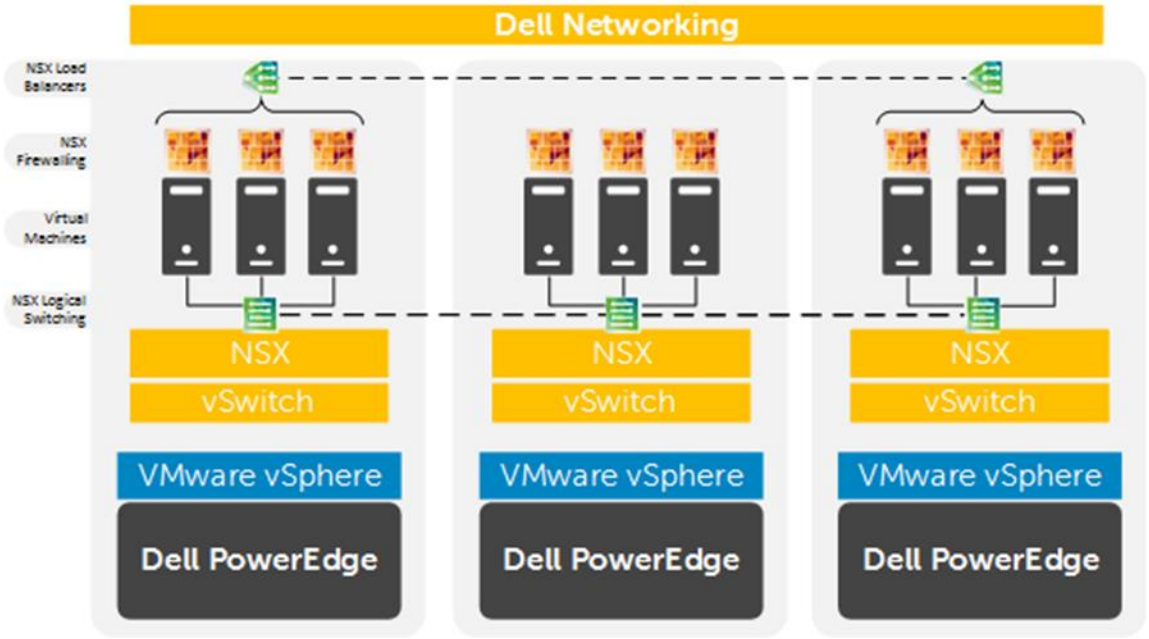

Key Features of Dell Open Networking and VMware NSX	
Power of Choice	Choose from best-of-breed open networking platforms, operating systems and applications.
Accelerated Innovation	Take advantage of open networking with open source standards-based tools and expertise to help accelerate innovation.
Open Networking Platform	All Dell Networking data center switches support the Open Network Install Environment (ONIE), allowing customers to choose between multiple operating systems and meet their unique needs.
Hardware VTEP Gateway	Layer 2 gateway through VXLAN Tunnel End Points (VTEP) bridges virtual and physical infrastructures.
Virtual Switching	VXLAN based network overlays enable logical layer 2 overlay extensions across a routed (L3) fabric within and across data center boundaries.
Virtual Routing	Dynamic routing between virtual networks performed in a distributed manner in the hypervisor kernel, and scale-out routing with active-active failover with physical routers.
Distributed Firewalling	Distributed stateful firewalling, embedded in the hypervisor kernel for up to 20 Gbps of firewall capacity per hypervisor host.
Load Balancing	L4-L7 load balancer with SSL offload and pass through, server health checks, and App Rules for programmability and traffic manipulation.

For more information on VMware NSX and integrated offers from Dell Networking please see the Dell Networking Solution Brief and the Reference architecture.

5.4 Scaling Guidance

The components are scaled either horizontally (by adding additional physical and virtual servers to each component of the solution architecture scales independently according to the desired number of supported users. Additional appliance nodes can be added at any time to expand the Virtual SAN SDS pool in a modular fashion. The scaling limit for Virtual SAN is restricted due to the limits of the Hypervisor so 64 Nodes in total per Cluster. The recommended limit by VMware with regards to the amount of VMs per Cluster is 2,000 so taking this into consideration we need 25 compute nodes for the lowest user density profile (FC430, Power User with 256 GB) which is 80 per Node. The image below shows a 2000 user Horizon with vSAN Block.

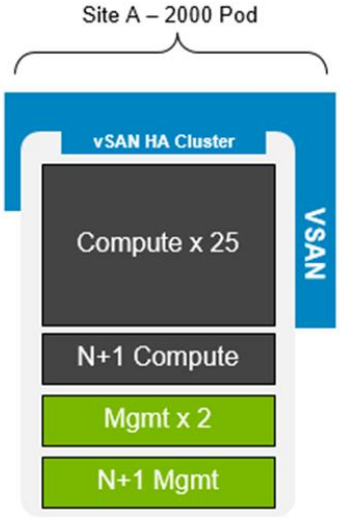

The recommendation from VMware is not to exceed 2,000 VM's/Sessions per Block and once we have more than 2000 VM's we would then split the management and compute into separate vSAN Clusters. The example below shows a scale out of a 10,000 user Horizon vSAN Pod with 2000 user blocks, each block contains its own Virtual Center, Horizon connection & composer servers.

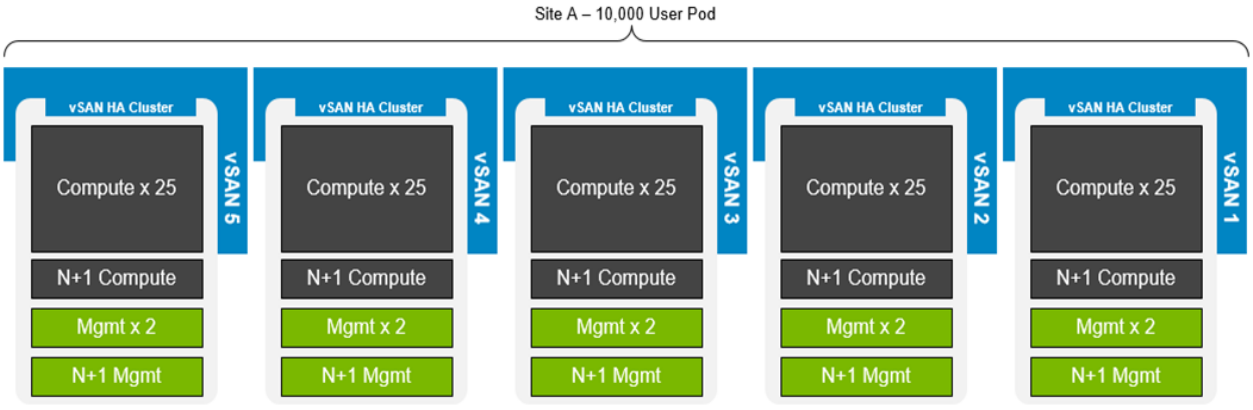

- The components are scaled either horizontally (by adding additional physical and virtual servers to the server pools) or vertically (by adding virtual resources to the infrastructure)
- Eliminate bandwidth and performance bottlenecks as much as possible
- Allow future horizontal and vertical scaling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

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

5.5 Solution high availability

High availability (HA) is offered to protect each layers of the solution architecture, individually if desired. Following the N+1 model, additional ToR switches for LAN, VMware vSAN are added to the Network layer and stacked to provide redundancy as required, additional compute and management hosts are added to their respective layers, vSphere clustering is introduced in the management layer, SQL is mirrored or clustered, an F5 device can be leveraged for load balancing.

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

Additional switches added to the existing thereby equally spreading each host's network connections across multiple switches.

Additional ESXi hosts added in the compute or management layers to provide N+1 protection.

Applicable VMware Horizon infrastructure server roles are duplicated and spread amongst management host instances where connections to each are load balanced via the addition of F5 appliances.

5.5.1 VMware vSAN HA/ FTT configuration

The minimum configuration required for Dell EMC VxRail is 3 ESXi hosts. The issue with having a 3-Node cluster is if one node fails there is nowhere to rebuild the failed components, so 3 node clusters should be used only for POC or non-production.

The virtual machines that are deployed via Horizon are policy driven and one of these policy settings is Number of failures to tolerate (FTT). The default value for FTT is FTT=1 so that will make a mirrored copy of the Virtual Machines VMDK, so if the VMDK is 40Gb in size then 80Gb of virtual machine space is needed.

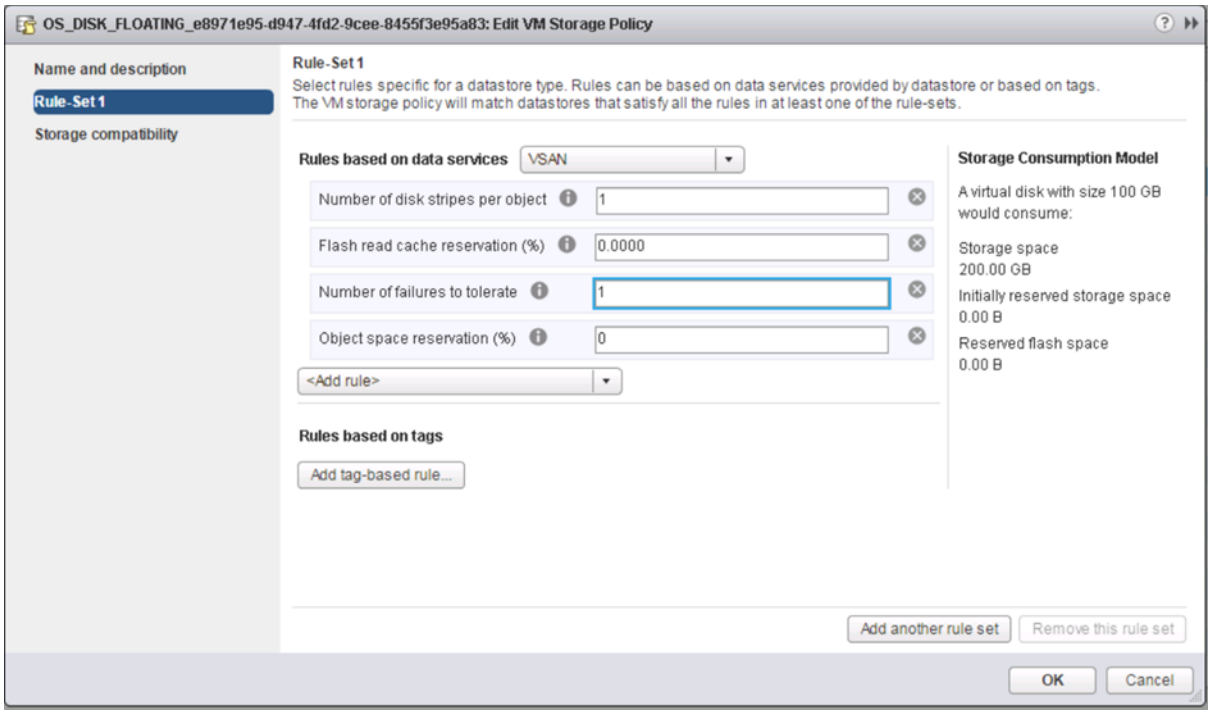

The recommended configuration by VMware for a VMware vSAN Cluster with FTT=1 and RAID 1 is four nodes and this ensures that the virtual machines are fully protected during operational & maintenance activities. This configuration can also survive another failure even when there is a host already in maintenance mode.

5.5.2 vSphere HA

Both compute and management hosts are identically configured, within their respective tiers. The management Tier leverages the shared VMware vSAN storage so can make full use of vSphere HA and VxRail vCompute nodes can be added to add HA to the configured storage policy. The hosts can be configured in an HA cluster following the boundaries of VMware vSAN 6.2 limits dictated by VMware (6,400 VMs per VMware vSAN Cluster). This will result in multiple HA clusters managed by multiple vCenter servers.

The number of supported VMs (200*) is a soft limit and this is discussed further in section 6 of this document.

VMware vSAN Limits	Minimum	Maximum
Number of supported ESXi hosts per VMware vSAN cluster	3	64
Number of supported VMs per host	n/a	200*
Number of supported VMs per VMware vSAN Cluster	n/a	6400
Disk groups per host	1	5
HDDs per disk group	1	7
SSDs per disk group	1	1
Components per host	n/a	9000
Components per object	n/a	64

5.5.3 Horizon infrastructure protection

VMware Horizon infrastructure data protection with Dell Data Protection – http://dell.to/1ed2dQf

5.5.4 Management server high availability

The applicable core Horizon roles will be load balanced via DNS by default. In environments requiring HA, F5 BigIp can be introduced to manage load-balancing efforts. Horizon, VCS and vCenter configurations (optionally vCenter Update Manager) are stored in SQL which will be protected via the SQL mirror.

If the customer desires, some Role VMs can be optionally protected further via the form of a cold stand-by VM residing on an opposing management host. A vSphere scheduled task can be used, for example, to clone the VM to keep the stand-by VM current. Note – In the HA option, there is no file server VM, its duties have been replaced by introducing a NAS head.

The following will protect each of the critical infrastructure components in the solution:

- The Management hosts will be configured in a vSphere cluster.
- SQL Server mirroring is configured with a witness to further protect SQL.

5.5.5 Horizon Connection Server high availability

The HCS role as a VM and running in a VMware HA Cluster, the HCS server can be guarded against a physical server failure.

For further protection in an HA configuration, deploy multiple replicated Horizon Connection Server instances in a group to support load balancing and HA. Replicated instances must exist on within a LAN connection environment it is not recommended VMware best practice to create a group across a WAN or similar connection.

5.5.6 SQL Server high availability

HA for SQL is provided via AlwaysOn using either Failover Cluster Instances or Availability Groups. This configuration protects all critical data stored within the database from physical server as well as virtual server problems. DNS is used to control access to the primary SQL instance. Place the principal VM that will host the primary copy of the data on the first Management host. Additional replicas of the primary database are placed on subsequent Management hosts.

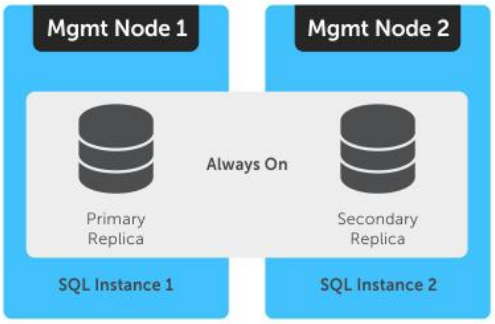

Please refer to these links for more information: LINK1 and LINK2.

