

# Dell EMC VMware vSAN Ready Nodes for VMware Horizon – Reference Architecture

Integration of VMware Horizon with vSAN Ready Node appliance clusters

Dell Engineering  
September 2017

## Revisions

Date	Description
June 2016	Initial release
July 2016	Updated with HY-8 R730, Thin Clients, added VxRail, NSX, GPU, RDSH and Horizon Air Hybrid-Mode
August 2016	All-Flash configuration added
March 2017	FX2 platform and SATADOM boot option for supported configurations. NVIDIA M10 support
April 2017	Section 6 Test results update
June 2017	Disk configuration Alignment, FC630 vSAN HCL addition
September 2017	Horizon Apps section added and Thin Clients section updated

THIS WHITE PAPER IS FOR INFORMATIONAL PURPOSES ONLY, AND MAY CONTAIN TYPOGRAPHICAL ERRORS AND TECHNICAL INACCURACIES. THE CONTENT IS PROVIDED AS IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND.

Copyright © 2016-2017 Dell Inc. All rights reserved. Dell and the Dell logo are trademarks of Dell Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

# Table of contents

Revisions.....	2
1 Introduction.....	6
1.1 Purpose .....	6
1.2 Scope.....	6
1.3 What's new .....	6
2 Solution architecture overview .....	7
2.1 Introduction .....	7
2.2 What is a vSAN Ready Node (vSRN)? .....	7
2.3 Physical architecture overview .....	8
2.4 Solution layers .....	8
2.4.1 Networking.....	8
2.4.2 Management and Compute Servers.....	10
2.4.3 Software Defined Storage (vSAN).....	11
3 Hardware components .....	13
3.1 Network.....	13
3.1.1 Dell Networking S3048 (1Gb ToR switch).....	13
3.1.2 Dell Networking S4048 (10Gb ToR switch).....	14
3.1.3 FN410S.....	14
3.2 Dell EMC VDI Optimized vSAN Ready Nodes .....	15
3.2.1 vSRN Server Configuration .....	19
3.2.2 vSRN R630.....	20
3.2.3 vSRN R730.....	24
3.2.4 vSRN R730XD.....	27
3.2.5 vSRN FC430.....	30
3.2.6 vSAN FC630.....	31
3.2.7 Summary comparison.....	33
3.3 GPUs .....	34
3.3.1 NVIDIA Tesla GPUs .....	34
3.4 Dell Wyse Thin Clients .....	36
3.4.1 Wyse 3040 Thin Client (ThinOS, ThinLinux).....	36
3.4.2 Wyse 5030 PCoIP Zero Client.....	36
3.4.3 Wyse 5040 AIO Thin Client with PCoIP .....	36

3.4.4	Wyse 5050 AIO PCoIP Zero Client .....	37
3.4.5	Wyse 7030 PCoIP Zero Client.....	37
3.4.6	Wyse 5060 Thin Client (ThinOS, ThinLinux, WES7P, WIE10) .....	37
3.4.7	Wyse 7020 Thin Client (Windows 10 IoT) .....	38
3.4.8	Wyse 7040 Thin Client with Windows Embedded Standard 7P.....	38
3.4.9	Latitude 3480 and 5280 Mobile Thin Clients (Win 10 IoT) .....	38
4	Software Components.....	40
4.1	VMware.....	40
4.1.1	VMware vSphere 6.x .....	40
4.1.2	VMware vSAN .....	41
4.1.3	vSAN Best Practices .....	42
4.1.4	VMware Horizon .....	45
4.1.5	VMware Horizon Apps.....	48
4.1.6	NUMA architecture considerations .....	49
4.2	NVIDIA GRID vGPU .....	51
4.2.1	vGPU Profiles .....	52
5	Solution architecture for vSRN with Horizon .....	59
5.1	Management Server Infrastructure.....	59
5.1.1	RDSH VM Configuration.....	59
5.1.2	NVIDIA GRID License Server Requirements .....	60
5.1.3	SQL Databases .....	61
5.1.4	DNS .....	61
5.2	Storage Architecture Overview.....	62
5.2.1	vSAN.....	62
5.3	Virtual Networking.....	62
5.3.1	VSRN Network Configuration .....	62
5.3.2	VMware NSX .....	65
5.4	Scaling Guidance.....	67
5.5	Solution High Availability .....	68
5.5.1	vSAN HA/ FTT configuration .....	69
5.5.2	vSphere HA .....	69
5.5.3	Management Server High Availability.....	70
5.5.4	Horizon HCS High Availability .....	70

5.5.5 SQL Server High Availability .....	71
5.6 VMware Horizon Communication Flow .....	71
6 Solution Performance and Testing .....	72
6.1 Test and Performance Analysis Methodology .....	73
6.1.1 Testing Process .....	73
6.1.2 Resource Monitoring .....	76
6.1.3 Resource Utilization .....	76
6.2 Test Configuration Details .....	77
6.2.1 Compute VM configurations .....	78
6.2.2 Platform Configuration .....	79
6.3 Test Results and Analysis .....	79
6.3.1 vSRN-B5 .....	81
6.3.2 vSRN-C7 .....	93
6.3.3 vSRN-FC430 B5 .....	112
6.3.4 vSAN-FC630 B5 .....	126
Acknowledgements .....	130
About the Authors .....	131

# 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 VMware vSAN Ready Nodes (vSRN) with vSphere 6.5 and VMware vSAN 6.5.