5.6 VMware Horizon communication flow

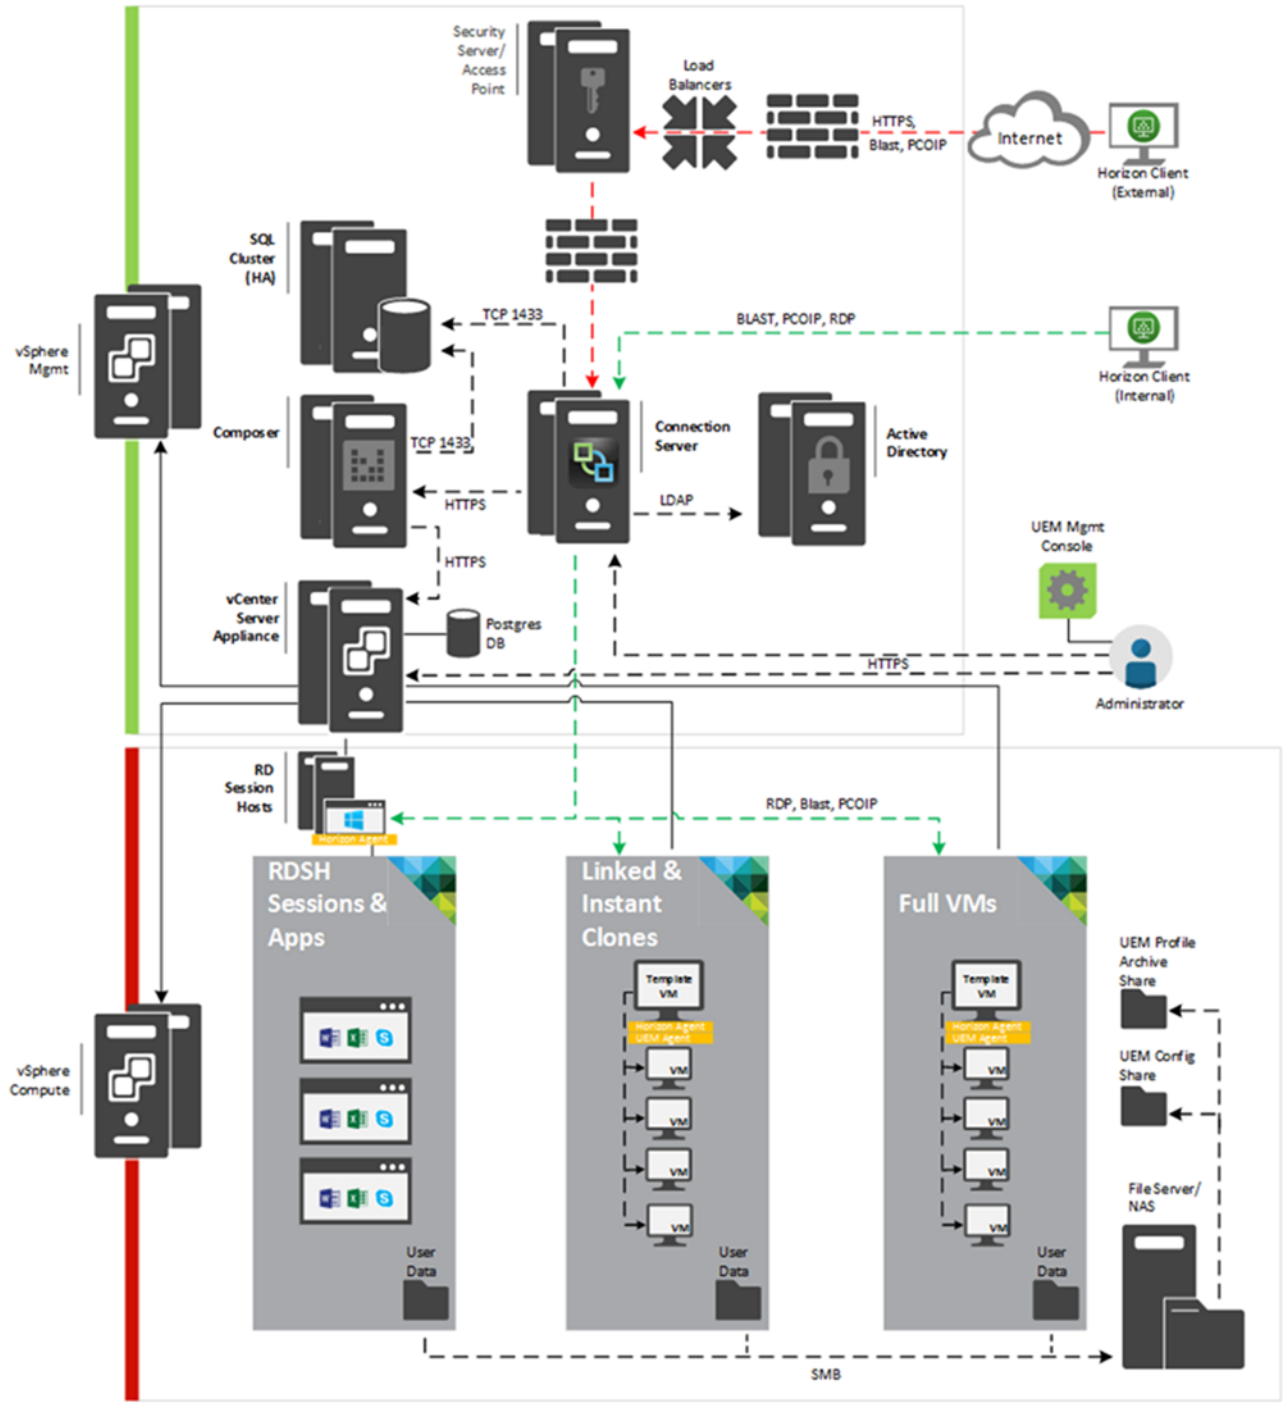

6 Solution performance and testing

At the time of publication, here are the available density recommendations. The below user densities were achieved by following the VMware best practices of FTT=1 and a reserved slack space of 30%.

*The soft limit for the amount of VMs supported per host is 200, this is due to number of total objects that are supported per cluster. This is a factor in very large clusters but for small to medium Cluster configurations this should not be an issue.The hardware configuration details are listed in Section 3.3

User density summary

Hypervisor	Provisioning	Profile	Template OS	Config	User Density
6.0 Update 2	Linked Clone	Task	Windows 10	V470/V470F-B5	150
6.0 Update 2	Linked Clone	Knowledge	Windows 10	V470/V470F-B5	130
6.0 Update 2	Linked Clone	Power	Windows 10	V470/V470F-B5	105
6.0 Update 2	Linked Clone	Task	Windows 10	V470/V470F-C7	230*
6.0 Update 2	Linked Clone	Knowledge	Windows 10	V470/V470F-C7	170
6.0 Update 2	Linked Clone	Power	Windows 10	V470/V470F-C7	140
6.0 Update 2	RDS	Task	W2K12R2	V470/V470F-C7	350
6.0 Update 2	Linked Clone	Knowledge	Windows 10	GPU-C7 M6Q-1Q	32

The detailed validation results and analysis of these reference designs are in the next section.

6.1 Test and Performance Analysis Methodology

6.1.1 Testing process

In order to ensure the optimal combination of end-user experience (EUE) and cost-per-user, performance analysis and characterization (PAAC) on Dell Wyse Datacenter solutions is carried out using a carefully designed, holistic methodology that monitors both hardware resource utilization parameters and EUE during load-testing.

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

At different times during testing, the testing team will complete some manual “User Experience” Testing while the environment is under load. This will involve a team member logging into a session during the run and

completing tasks similar to the User Workload description. While this experience will be subjective, it will help provide a better understanding of the end user experience of the desktop sessions, particularly under high load, and ensure that the data gathered is reliable.

6.1.1.1 Load Generation

Login VSI by Login Consultants is the de-facto industry standard tool for testing VDI environments and server-based computing (RDSH environments). It installs a standard collection of desktop application software (e.g. Microsoft Office, Adobe Acrobat Reader) on each VDI desktop; it then uses launcher systems to connect a specified number of users to available desktops within the environment. Once the user is connected, the workload is started via a logon script which starts the test script once the user environment is configured by the login script. Each launcher system can launch connections to a number of 'target' machines (i.e. VDI desktops). The launchers and Login VSI environment are configured and managed by a centralized management console.

Additionally, the following login and boot paradigm is used:

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

6.1.1.2 Profiles and Workloads

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

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

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

Profile to workload mapping

Profile Name	Workload
Task Worker	Login VSI Task worker
Knowledge Worker	Login VSI Knowledge worker
Power Worker	Login VSI Power worker
Graphics LVSI Power + ProLibrary	Graphics - Login VSI Power worker with ProLibrary
Graphics LVSI Custom	Graphics – LVSI Custom

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

Login VSI Task Worker Workload

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

Login VSI Knowledge Worker Workload

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

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

Login VSI Power Worker Workload

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

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

Graphics - Login VSI Power Worker with ProLibrary workload

For lower performance graphics testing where lower amounts of graphics memory are allocated to each VM, the Power worker + Pro Library workload is used. The Login VSI Pro Library is an add-on for the Power worker workload which contains extra content and data files. The extra videos and web content of the Pro Library utilizes the GPU capabilities without overwhelming the lower frame buffer assigned to the desktops. This type of workload is typically used with high density vGPU and sVGA or other shared graphics configurations.

Graphics – LVSI Custom workload

This is a custom Login VSI workload specifically for higher performance, intensive graphics testing. For this workload, SPECwpc benchmark application is installed to the client VMs. During testing, a script is started that launches SPECwpc which executes the Maya and sw-03 modules for high performance tests and module sw-03 only for high density tests. The usual activities such as Office application execution are not performed with this workload. This type of workload is typically used for lower density/high performance pass-through, vGPU, and other dedicated, multi-user GPU configurations.

6.1.2 Resource Monitoring

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

6.1.2.1 GPU Resources

ESXi hosts

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

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

6.1.2.2 VMware vCenter

VMware vCenter is used for VMware vSphere-based solutions to gather key data (CPU, Memory, Disk and Network usage) from each of the compute hosts during each test run. This data is exported to .csv files for single hosts and then consolidated to show data from all hosts (when multiple are tested). While the report does not include specific performance metrics for the Management host servers, these servers are monitored during testing to ensure they are performing at an expected performance level with no bottlenecks.

6.1.3 Resource Utilization

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

Resource utilization thresholds

Parameter	Pass/Fail Threshold
Physical Host CPU Utilization (AHV & ESXi hypervisors)*	100%
Physical Host CPU Utilization (Hyper-V)	85%
Physical Host Memory Utilization	85%
Network Throughput	85%
Storage IO Latency	20ms

6.2 Test Configuration Details

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

Hardware and software test components.

Component	Description/Version
Hardware platform(s)	VxRail V470 B5 & C7
Hypervisor(s)	ESXi 6.0 Update 2

Broker technology	Horizon 7
Broker database	Microsoft SQL 2012
Management VM OS	Windows Server 2012 R2 (Connection Server & Database)
Virtual desktop OS	Windows 10 Enterprise
Office application suite	Office Professional 2016
Login VSI test suite	Version 4.1

6.2.1 Compute VM configurations

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

Desktop VM specifications

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

Screen resolutions

User Profile	Screen Resolution
Task Worker	1280 X 720
Knowledge Worker	1920 X 1080
Power Worker	1920 X 1080

Graphics LVSI Power + ProLibrary	1920 X 1080
Graphics LVSI Custom – Density	1920 X 1080
Graphics LVSI Custom - Performance	1920 X 1080

6.2.2 Platform Configuration

The hardware configuration details are listed in Section 3.3

6.3 Test results and analysis

The following table summarizes the test results for the compute hosts using the various workloads and configurations. Refer to the prior section for platform configuration details.

Test result summary

Platform Config	Hypervisor	Broker & Provisioning	Login VSI Workload	Density Per Host	Avg CPU	Avg Mem Consumed	Avg Mem Active	Avg IOPS / User
V470/V470F-B5	ESXi 6.0 U2	Horizon 7 & Linked Clone	Task Worker	150	68%	337 GB	211 GB	6
V470/V470F-B5	ESXi 6.0 U2	Horizon 7 & Linked Clone	Knowledge Worker	130	88%	329 GB	312 GB	6.5
V470/V470F-B5	ESXi 6.0 U2	Horizon 7 & Linked Clone	Power Worker	105	82%	372 GB	296 GB	7
V470/V470F-C7	ESXi 6.0 U2	Horizon 7 & Linked Clone	Task Worker	230*	80%	495 GB	180 GB	9.8
V470/V470F-C7	ESXi 6.0 U2	Horizon 7 & Linked Clone	Knowledge Worker	170	85%	496 GB	180 GB	9
V470/V470F-C7	ESXi 6.0 U2	Horizon 7 & Linked Clone	Power Worker	140	85%	506 GB	196 GB	11.75
V470/V470F-C7	ESXi 6.0 U2	Horizon 7 & RDS	Task Worker	350	92%	211 GB	132 GB	5.7
GPU-C7 M60Q-1Q	ESXi 6.0 U2	Horizon 7 & Linked Clone	Task Worker	32	83%	171 GB	110 GB	18

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

Avg CPU: This is the average CPU usage over the steady state period. For clusters, this represents the combined average CPU usage of all compute hosts. On the latest Intel series processors, the ESXi host CPU metrics will exceed the rated 100% for the host if Turbo Boost is enabled (by default). An additional 35% of

CPU is available from the Turbo Boost feature but this additional CPU headroom is not reflected in the VMware vSphere metrics where the performance data is gathered. Therefore, CPU usage for ESXi hosts is adjusted and a line indicating the potential performance headroom provided by Turbo boost is included in each CPU graph.

Avg Consumed Memory: Consumed memory is the amount of host physical memory consumed by a virtual machine, host, or cluster. For clusters, this is the average consumed memory across all compute hosts over the steady state period.

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

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

Avg Net Mbps/User: Amount of network usage over the steady state period divided by the number of users. For clusters, this is the combined average of all compute hosts over the steady state period divided by the number of users on a host.

6.3.1 VxRail-V470/V470F B5

The results shown in the below tables focus on the compute hosts, the management hosts will have 30% less due to amount of resources needed for the management VMs which include the VxRail Manager.

6.3.1.1 Task Worker, 150 users, ESXi 6.0 U2, Horizon 7 linked clones

The following graph shows the output from the Login VSI Analyzer. VSI Max was not reached on this test run.

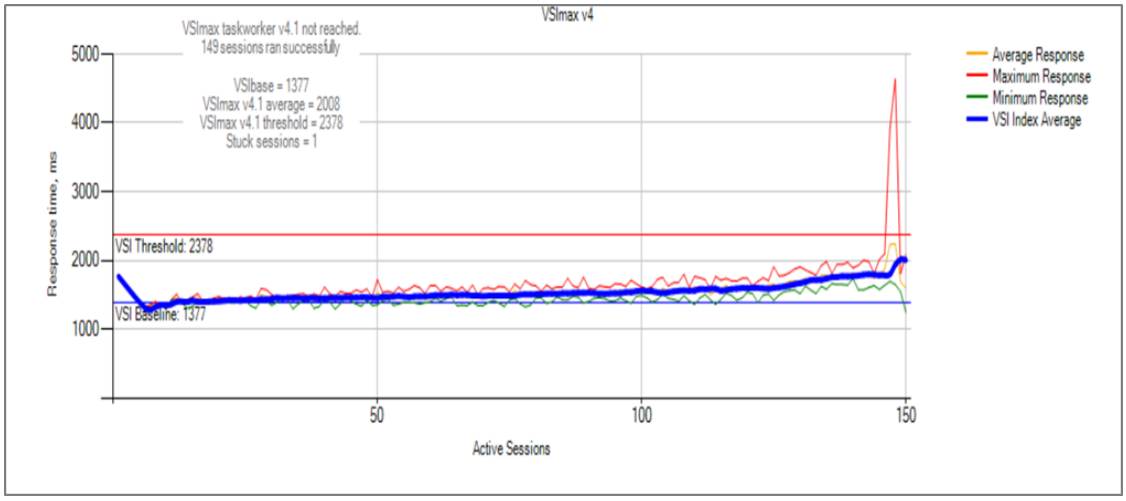

The Task Worker workload didn't quite reach the 85% mark but VMs configured with only one vCPU tend to load the CPU less. The E5-2660 v4 processor was used for the V470/V470F-B5 testing.

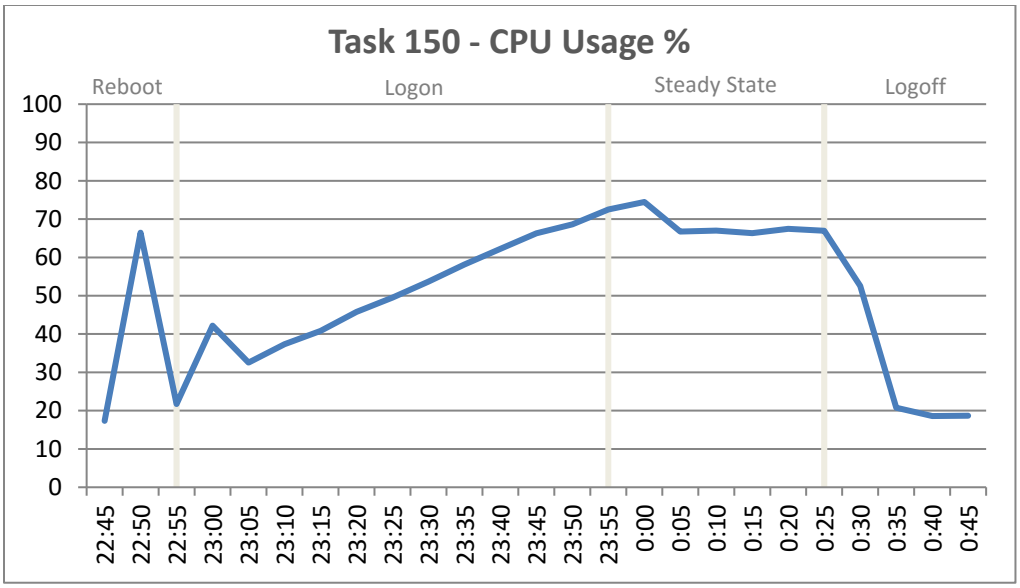

Latency on the datastore spiked temporarily during the boot phase of the test but quickly settled once all the VMs were booted. For the logon and steady state phase of each test, the latency remained well below the 20ms threshold reaching a max of 2-3 ms during the test run.

The IOPS peaked during the boot phase and for each profile test and then settled thereafter during the login phase and reduced once steady state was reached.

This chart was captured from within vSphere and was a feature released with VSAN6.2 so we do not need to use VSAN Observer as was previously the case with past VSAN validations.

The statistics below are on a per host basis so as VSAN scales linearly; to calculate the total IOPs for a three node cluster you would multiple by three.

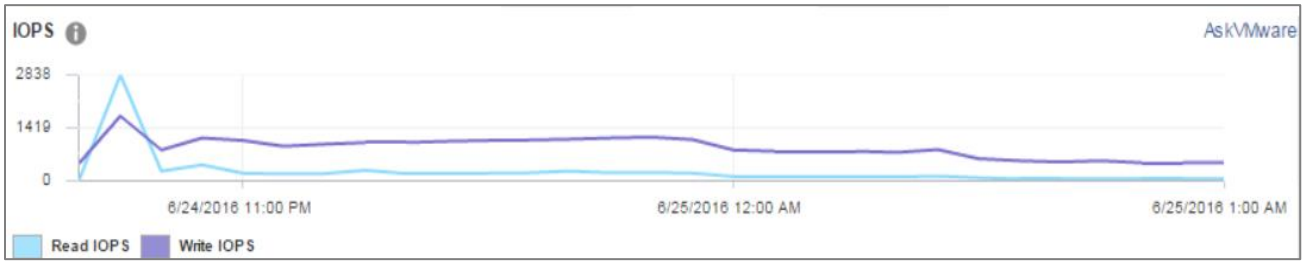

Memory usage is monitored on ESXi host, memory usage monitored are consumed, active, balloon and swap used, as swap and ballooning usage would indicate host memory reached saturation point and the VM performance may start deteriorating. All tests were carried out on hosts with 384GB physical memory installed and no swapping or ballooning was experienced during the tests.

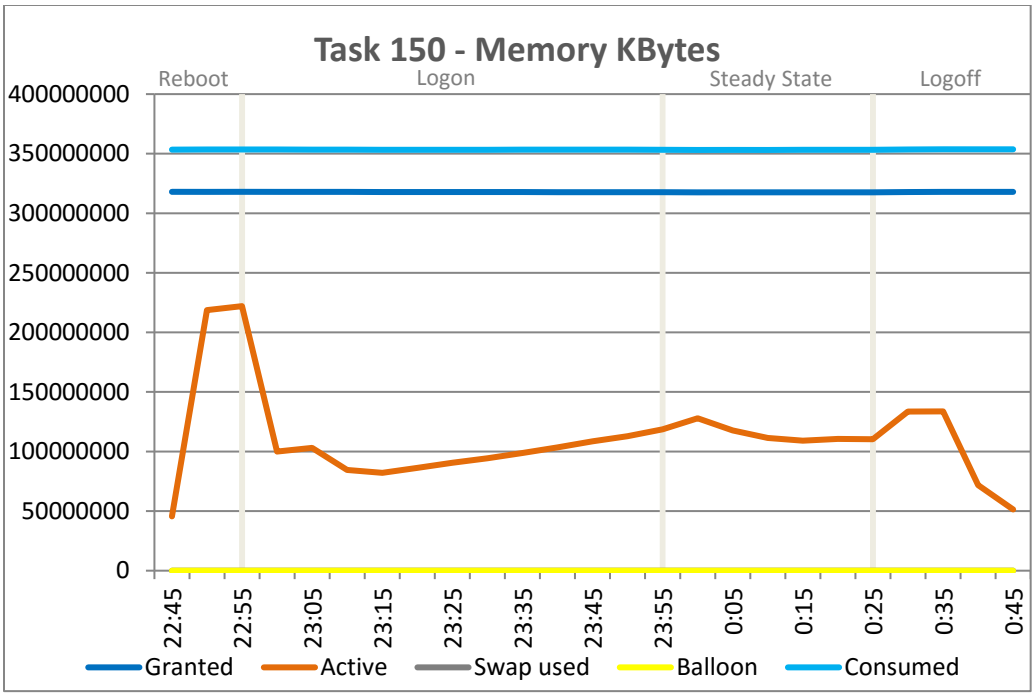

There were no issues with network usage on any of the test runs. All tests showed that the reboot of the VM pool before beginning Login VSI testing produced the highest spike in network activity. There is a significant reduction in activity once the steady state phase is reached after all machines have logged on.

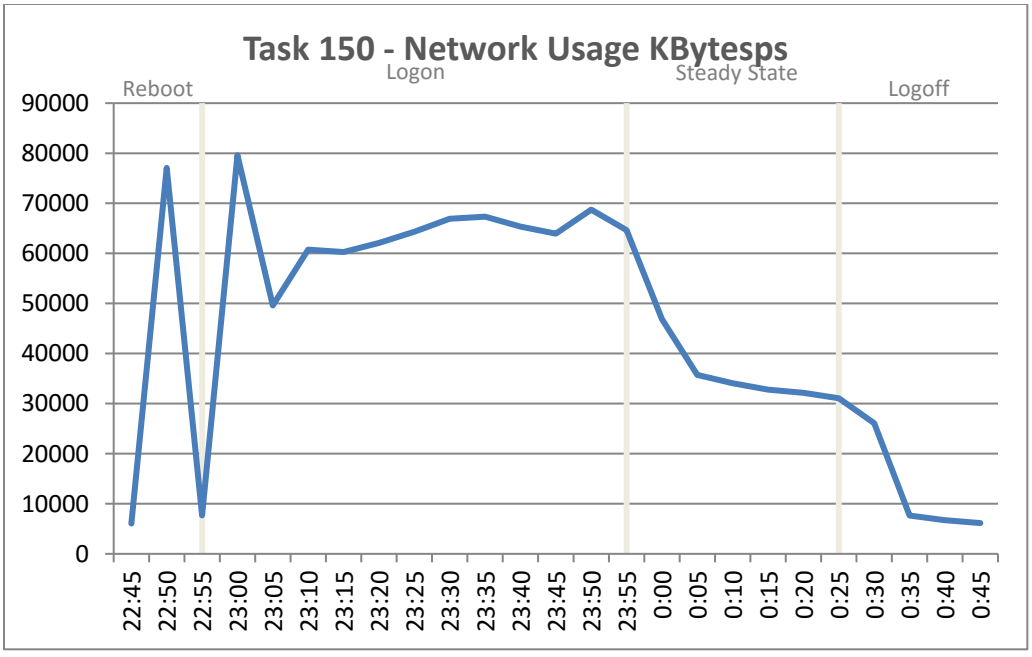

6.3.1.2 Knowledge Worker, ESXi 6.0U2, Horizon 7 linked clones

The following graph shows the output from the Login VSI Analyzer. VSI Max was not reached on this test run.

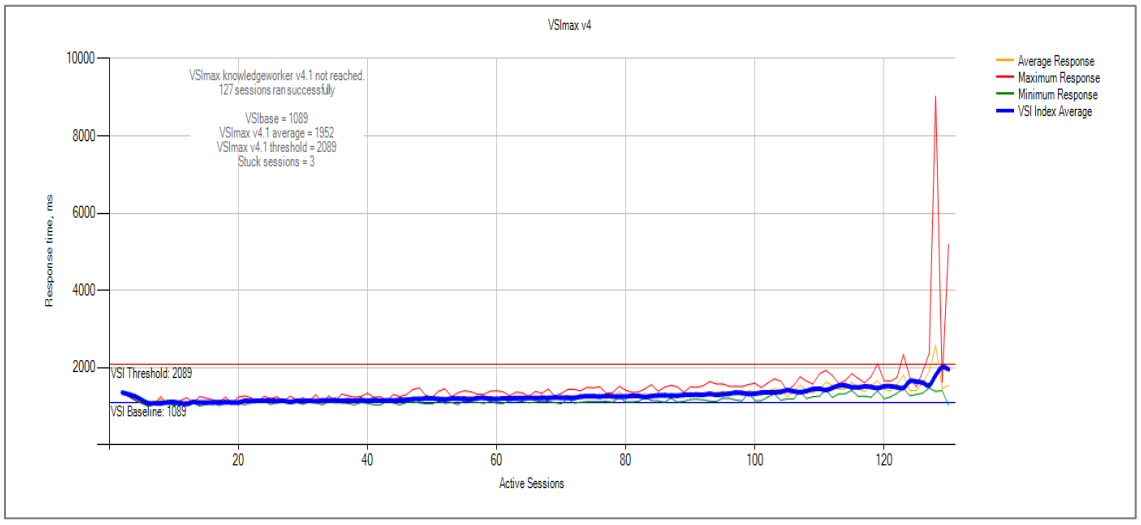

Maximum CPU utilization for the Knowledge workers was in the region of the 90% threshold indicating the number of users tested was appropriate.

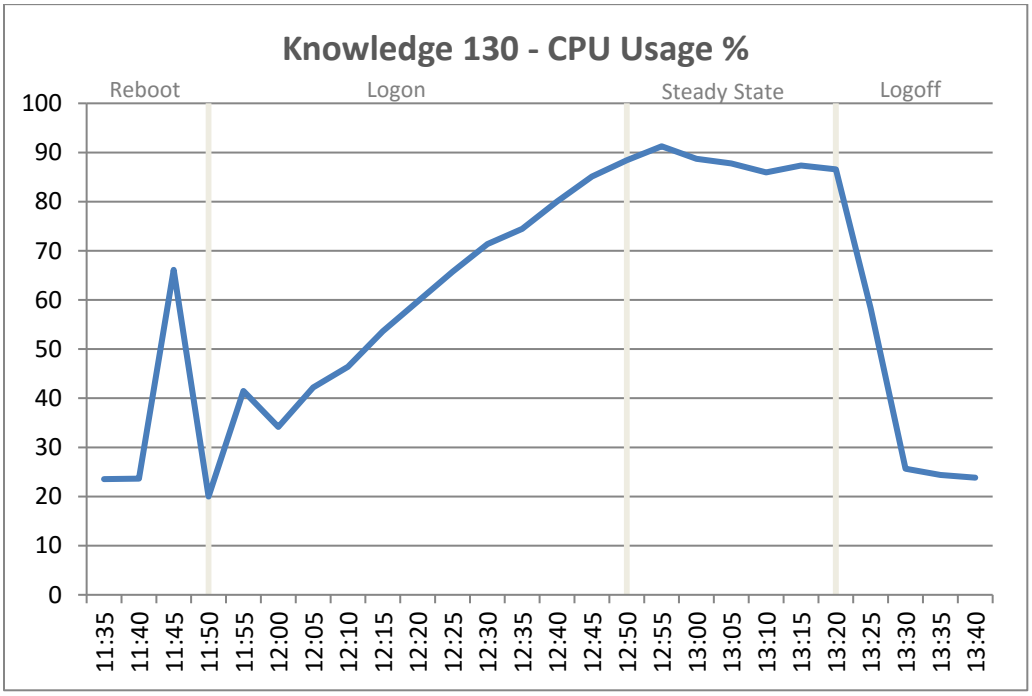

Latency on the datastore spiked temporarily during the boot phase of the test but quickly settled once all the VMs were booted. For the logon and steady state phase of each test, the latency remained well below the 20ms threshold reaching a max of 2-3 ms during the test run.

The IOPS peaked during the boot phase and for each profile test and then settled thereafter during the login phase and reduced once steady state was reached.

This chart was captured from within vSphere and was a feature released with VSAN6.2 so we do not need to use VSAN Observer as was previously the case with past VSAN validations.

The statistics below are on a per host basis so as VSAN scales linearly; to calculate the total IOPs for a three node cluster you would multiple by three.

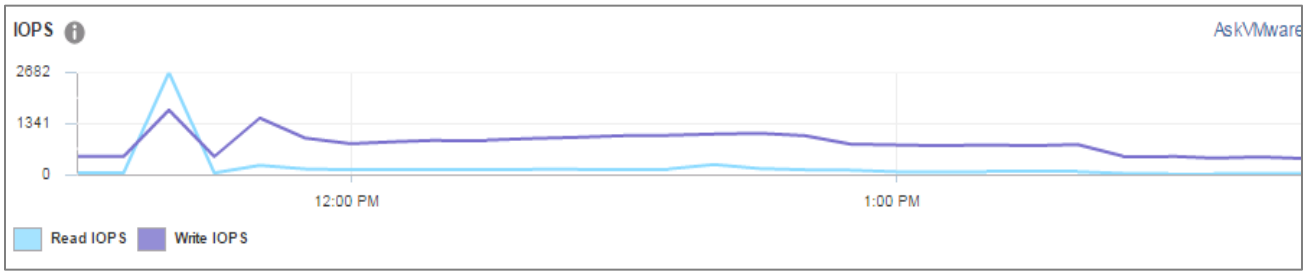

Memory usage is monitored on ESXi host, memory usage monitored are consumed, active, balloon and swap used, as swap and ballooning usage would indicate host memory reached saturation point and the VM performance may start deteriorating. All tests were carried out on hosts with 384GB physical memory installed and no swapping or ballooning was experienced during the tests.

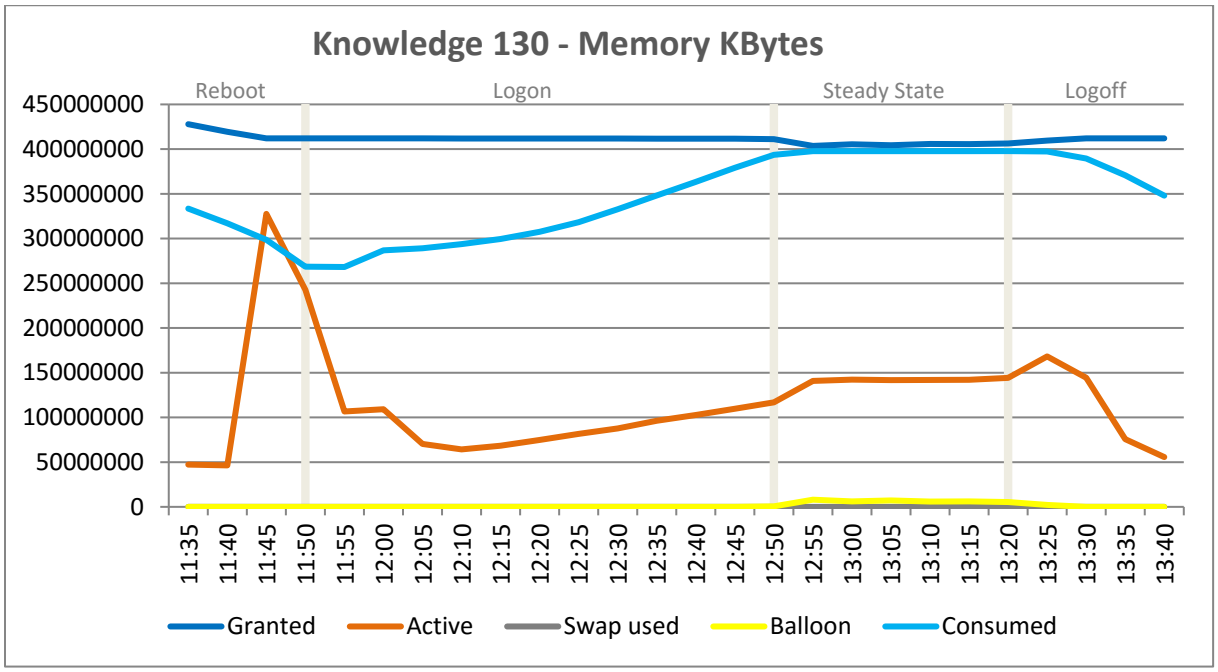

There were no issues with network usage on any of the test runs. All tests showed that the reboot of the VM pool before beginning Login VSI testing produced the highest spike in network activity. There is a significant reduction in activity once the steady state phase is reached after all machines have logged on.

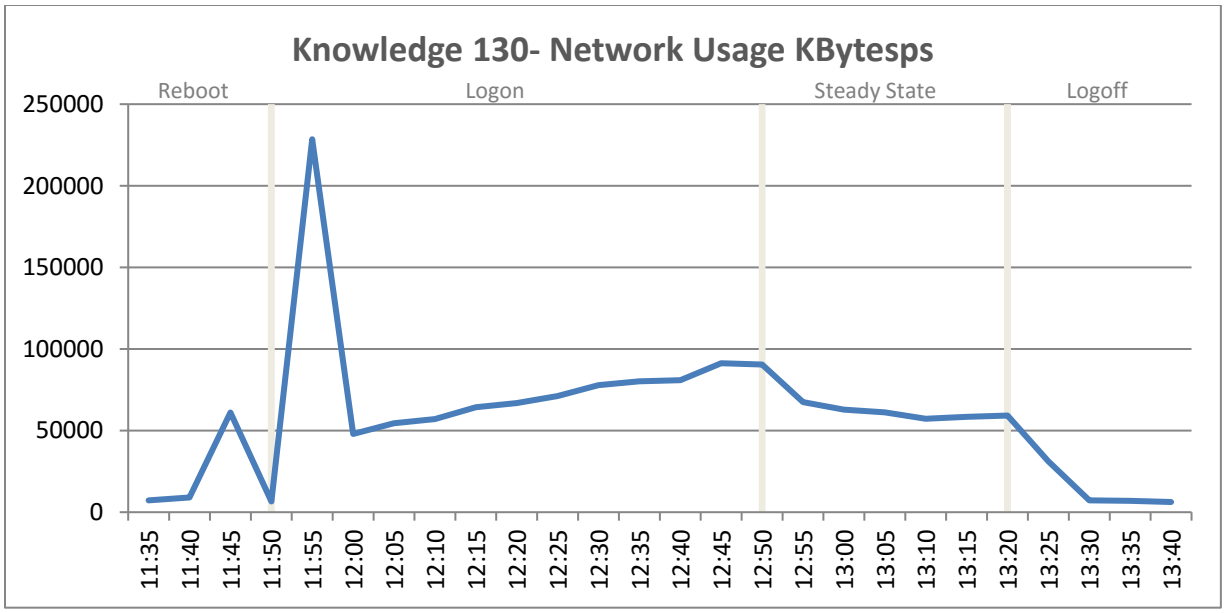

6.3.1.3 Power Worker, 100 users, ESXi 6.0U2, Horizon 7 linked clones

The following graph shows the output from the Login VSI Analyzer. VSI Max was not reached on this test run.