## 1.2 Scope

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

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

## 1.3 What's new

- Introduce VMware vSAN Ready Nodes
- Introduce Hybrid configuration for vSAN
- Introduce vSAN new features
- Introduce All-Flash configuration for vSAN
- Introduce FX2 FC430 configuration
- Introduce FX2 FC630 vSAN configuration

## 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 (application/ session virtualization)

### 2.2 What is a vSAN Ready Node (vSRN)?

A vSRN is a validated Dell Server configuration in a tested, verified Hardware form factor for vSAN deployments, jointly recommended by Dell and VMware. This makes the process easier for the customer and from the vSRN compatibility page link they can select any of the Dell Hybrid and All-Flash configurations depending on their requirements. The vSRN Nodes also comes with pre-selected Horizon OEM VMware SKUs for licensing and support of Horizon on vSAN.

Home > Resources > Compatibility Guides

## VMware Compatibility Guide

Search Compatibility Guide:  All Listings

Looking for a simplified search? [Use the Guided Search Wizard](#)

What are you looking for:  Compatibility Guides  Current Results:

**Need Help? Try out the [Virtual SAN Ready Node Configurator](#).**

**STEP 1: Refer to the "[Virtual SAN Hardware Quick Reference Guide](#)" for guidance on how to build a Virtual SAN Ready Node.**

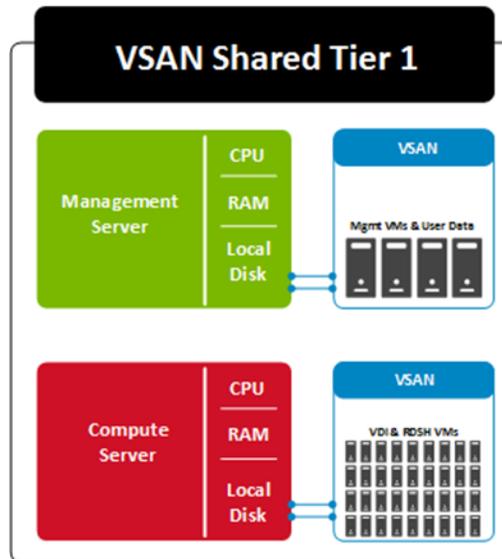
**STEP 2: To build a Virtual SAN Ready Node:**

Select your Virtual SAN Ready Node of choice based on following certified Ready Nodes.

<b>Ready Node Types:</b> All	<b>Ready Node Vendors:</b> All Cisco <b>DELL</b> Fujitsu Hewlett Packard Enterprise Hitachi	<b>Ready Node Profile:</b> All HY-2 Series HY-4 Series HY-6 Series HY-8 Series AF-6 Series
<b>Ready Node Supported Releases:</b> All ESXi 6.0 U2 ESXi 6.0 U1 ESXi 6.0 ESXi 5.5 U3	<b>Ready Node Generation:</b> All Gen1 - 6G Gen2 - 12G	<b>Ready Node Server Type:</b> All Blade Rackmount
<b>Pre-Install Options:</b> ESXi Pre-Installed ESXi Not Pre-Installed	<b>Posted Date Range:</b> All	<b>Raw Storage Capacity (TB):</b> All
<b>Keyword:</b>		

## 2.3 Physical architecture overview

The core vSRN architecture consists of a software-defined Shared Tier1 model. This consists of a Cache and a Capacity Tier, the minimum requirements for which are 1 x SSD for the Cache Tier and 1 xHDD or SSD for the Capacity Tier. The management and compute nodes are configured in the same vSRN Cluster and share the vSAN datastore. The user data can be hosted via a file server residing within the VSAN file system.



## 2.4 Solution layers

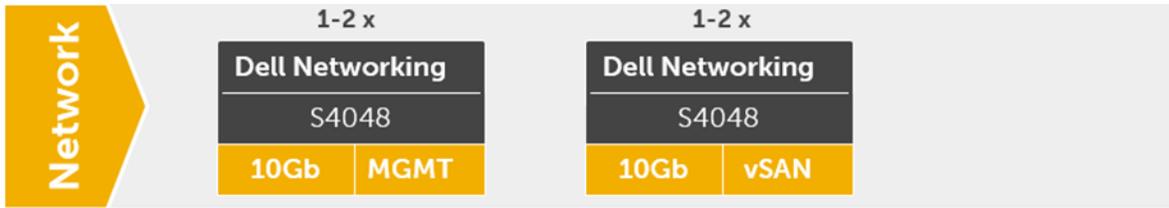
The vSRN Solution leverages a core set of hardware and software components consisting of five primary layers:

- Networking Layer
- Management and Compute Server Layer
- Storage Layer (vSAN)
- Thin Client Layer (please refer to section 3.4)

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

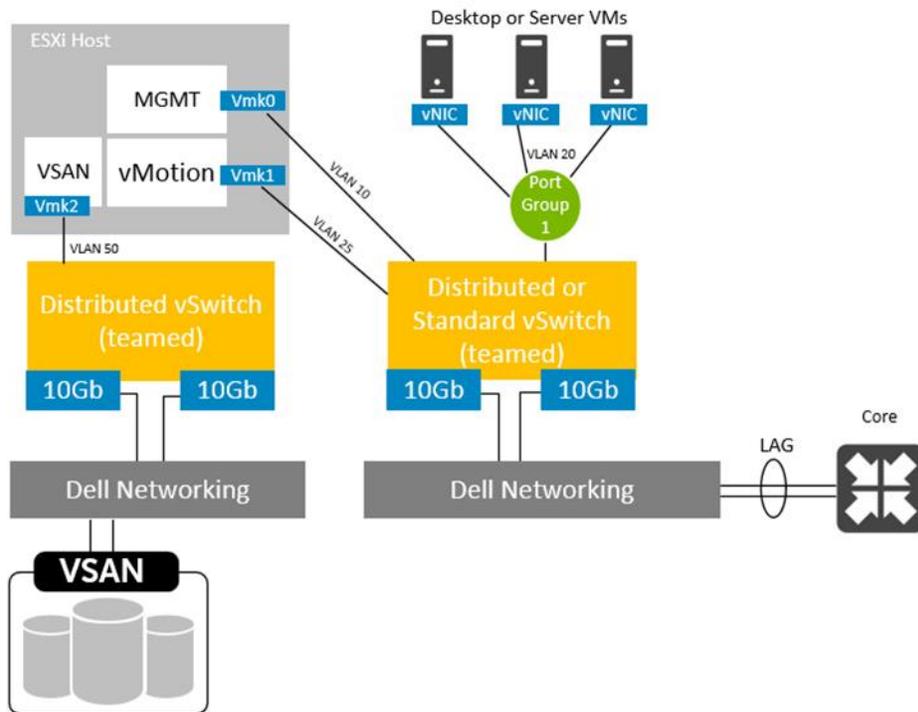
### 2.4.1 Networking

Dell recommends 10Gb networking be used with vSAN traffic separated into discrete Switching Fabrics for HA. Management traffic can be optionally separated physically or converged via VLANs with all other traffic leveraging the same interfaces.



Additional switches are added and stacked as required to provide High Availability for the Network layer. When there is a requirement for 1Gb connectivity for DRAC/remote management, we can use an existing 1Gb ToR or add a Dell Networking 1Gb 48-port switch for this function.

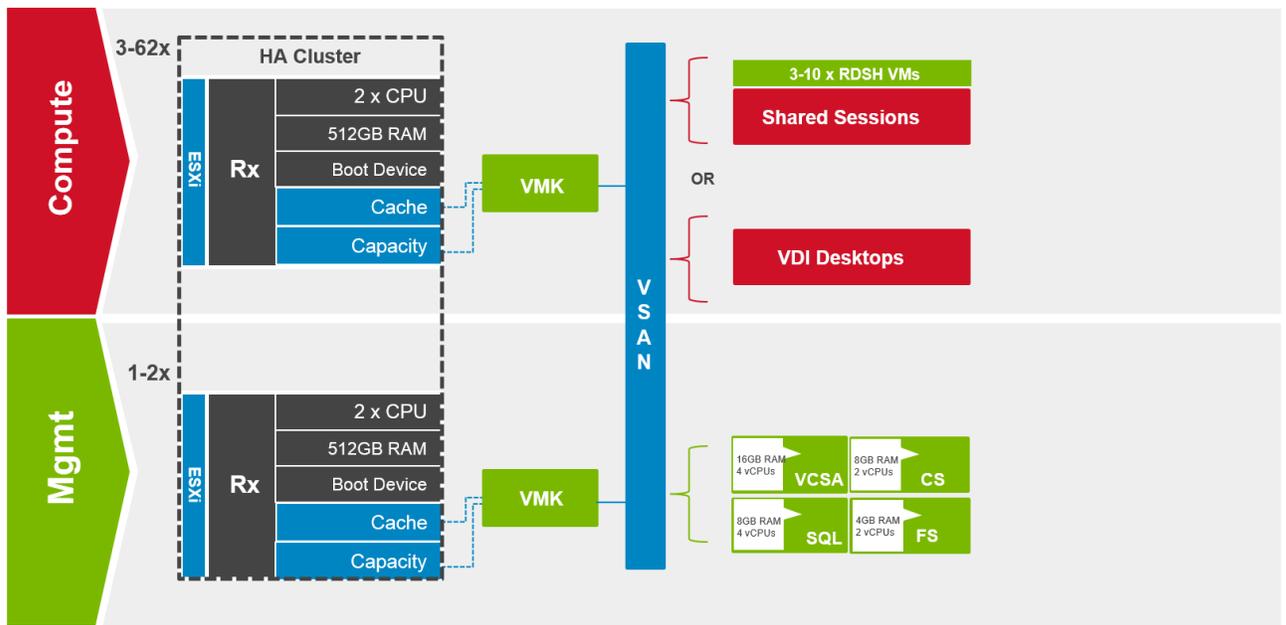
In the Shared Tier 1 architecture, a single Dell Networking switch is shared among all network connections for both management and compute, the management and vSAN traffic are separated out via VLAN across 2 x 10Gb NICs. When deploying larger cluster configurations it may be optimal to split out the management and vSAN traffic. The Server configurations are equipped with 4 x 10Gb connections to facilitate this. All Top of Rack (ToR) traffic is layer 2 (L2)/ switched locally, with all layer 3 (L3)/ routable VLANs trunked from a core or distribution switch. The following diagrams illustrate the logical data flow in relation to the core switch.