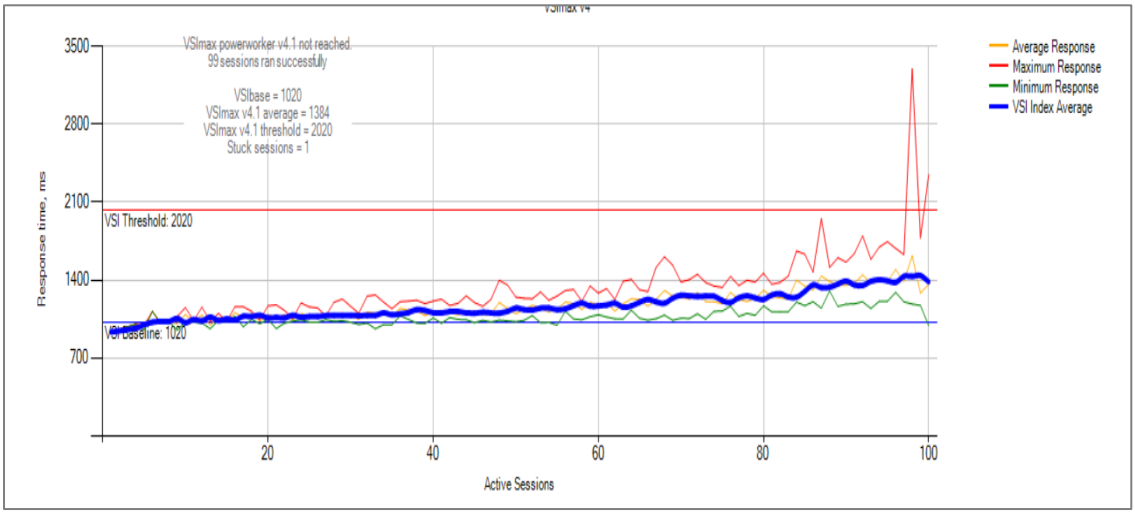

Maximum CPU utilization for Power workers was under 85% threshold indicating the number of users tested was appropriate.

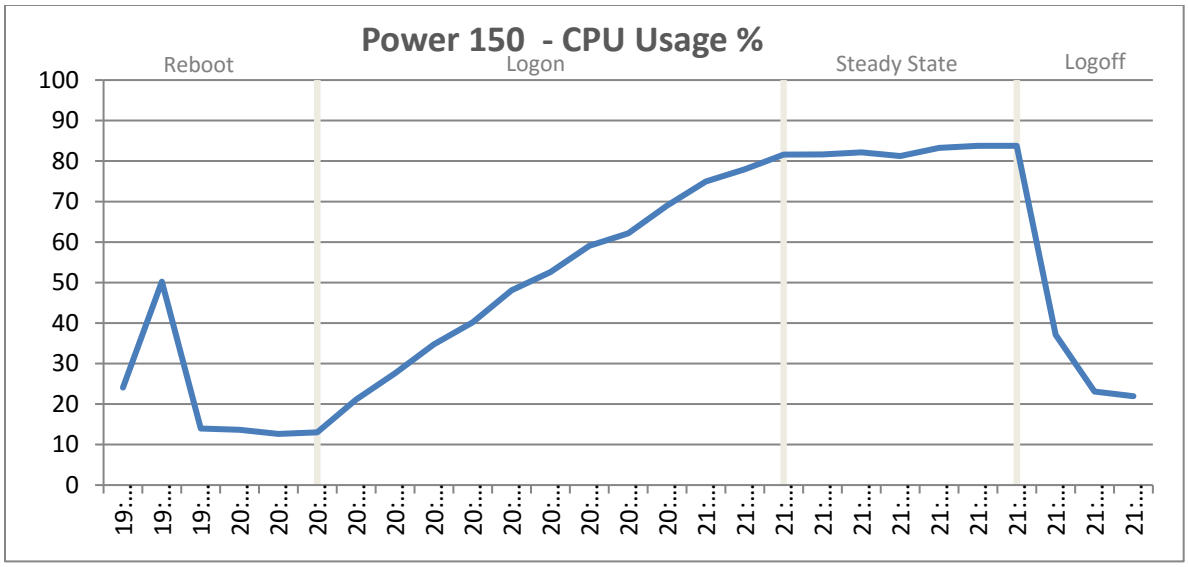

Latency on the datastore spiked temporarily during the boot phase of the test but quickly settled once all the VMs were booted. For the logon and steady state phase of each test, the latency remained well below the 20ms threshold reaching a max of 2-3 ms during the test run.

The IOPS peaked during the boot phase and for each profile test and then settled thereafter during the login phase and reduced once steady state was reached.

This chart was captured from within vSphere and was a feature released with VSAN6.2 so we do not need to use VSAN Observer as was previously the case with past VSAN validations.

The statistics below are on a per host basis so as VSAN scales linearly; to calculate the total IOPs for a three node cluster you would multiple by three.

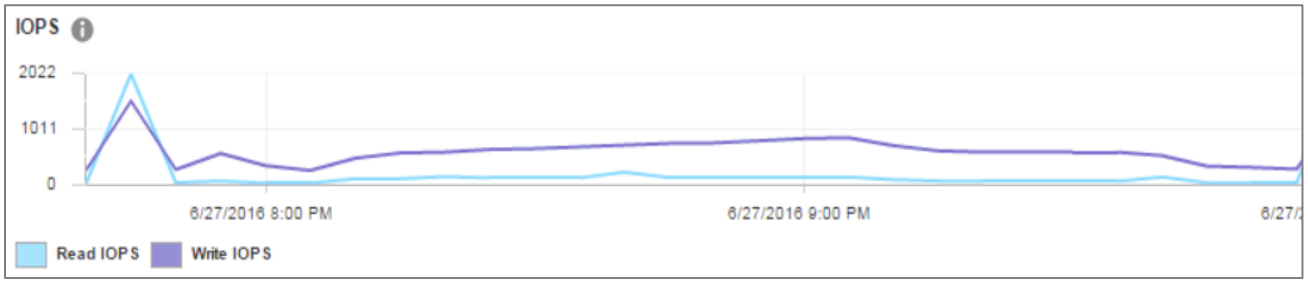

Memory usage is monitored on ESXi host, memory usage monitored are consumed, active, balloon and swap used, as swap and ballooning usage would indicate host memory reached saturation point and the VM performance may start deteriorating. All tests were carried out on hosts with 384GB physical memory installed and no swapping or ballooning was experienced during the tests.

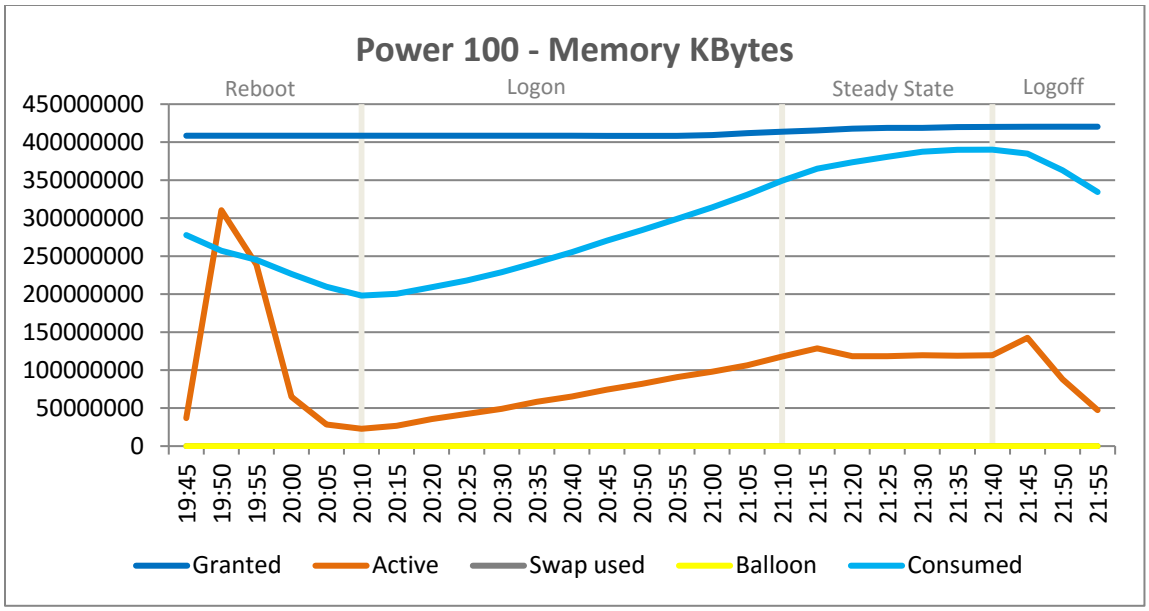

There were no issues with network usage on any of the test runs. All tests showed that the reboot of the VM pool before beginning Login VSI testing produced the highest spike in network activity. There is a significant reduction in activity once the steady state phase is reached after all machines have logged on.

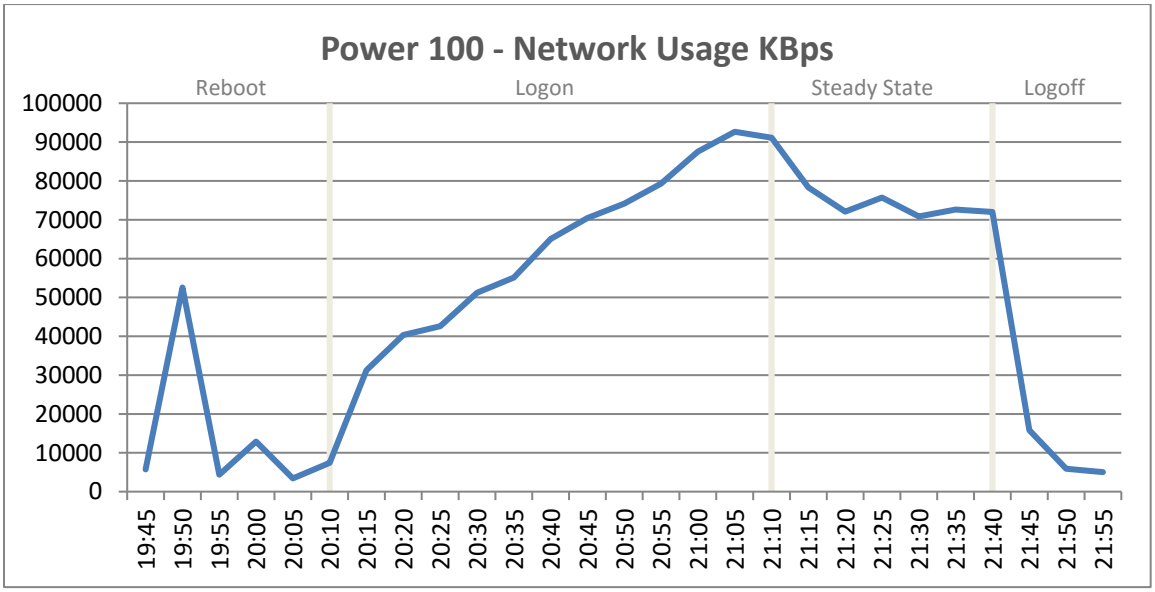

6.3.2 VxRail-V470/V470F C7

6.3.2.1 Task Worker, 230 users, ESXi 6.0 U2, Horizon 7 linked clones

The following graph show the output from the Login VSI Analyzer for each V470-C7 test run. VSI Max was not reached reached

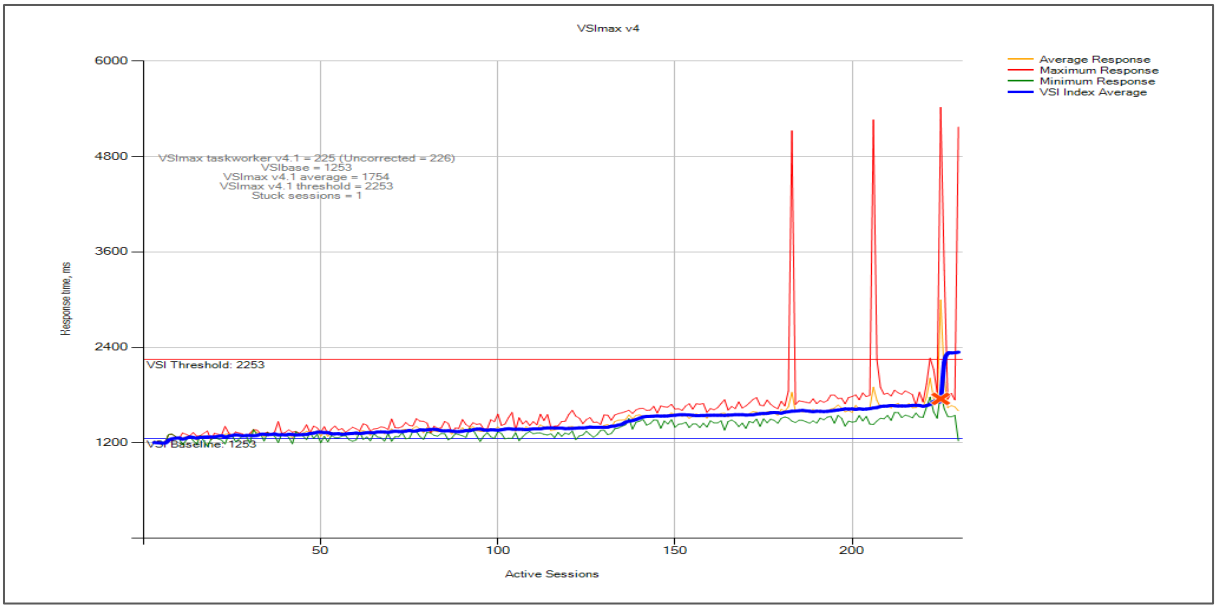

Maximum CPU utilization for the Knowledge workers was in the region of the 80% threshold indicating the number of users tested was appropriate.

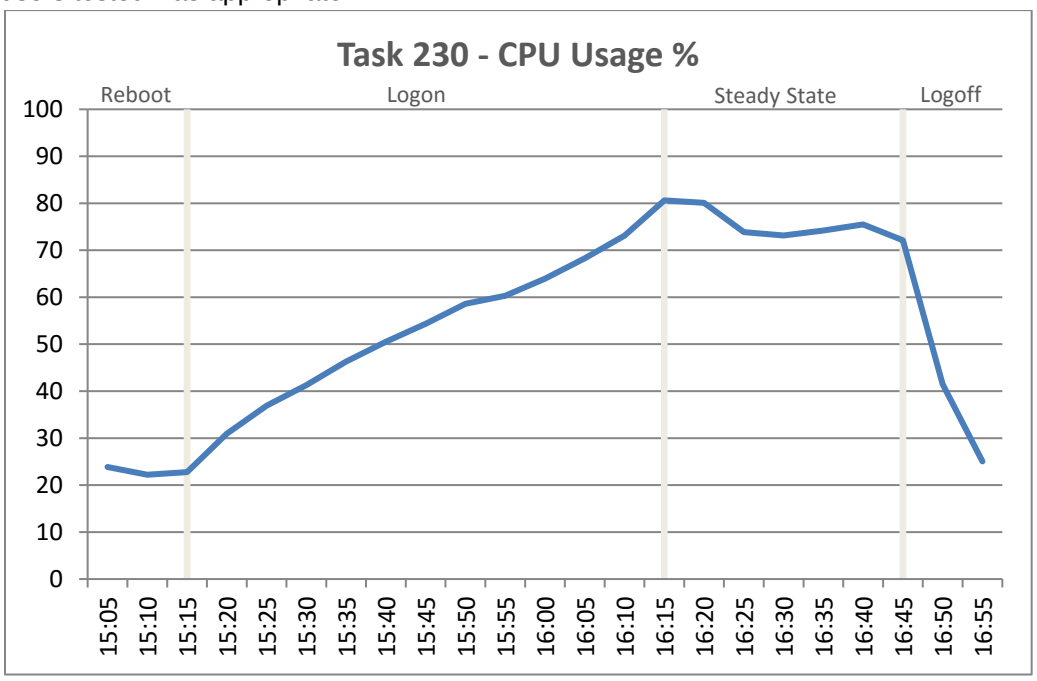

Latency on the datastore spiked temporarily during the boot phase of the test but quickly settled once all the VMs were booted. For the logon and steady state phase of each test, the latency remained well below the 20ms threshold reaching a max of 2-3 ms during the test run.