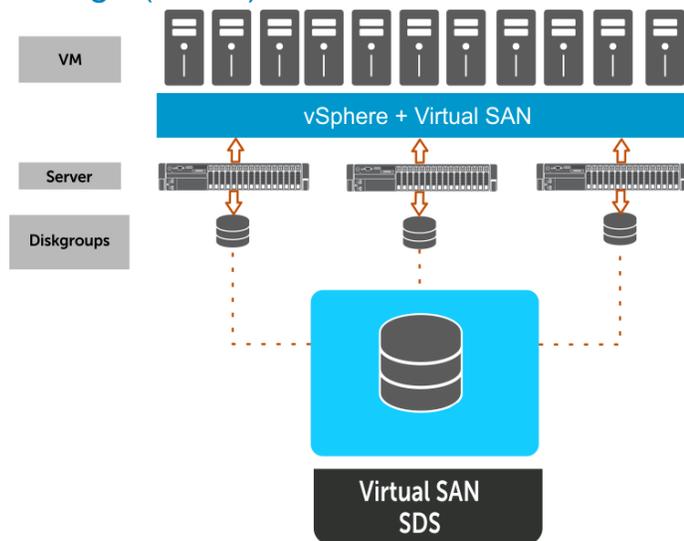
## 2.4.2 Management and Compute Servers

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

Dell recommends that the VDI management infrastructure nodes be physically separated from the compute resources. In this configuration both management and compute exist in the same vSAN Cluster but the management node is reserved for management server VMs only and this will be expanded as needed depending on the size of the cluster. Compute hosts can be used interchangeably for Horizon or RDSH as required.



## 2.4.3 Software Defined Storage (vSAN)



VMware vSAN is software-defined storage solution fully integrated into vSphere. Once enabled on a cluster, all the disk devices present in these hosts are pooled together to create a shared data store that will be accessible by all hosts in the 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 vSAN 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 known as the FailureToTolerate (FTT); 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.

HDDs play two roles in a hybrid vSAN: they make up the capacity of the vSAN data store as well as making up components for a stripe width. SAS, NL-SAS and SATA drives 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 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.x 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.

There are two ways to build a vSAN cluster, build your custom configuration using the HCL [link](#) or choose the Dell VMware Certified Ready Nodes [link](#).

## 3 Hardware components

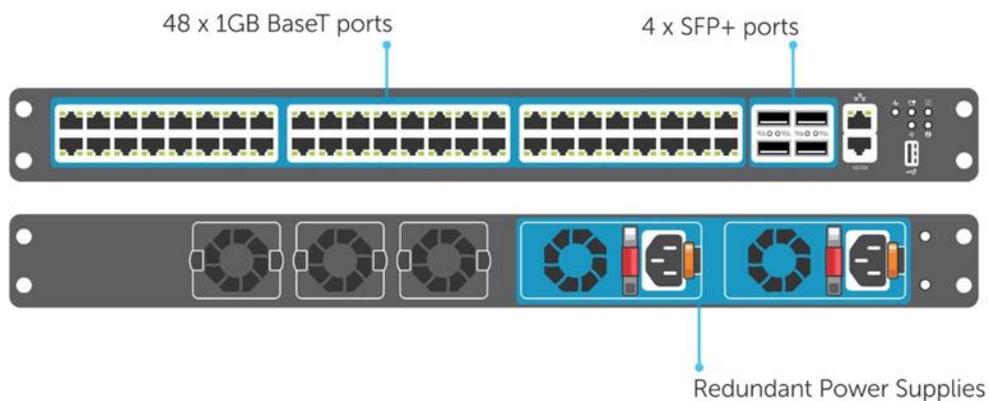
### 3.1 Network

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

#### 3.1.1 Dell Networking S3048 (1Gb ToR switch)

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

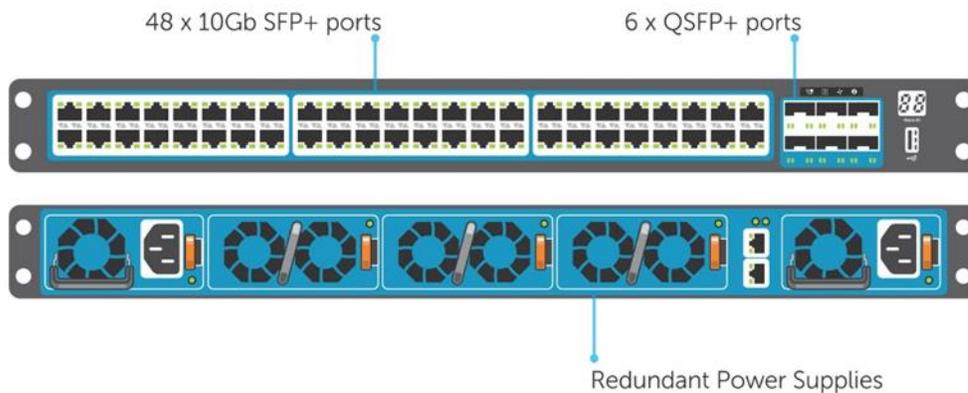
Model	Features	Options	Uses
Dell Networking S3048-ON	48 x 1000BaseT 4 x 10Gb SFP+	Redundant hot-swap PSUs & fans	1Gb connectivity
	Non-blocking, line-rate performance	VRF-lite, Routed VLT, VLT Proxy Gateway	
	260Gbps full-duplex bandwidth	User port stacking (up to 6 switches)	
	131 Mpps forwarding rate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



### 3.1.2 Dell Networking S4048 (10Gb ToR switch)

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

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



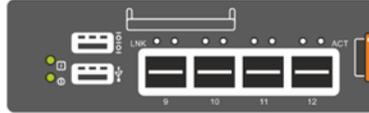
For more information on the S3048, S4048 switches and Dell Networking, please visit: [link](#)

### 3.1.3 FN410S

The FN IO Module (FN IOM) is designed specifically for the PowerEdge FX2 converged-infrastructure chassis, part of the PowerEdge FX architecture. Supporting up to two FN IOMs per chassis, the FX2 converged infrastructure also includes up to 8 x 10GbE internal ports, plus redundant cooling fans and power supplies.

The FN IOM can simplify FX2 connectivity by as much as 8-to-1, greatly reducing cabling complexity. Plug-and-play convenience with zero-touch and preconfigured features enables layer 2 control. This allows fast, easy network deployment using the Dell Blade IO Manager, with additional management of the FN IOM through a CLI using the Dell Networking OS 9, enabling full layer 3 switching functions.

Model	Features	Options	Uses
PowerEdge FN410S	4 x external SFP+	4-port SFP+ 10Gb module	10Gb connectivity



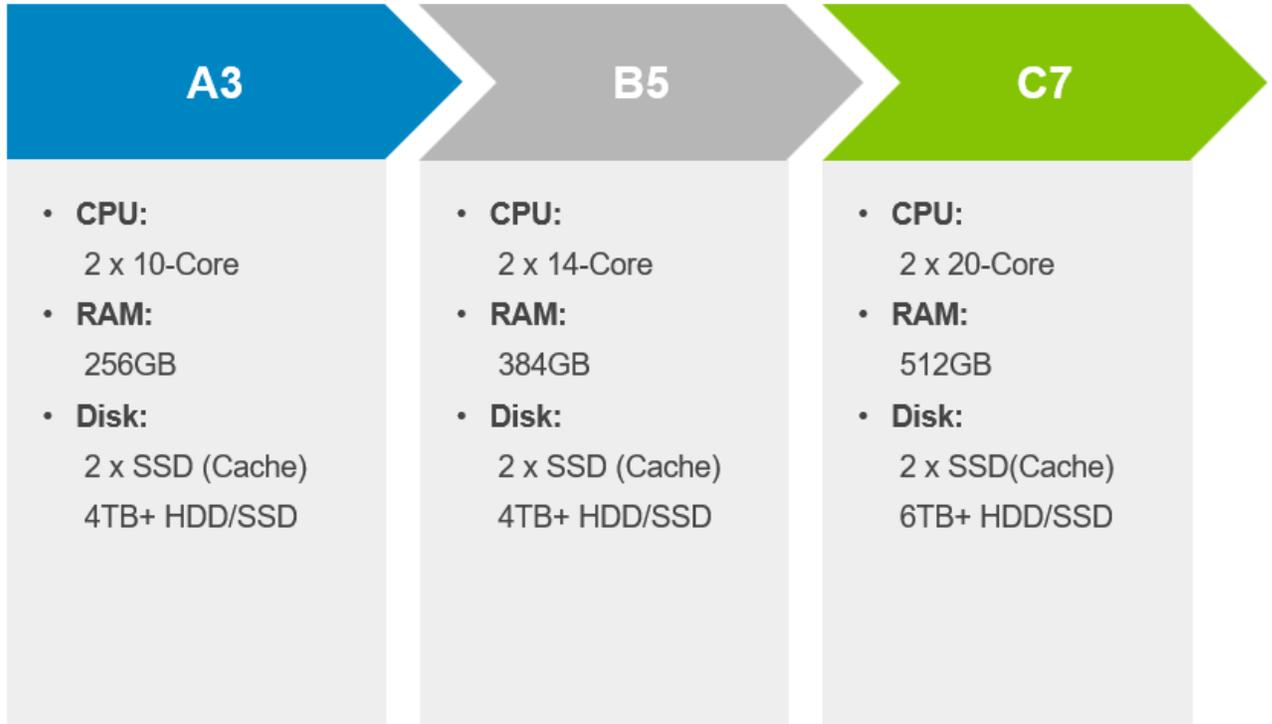
For more information on the FN410S IO Module for FX2, please visit [link](#)