The IOPS peaked during the boot phase and for each profile test and then settled thereafter during the login phase and reduced once steady state was reached.

This chart was captured from within vSphere and was a feature released with VSAN6.2 so we do not need to use VSAN Observer as was previously the case with past VSAN validations.

The statistics below are on a per host basis so as VSAN scales linearly; to calculate the total IOPs for a three node cluster you would multiple by three.

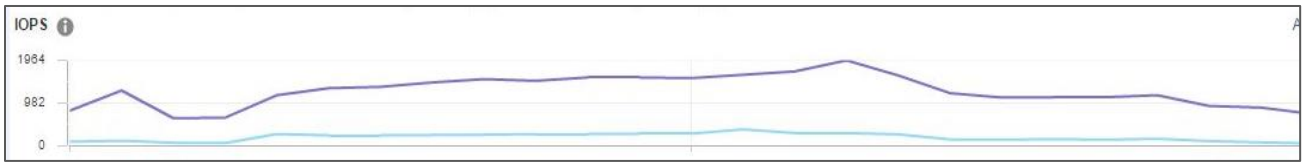

Memory usage is monitored on the ESXi host, memory usage metrics monitored are consumed, active, balloon and swap used, as swap and ballooning usage would indicate host memory reached saturation point and the VM performance may start to deteriorating. All tests were carried out on hosts with 512 GB of physical memory installed.

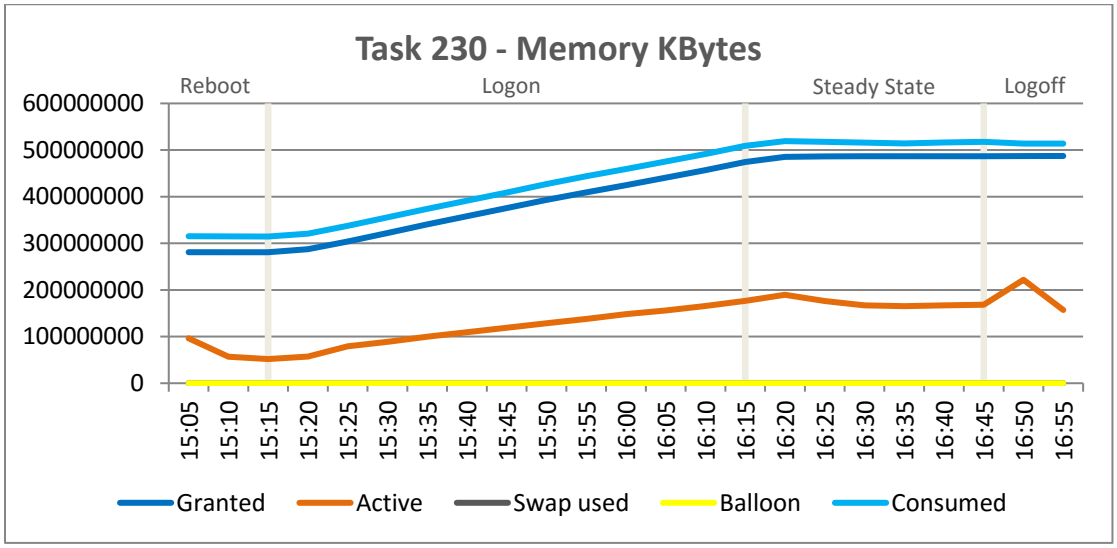

There were no issues with network usage on any of the test runs. There is a significant reduction in activity once the steady state phase is reached after all machines have logged on.

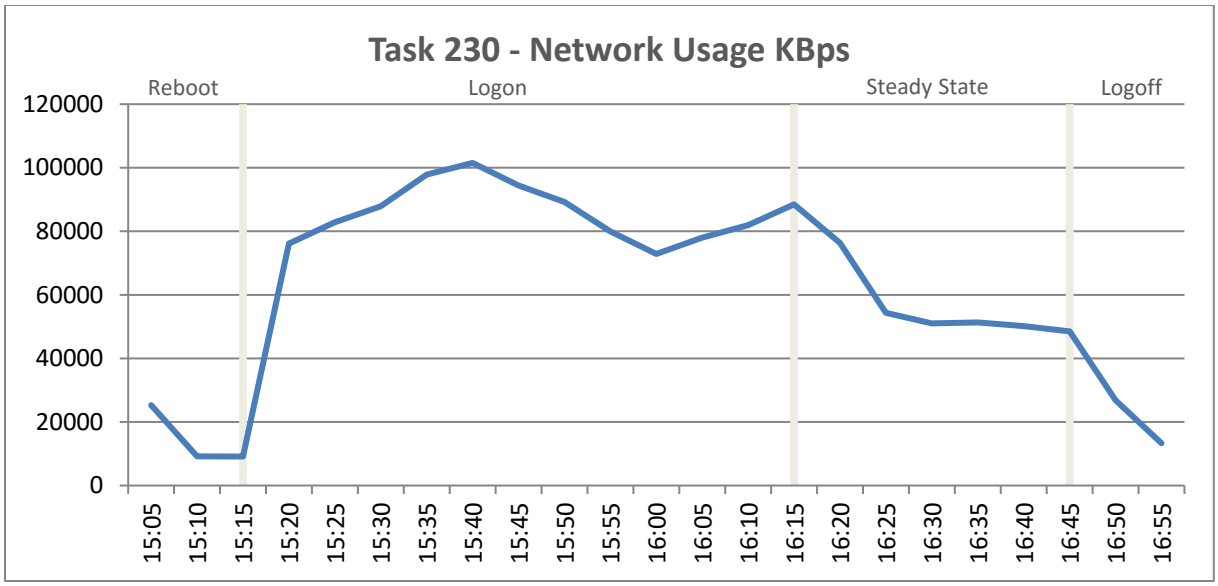

6.3.2.2 Knowledge Worker, ESXi 6.0U2, Horizon 7 linked clones

The following graph show the output from the Login VSI Analyzer for each V470-C7 test run. VSI Max was not reached reached

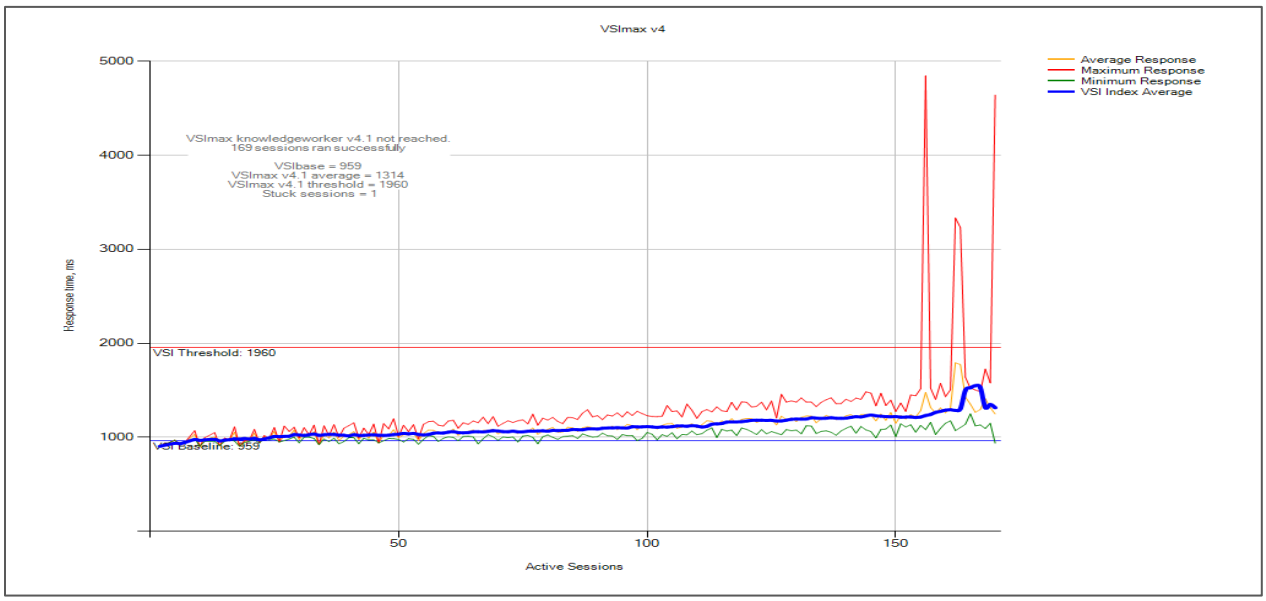

Maximum CPU utilization was in the region of the 85% threshold indicating the number of users tested was appropriate.

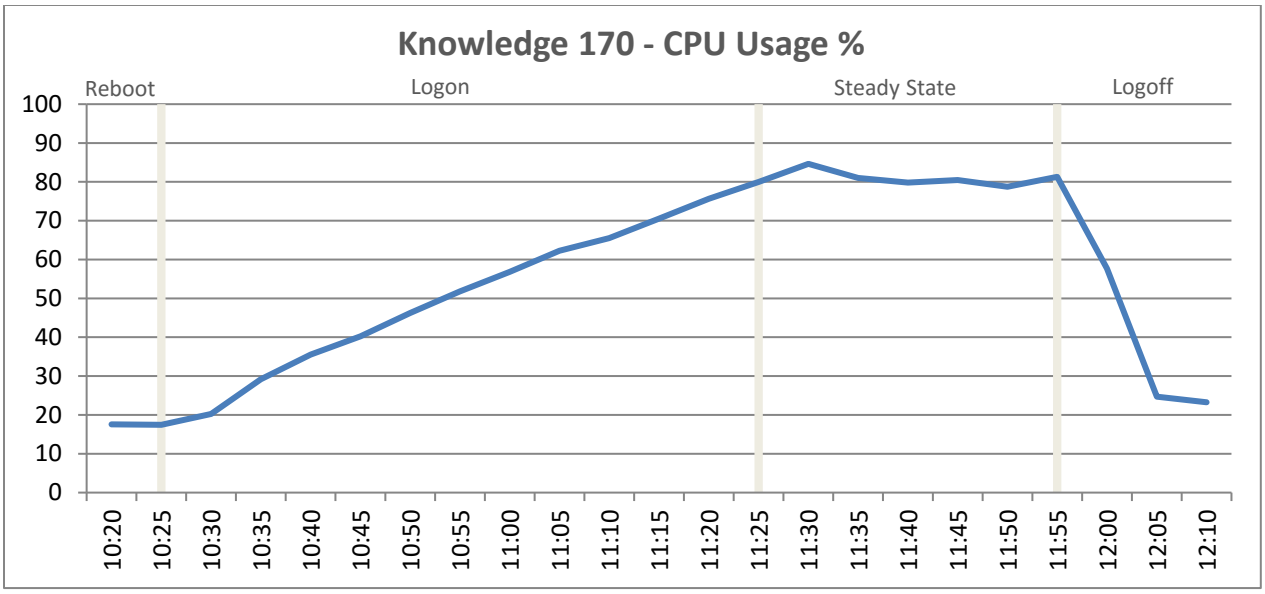

Latency on the datastore spiked temporarily during the boot phase of the test but quickly settled once all the VMs were booted. For the logon and steady state phase of each test, the latency remained well below the 20ms threshold reaching a max of 2-3 ms during the test run.

The IOPS peaked during the boot phase and for each profile test and then settled thereafter during the login phase and reduced once steady state was reached.

This chart was captured from within vSphere and was a feature released with VSAN6.2 so we do not need to use VSAN Observer as was previously the case with past VSAN validations.

The statistics below are on a per host basis so as VSAN scales linearly; to calculate the total IOPs for a three node cluster you would multiple by three.

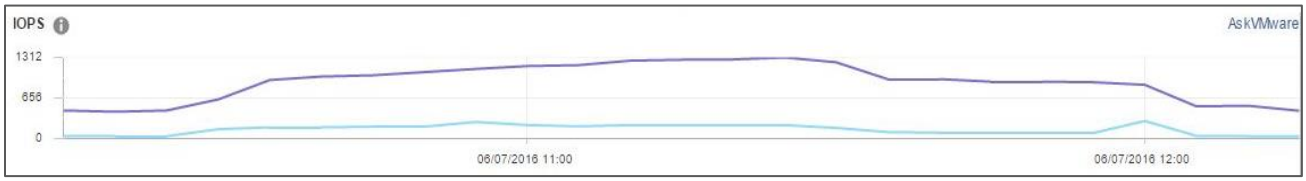

Memory usage is monitored on the ESXi host, memory usage metrics monitored are consumed, active, balloon and swap used, as swap and ballooning usage would indicate host memory reached saturation point and the VM performance may start to deteriorating. All tests were carried out on hosts with 512 GB of physical memory installed.

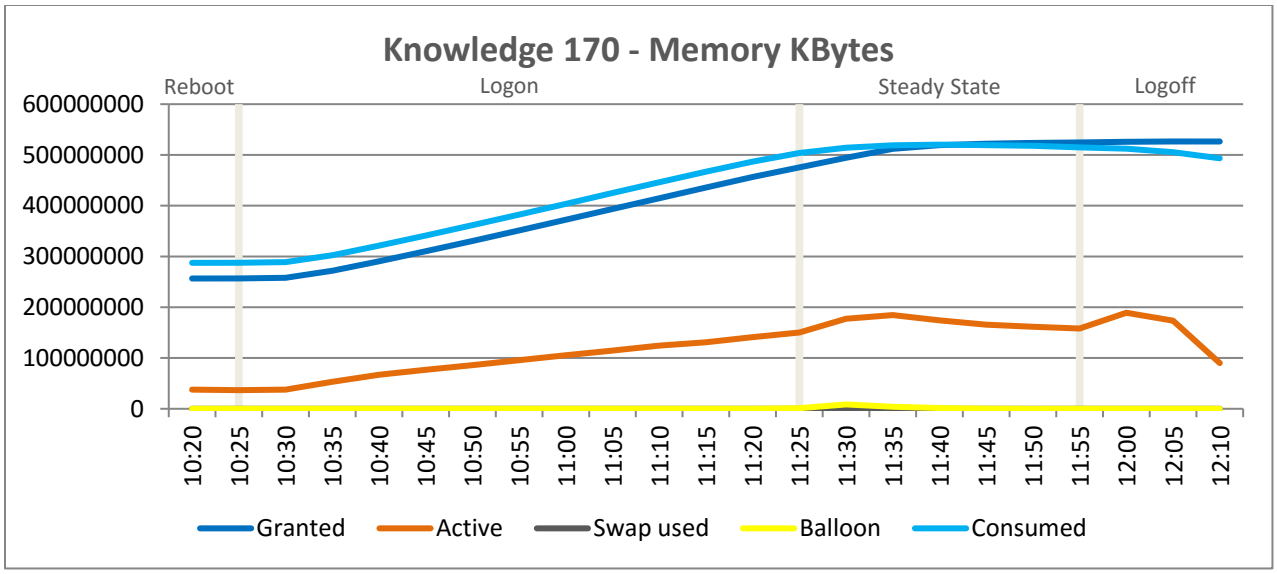

There were no issues with network usage on any of the test runs. There is a significant reduction in activity once the steady state phase is reached after all machines have logged on.

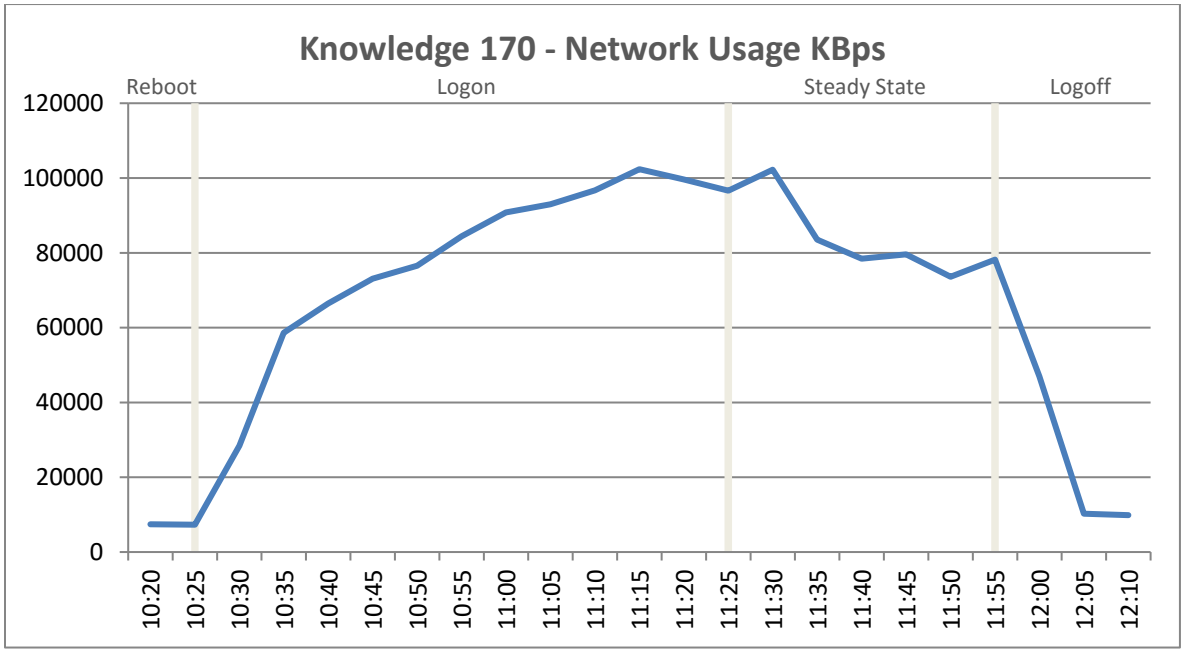

6.3.2.3 Power Worker, 100 users, ESXi 6.0U2, Horizon 7 linked clones

The following graph show the output from the Login VSI Analyzer for each V470-C7 test run. VSI Max was not reached reached

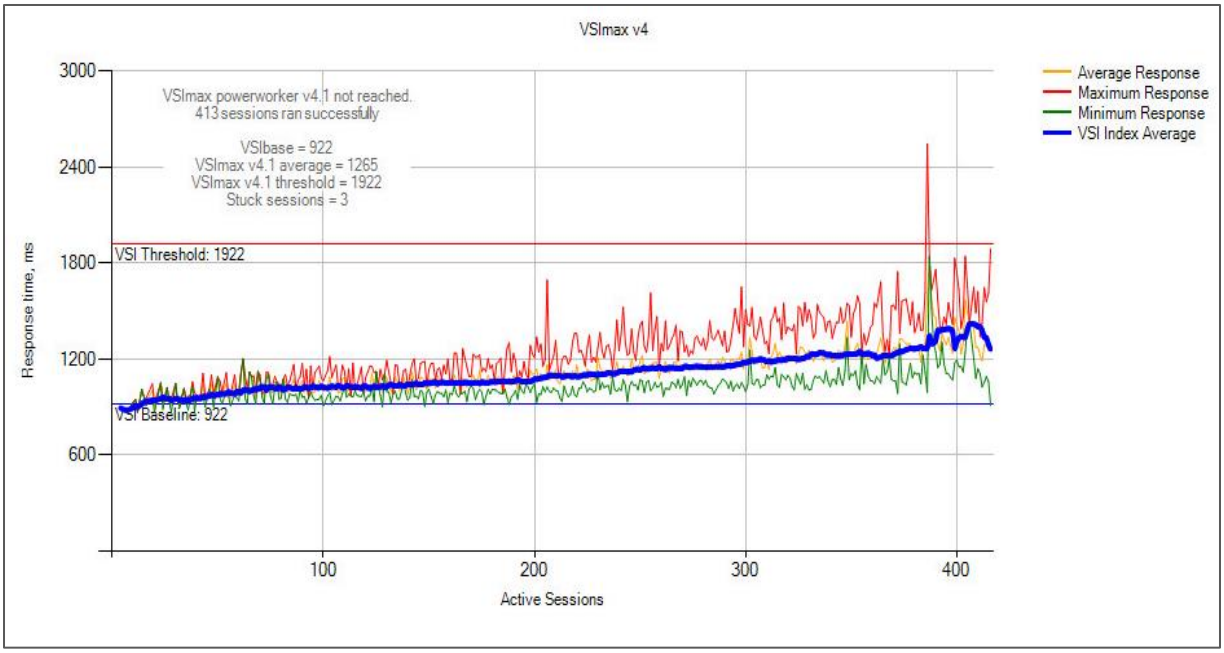

Maximum CPU utilization was in the region of the 85% threshold indicating the number of users tested was appropriate.

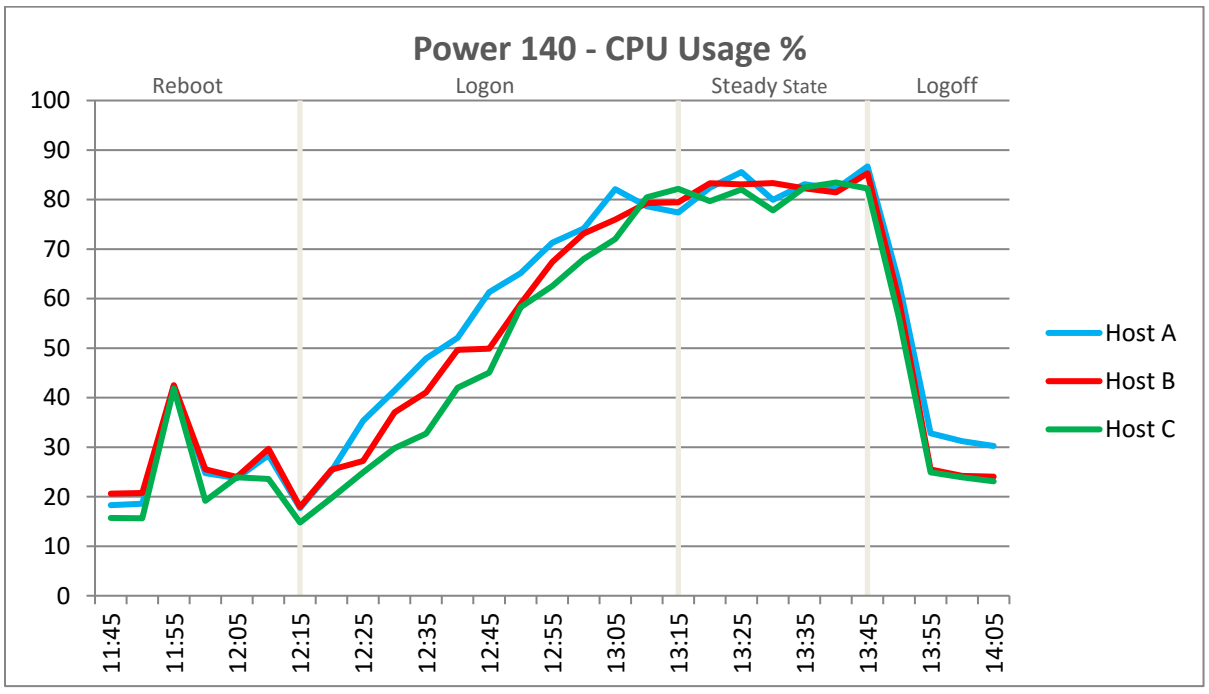

Latency on the datastore spiked temporarily during the boot phase of the test but quickly settled once all the VMs were booted. For the logon and steady state phase of each test, the latency remained well below the 20ms threshold reaching a max of 2-3 ms during the test run.

The IOPS peaked during the boot phase and for each profile test and then settled thereafter during the login phase and reduced once steady state was reached.

This chart was captured from within vSphere and was a feature released with VSAN6.2 so we do not need to use VSAN Observer as was previously the case with past VSAN validations.

The statistics below are on a per host basis so as VSAN scales linearly; to calculate the total IOPs for a three node cluster you would multiple by three.

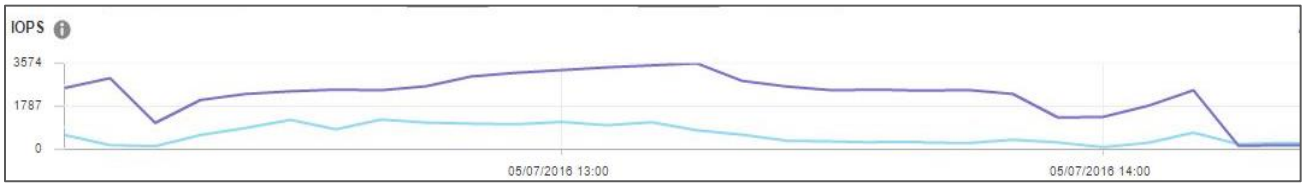

Memory Utilization

Memory usage is monitored on the ESXi host, memory usage metrics monitored are consumed, active, balloon and swap used, as swap and ballooning usage would indicate host memory reached saturation point and the VM performance may start to deteriorating. All tests were carried out on hosts with 512 GB of physical memory installed.

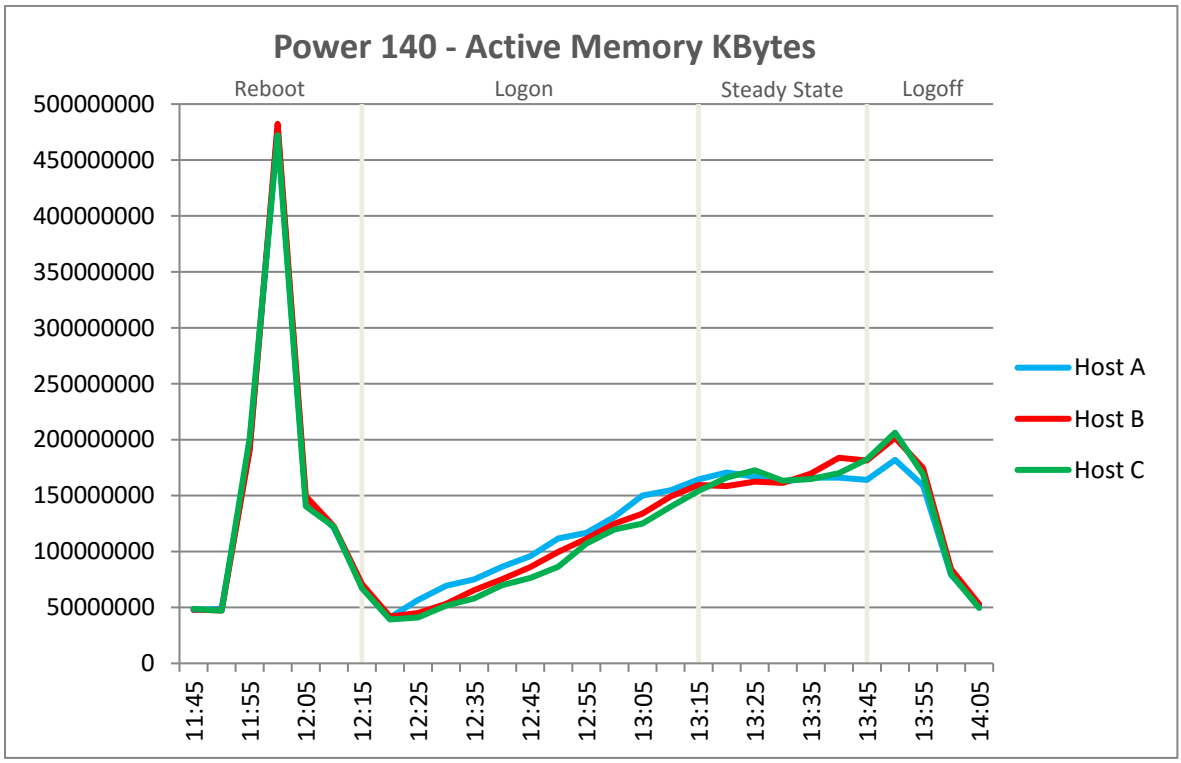

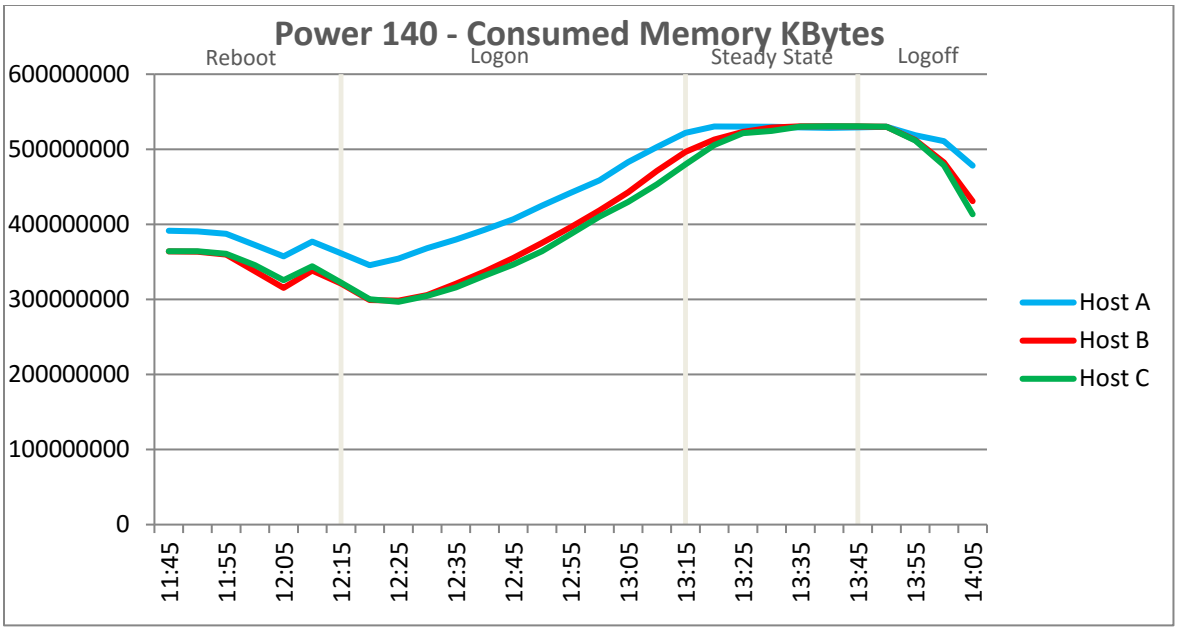

There were no issues with network usage on any of the test runs. There is a significant reduction in activity once the steady state phase is reached after all machines have logged on.

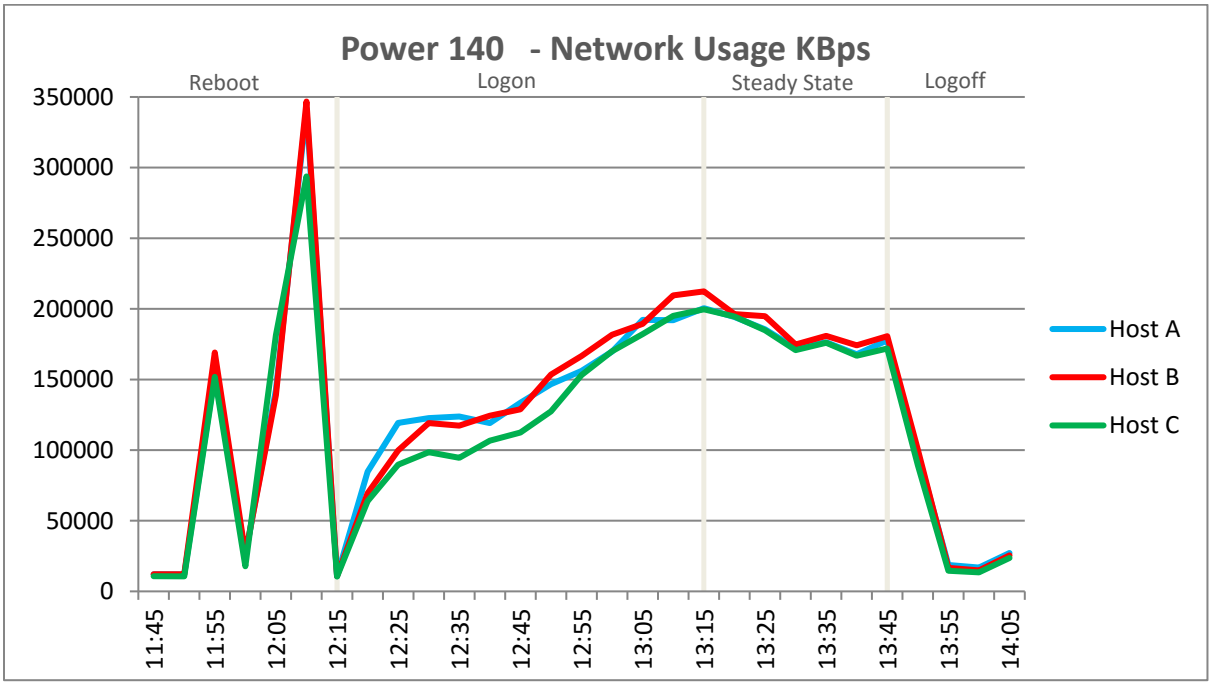

6.3.2.4 V470-C7-GPU-M60-1Q

The following testing was completed on the V470-C7 configuration with 2 x M60 GPU cards added to the configuration. The Knowledge Worker workload was used with the M60-1Q GPU profile and with a 2 x M60 configuration the maximum amount of VMs is 32. We are currently unable to test with the maximum supported display resolution for this profile which is 4096x2160 so 2560x1600

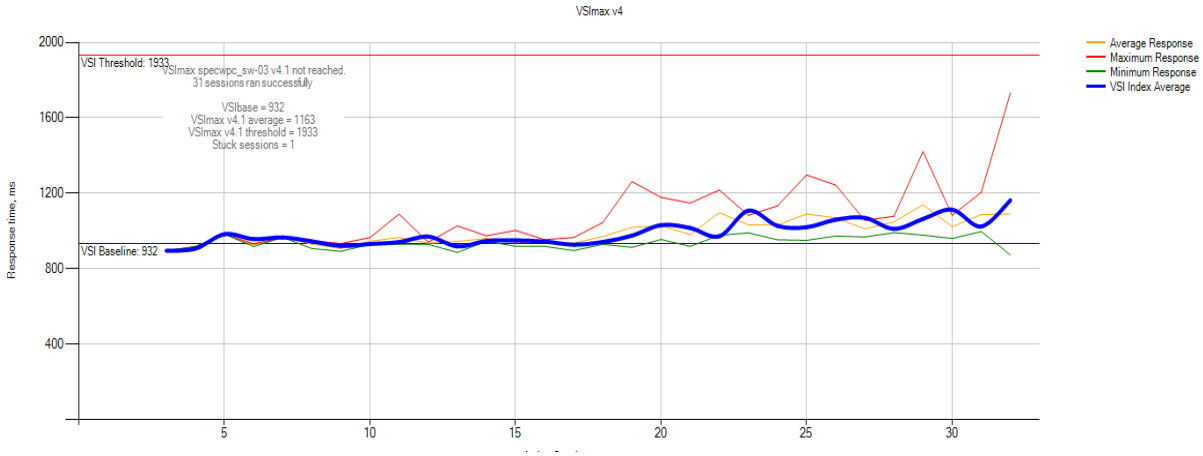

Maximum CPU utilization was in the region of the 85% threshold indicating the number of users tested was appropriate.