## 3.2 Dell EMC VDI Optimized vSAN Ready Nodes

This update of the RA has streamlined the disk configurations for each of the platform configurations and by default each configuration has two diskgroups.

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

- A3 configuration is perfect for small scale, POC or low density cost-conscious 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.



In the VSAN Shared Tier 1 model the VDI desktops execute from the local storage on each of the Compute servers. The hypervisor used in this solution is vSphere. In this model, both the Compute and Management server hosts access vSAN storage.

The below table gives an overview of the Hybrid configurations in this RA. The usable capacity is calculated using VMware vSAN best practices so includes a reduction of 30% for slack space and FTT=1. The types of drives listed for the caching and capacity tier need to be adhered to and if these are altered it will invalidate the vSRN HCL configuration. This would then be a vSAN DIY HCL configuration if they were changed but all parts need to be on the vSAN HCL. The calculation for VM's per Node is using the task worker profile and for more details on this please refer to section 6.

Hybrid	A3	B5	C7
VMs per Node	120	170	230
Useable capacity per Node	1.7TB	1.7TB	2.5TB
CPU	2 x 10 Core	2 x 14 Core	2 x 20 Core
Memory	256GB	384GB	512GB
Caching Tier Flash	Performance: Class E: 30,000-100,000 writes per second Endurance: Class D >=7300 TBW	Performance: Class E: 30,000-100,000 writes per second Endurance: Class D >=7300 TBW	Performance: Class E: 30,000-100,000 writes per second Endurance: Class D >=7300 TBW
Capacity Tier	1.2TB 10K RPM SAS	1.2TB 10K RPM SAS	1.2TB 10K RPM SAS
Storage Controller	H730	H730	H730
NIC	10GB	10GB	10GB

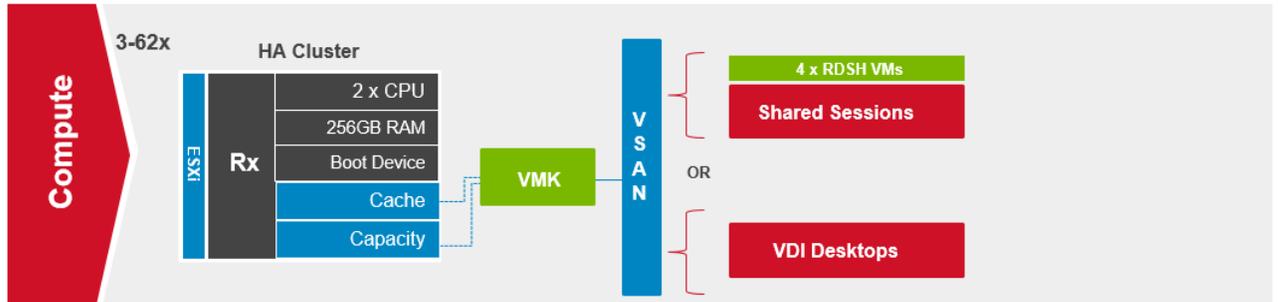
The All-Flash configurations in the below table have also the same best practices implemented, 30% slack space and a storage policy of FTT=1.