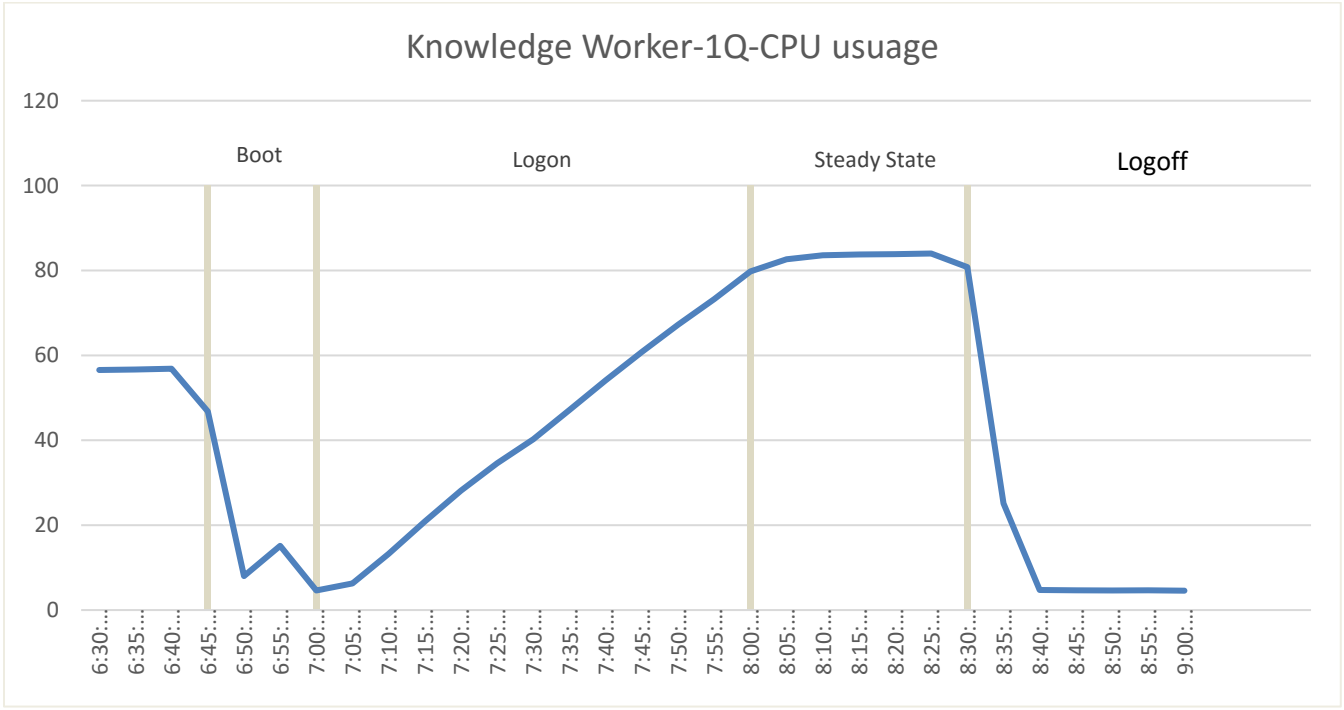

Memory usage is monitored on the ESXi host, memory usage metrics monitored are consumed, active, balloon and swap used, as swap and ballooning usage would indicate host memory reached saturation point and the VM performance may start to deteriorating. All tests were carried out on hosts with 512 GB of physical memory installed.

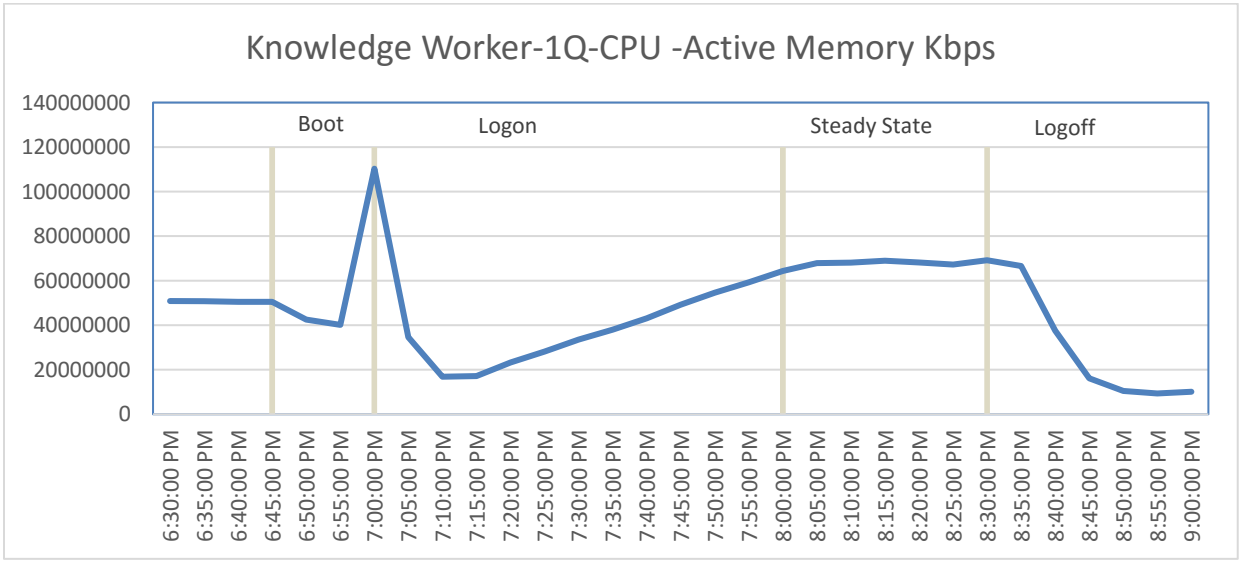

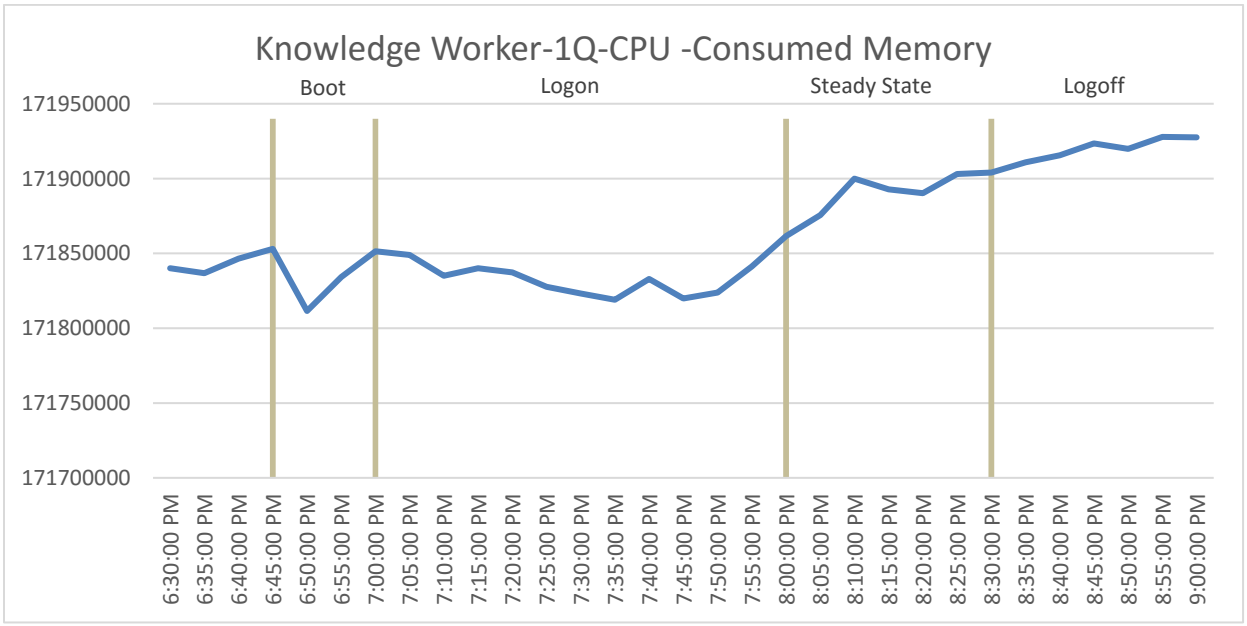

There were no issues with network usage on any of the test runs. There is a significant reduction in activity once the steady state phase is reached after all machines have logged on.

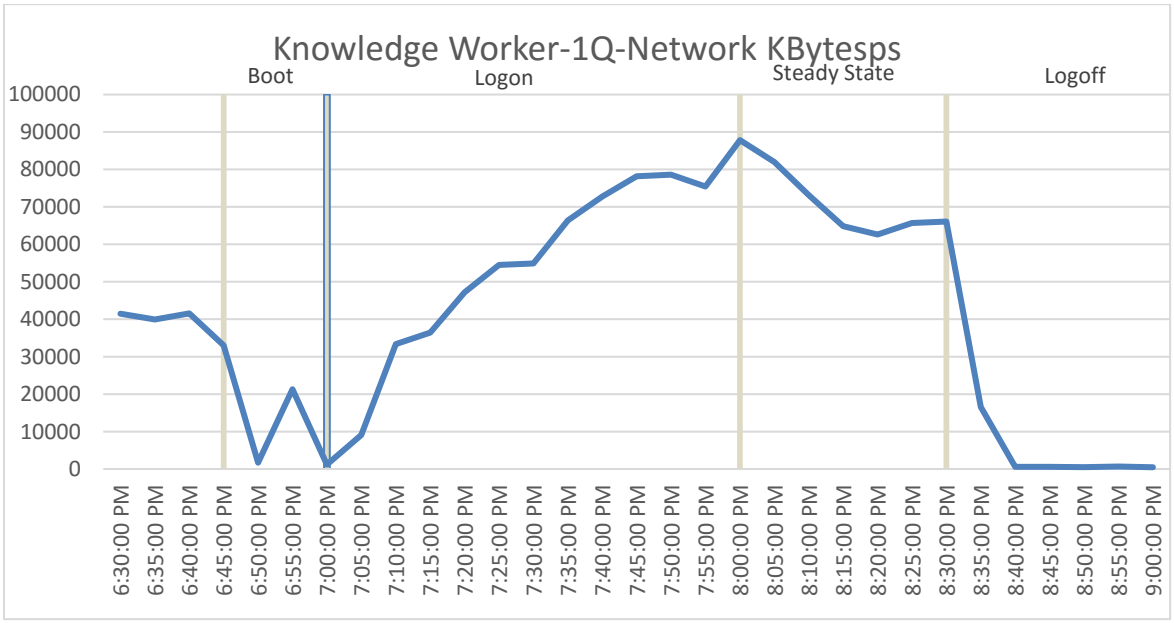

Datastore IOPS

Latency on the datastore spiked temporarily during the boot phase of the test but quickly settled once all the VMs were booted. For the logon and steady state phase of each test, the latency remained well below the 20ms threshold reaching a max of 2-3 ms during the test run.

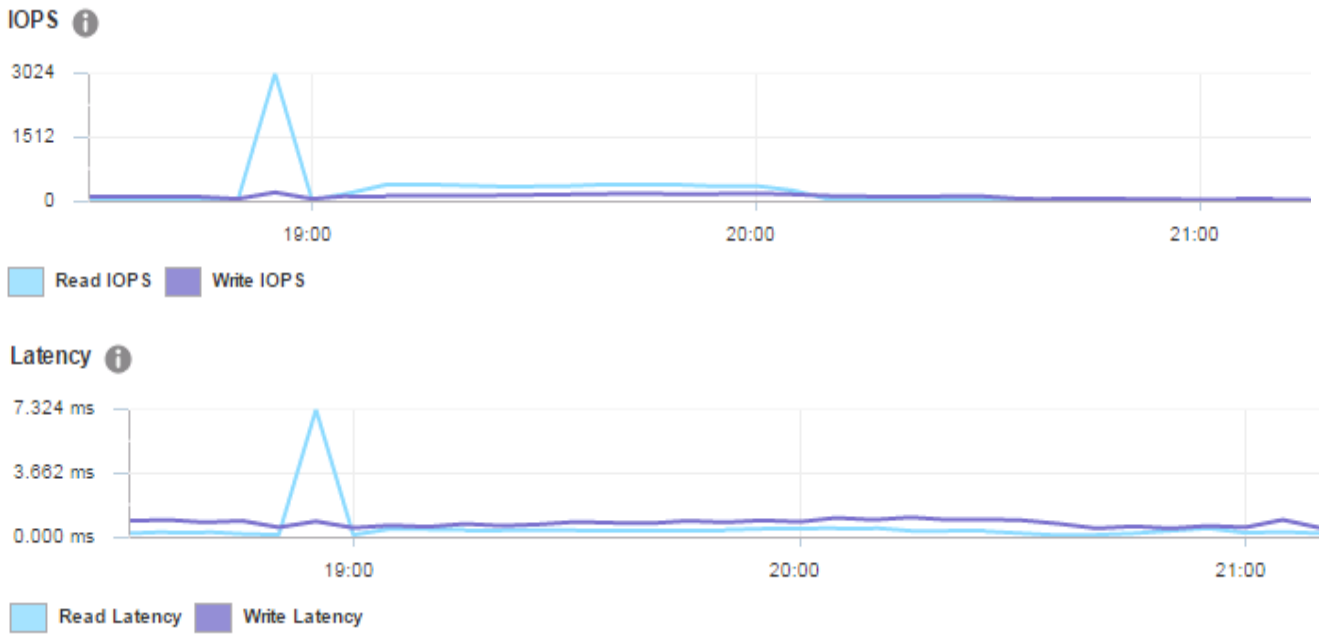

6.3.2.5 V470-C7, RDSH

LoginVSI VSI Max Result

VSI Max was not reached on the test run indicating there was no decrease in user experience.

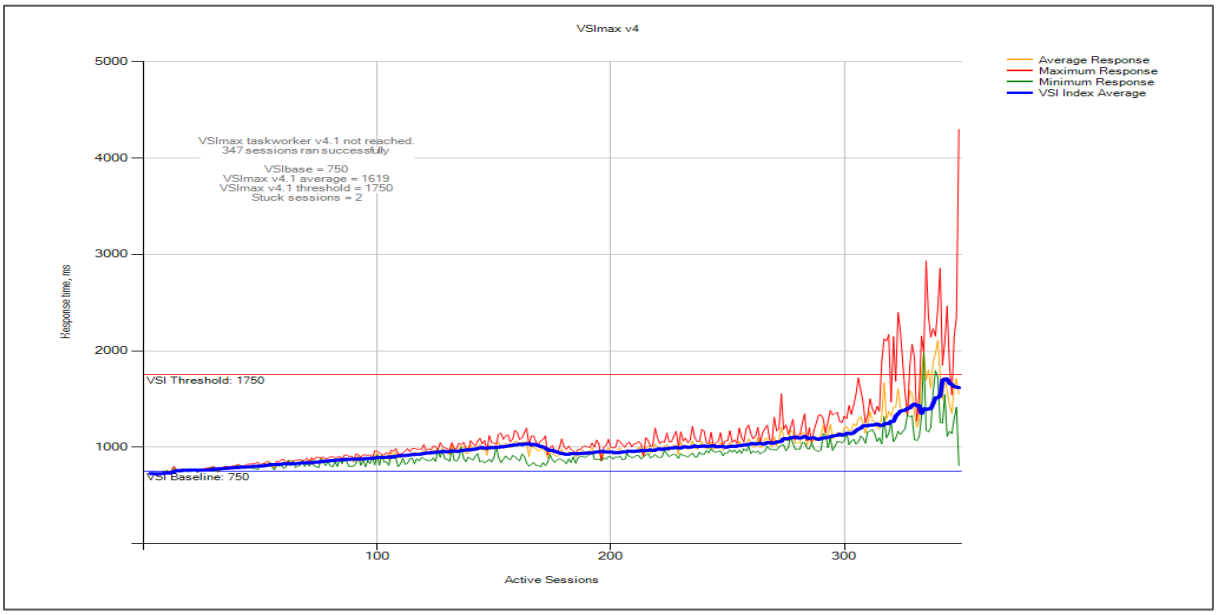

The host CPU was pushed to a 92% steady state average during this test run and the two sampled RDSH VMs to approximately 80% each. The data suggests that 350 Task Worker users are at the upper end of the capabilities of this config.

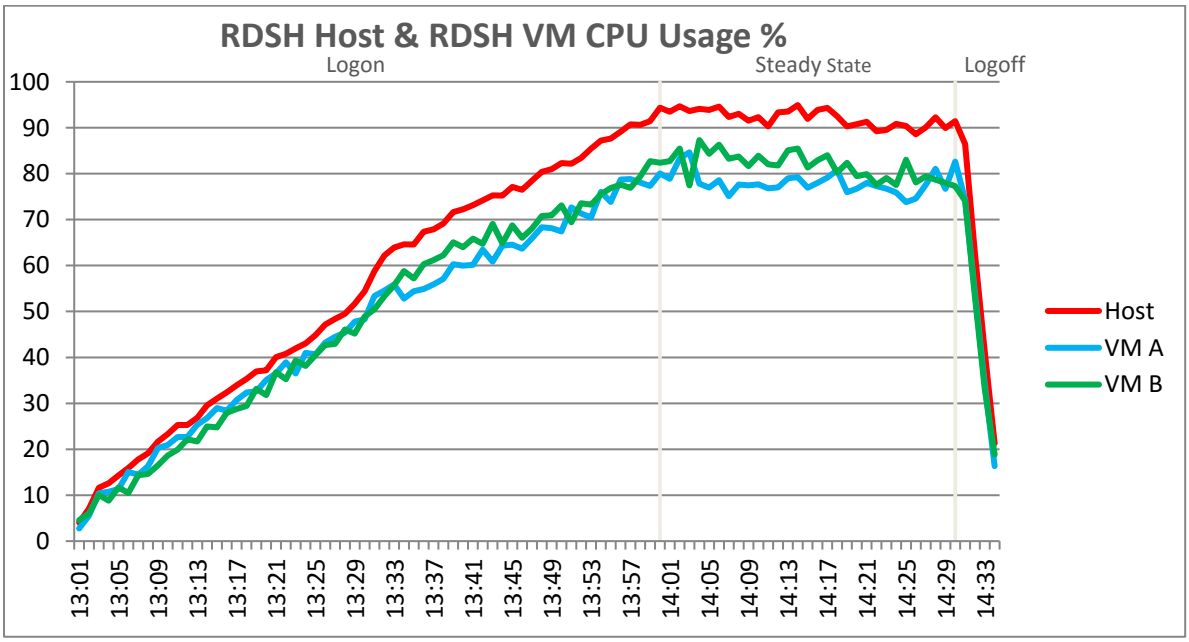

Latency on the datastore did not approach the 20ms threshold. There were no latency spikes during the logon or steady states phases of the test run with 4.5 ms the maximum value reached and 6.4 ms the maximum during the logoff phase.

As there was no reboot spike during RDSH testing, IOPS increased steadily as more users logged on reaching a peak just as the steady state phase began. At peak, IOPS reached approximately 2000 resulting in approximately 5.7 IOPS per user.

These charts are captured from within vSphere and were a feature released with VSAN6.2 so we do not need to use VSAN Observer as was previously the case with past VSAN validations.

The statistics below are on a per host bases so as VSAN scales linearly to calculate the total IOPS for a three node cluster you would multiple by three.

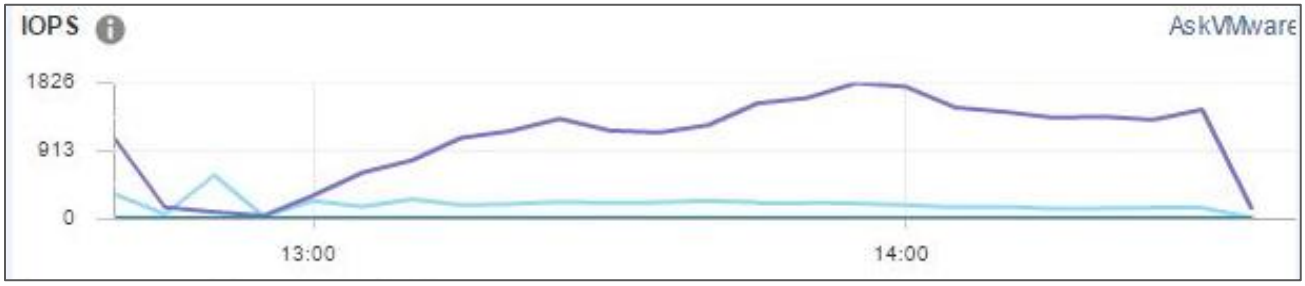

Memory Utilization

Memory usage is monitored on the ESXi host and the two sampled RDSH VMs. There was no ballooning on the physical host or the two samples RDSH VMs. No swapping took place on the physical host.

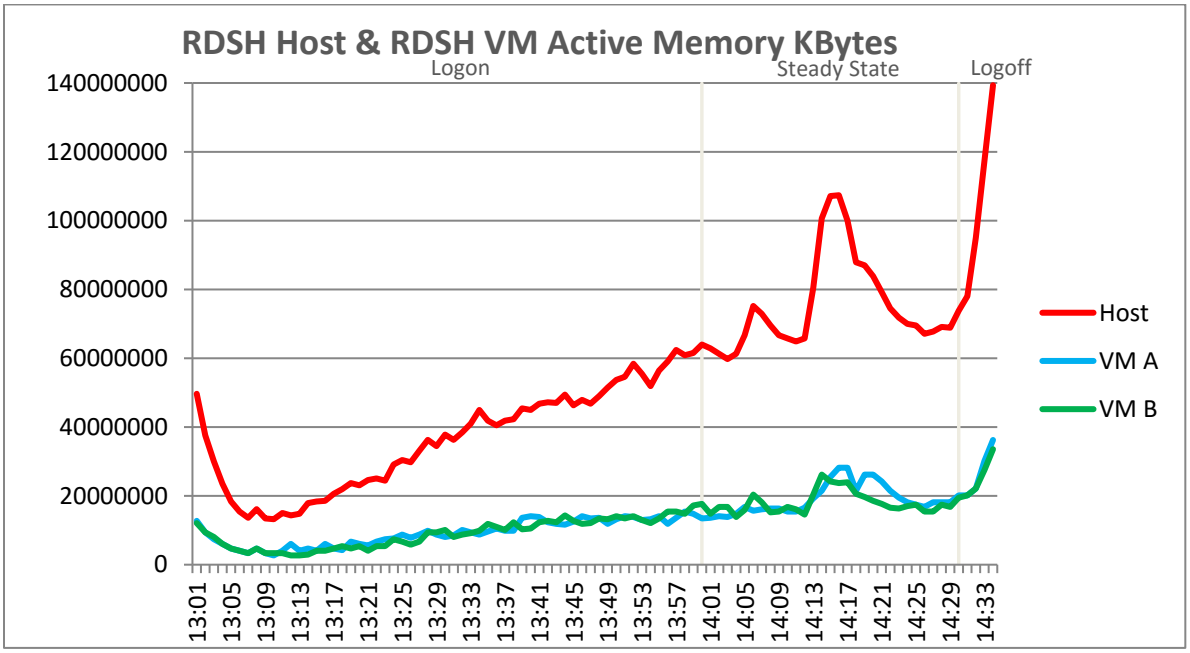

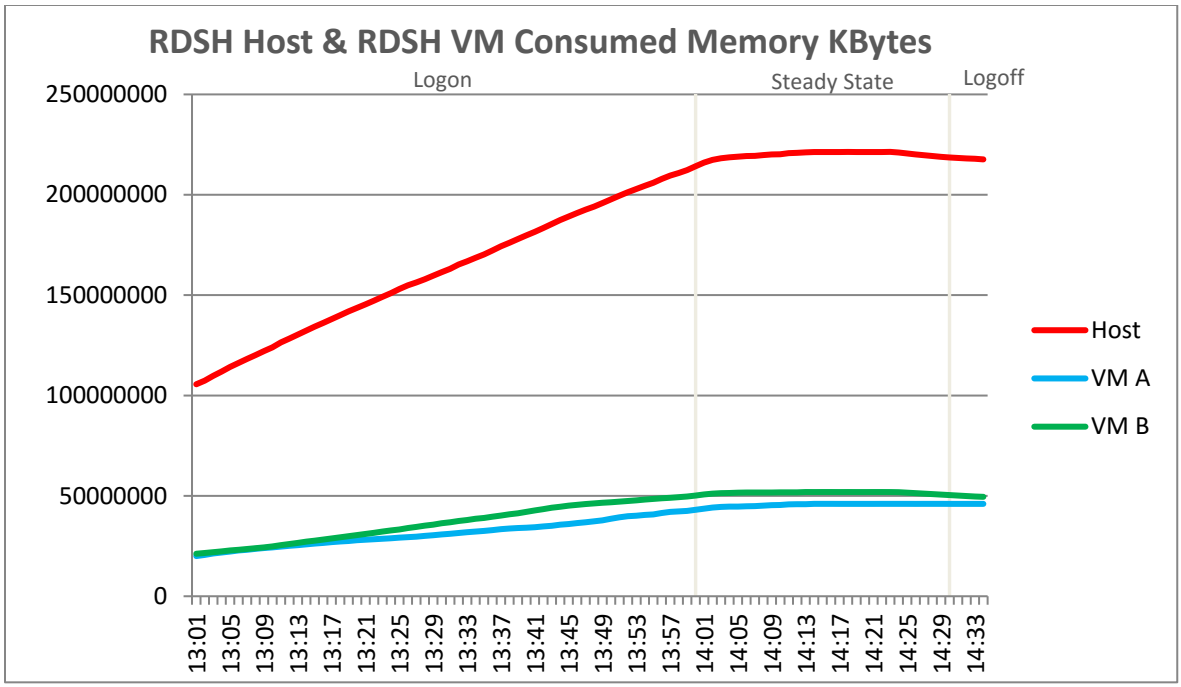

Network utilization was not an issue in this test with host usage reaching a maximum of approximately 65,000 KBytesps.

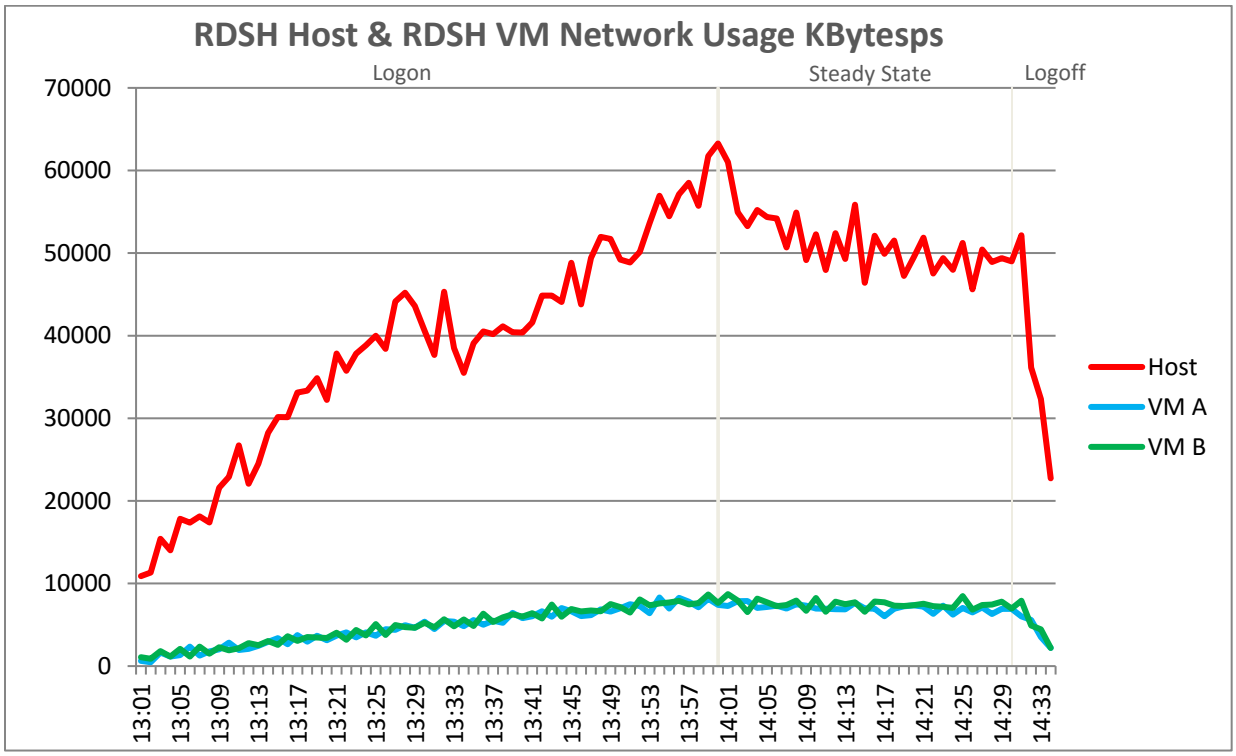

Acknowledgements

Thanks to David Hulama of the Wyse Technical Marketing team for his support and assistance with VMware data center EUC programs at Dell. David is a Senior Technical Marketing Advisor for VMware VDI solutions at Dell. David has a broad technical background in a variety of technical areas and expertise in enterprise-class virtualization solutions.

Thanks to Mike Hayes from the Limerick CSC team for his help and support with the Graphics Functionality Testing that was completed on VxRail. Mike is a Solutions Architect working at the Dell Customer Solution Center in Limerick, Ireland. Responsible for Client Solutions and VDI engagements at the Center in EMEA, Mike has a strong background in Desktop and Server Virtualization with over 15 years' experience working in Enterprise class IT environments. Highly skilled in Microsoft, VMware and Citrix platforms, Mike primarily works on design workshop and Proof-Of-Concept activity around VDI and high performance graphics, including Workstation and VR Technology. Twitter:@MikeJAtDell

Thanks to Kevin Corey from the Limerick CSC team for his help and support with the network setup for this validation. Kevin is a Network Solution Architect with over 17 years' experience in working with enterprise environments. Primarily focusing on data center networking, Kevin has experience working with technology from all major network vendors.

Thanks to Gus Chavira for his continued guidance in support for this program, Gus is the Dell CCC Alliance Director to VMware. Gus has worked in capacities of Sys Admin, DBA, Network and Storage Admin, Virtualization Practice Architect, Enterprise and Solutions Architect. In addition, Gus carries a B.S. in Computer Science.

Thanks to Andrew Mc Daniel for his support during this program, Andrew is the CTO/ Strategy Director with CCC- who is responsible for managing team that is responsible for examining new technologies and research projects to evaluate potential benefit of internal and external partners' hardware and software to Dell's E2E solutions for EUC and their strategic integration.

Thanks to Rick Biedler for his support during this program, Rick is the Engineering Director for Datacenter Appliances at Dell, managing the development and delivery of enterprise class desktop virtualization solutions based on Dell Datacenter components and core virtualization platforms.

About the Authors

Keith Keogh is the Lead Architect for VMware solutions in the Cloud Client Engineering Team at Dell. Keith has extensive experience with the Design, Implementation and Validation of EUC Enterprise Solutions.

Peter Fine is the Chief Architect and CTO of EUC Enterprise Engineering at Dell. Peter owns the strategy, architecture and leads the engineering of the datacenter EUC product and solutions portfolio. Peter also has extensive experience building and managing data centers for Fortune500 and financial services companies. Follow Peter @ExitTheFastLane or www.ExitTheFastLane.com.