All-Flash	A3	B5	R730xd AF-8C7
VMs per Node	120	170	230
Useable capacity per Node	1.3TB	2.7TB	4TB
CPU	2 x 14 Core	2 x 20 Core	2 x 20 Core
Memory	256GB	384GB	512GB
Caching Tier Flash	Performance: Class E: 30,000-100,000 writes per second  Endurance: Class D >=7300 TBW	Performance: Class E: 30,000-100,000 writes per second  Endurance: Class D >=7300 TBW	Performance: Class E: 30,000-100,000 writes per second  Endurance: Class D >=7300 TBW
Capacity Tier	Performance: Class D: 20,000-30,000 writes per second  Endurance: Class D >=7300 TBW	Performance: Class E: 30,000-100,000 writes per second  Endurance: Class D >=7300 TBW	Performance: Class D 20,000-30,000 writes per second  Endurance: Class B >=1825 TBW
Storage Controller	HBA330	HBA330	HBA330
NIC	10GB	10GB	10GB

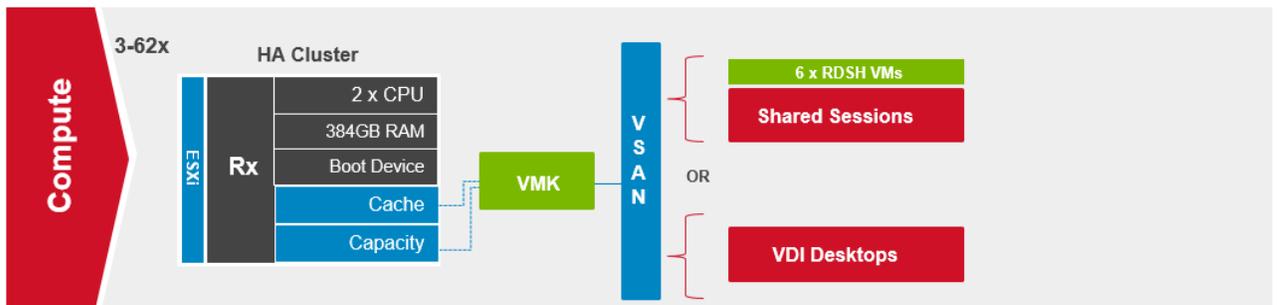
### 3.2.1 vSRN Server Configuration

The Shared Tier 1 vSRN Hybrid configuration model provides a scalable rack-based configuration that hosts user VDI sessions on local SSD (cache) and spinning disk (capacity) in the compute layer. The All-Flash configuration uses SSDs for the capacity layer.

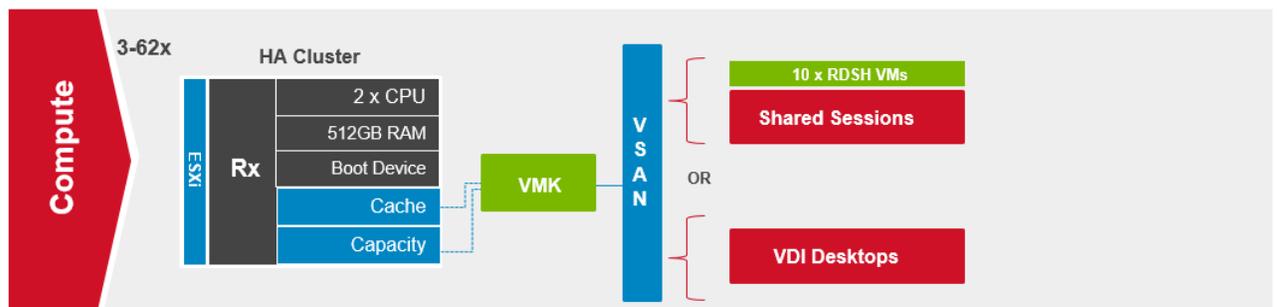
The A3 configuration below consists of two disk group, each diskgroup consists of 1 x SSD for cache & 1 x HDD/SSD for capacity.



The B5 configuration below consists of two disk groups and each diskgroup consists of 1 x SSD for cache & 2 x HDD/SSD for capacity.



The C7 configuration below consists of two disk group and each diskgroup consists of 1 x SSD for cache & 3 x HDD/SSD for capacity.

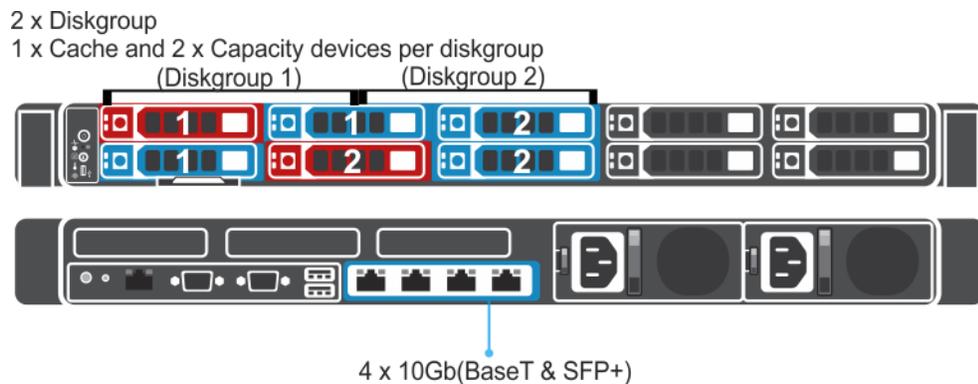


### 3.2.2 vSRN R630

The Dell R630 is a 1U platform with a broad range of configuration options. Each appliance comes equipped with dual CPUs, 10 to 20 cores, and up to 512GB of high-performance RAM by default. A minimum of two disks are required in each host, 1 x SSD for the Cache tier (Tier1) and 1 x HDD/SSD for the Capacity tier (Tier2).

#### 3.2.2.1 vSRN R630-A3 Configuration

There are two diskgroups in this configuration, diskgroup one (1) and diskgroup two (2) as depicted below, which consists of 1 x Cache SSD and 2 x Capacity HDD/SSD disks per diskgroup. There are two boot options, a single 64GB Satadom or a 16GB mirrored pair of SD modules and the ESXi hypervisor boots from here. Each platform can be outfitted with SFP+ or BaseT NICs.

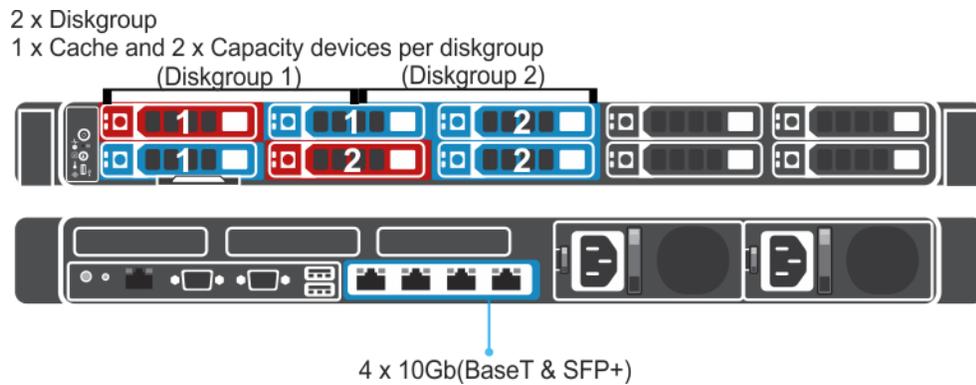


This would be classed as a vSRN HY-6 (Hybrid) or AF-4 (All-Flash) configuration on the vSRN HCL.

R630	A3
CPU	2 x E5-2630v4 (10C, 2.2GHz)
Memory	16 x 16GB 2400MT/s RDIMMs <b>Effective speed: 2400MT/s @ 256GB</b>
Storage Ctrls	PERC H730/H330 – no RAID
Boot Device	64GB Satadom or 2 X 16GB SD Module
Storage Hybrid	2 x 400GB SSD ( <b>Cache</b> ) 4 x 1.2TB HDD ( <b>Capacity</b> )
Storage All-Flash	2 x 400GB SSD ( <b>Cache</b> ) 2 x 1.92TB SSD ( <b>Capacity</b> )
Network	4 x 10Gb SFP+
iDRAC	iDRAC8 Ent w/ vFlash, 8GB SD
Power	2 x 750W PSUs

### 3.2.2.2 vSRN R630-B5 Configuration

There are two diskgroups in this configuration, diskgroup one (1) and diskgroup two (2) as depicted below, which consists of 1 x Cache SSD and 2 x Capacity HDD/SSD disks per diskgroup.

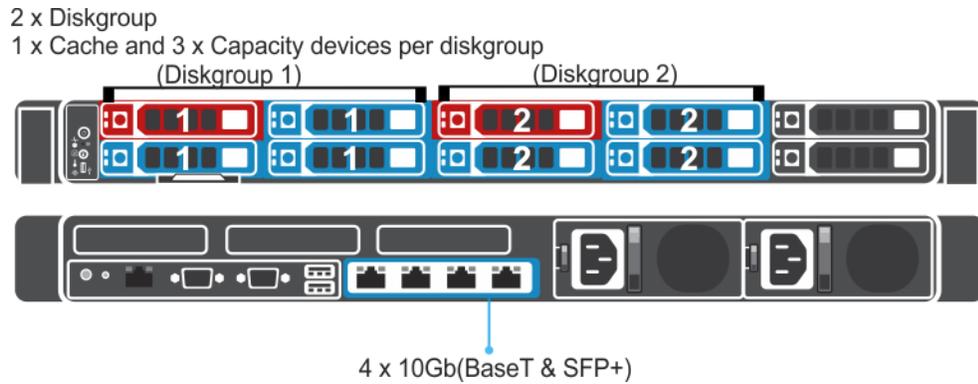


This would be classed as a vSRN HY-6 (Hybrid) or AF-4 (All-Flash) configuration on the vSRN HCL.

R630	B5
CPU	2 x E5-2660v4 (14C, 2.0GHz)
Memory	24 x 16GB 2400MT/s RDIMMs Effective speed: 2133MT/s @ 384GB
Storage Ctrlrs	PERC H730/H330 – no RAID
Boot Device	64GB Satadom or 2 X 16GB SD Module
Storage Hybrid	2 x 400GB SSD ( <b>Cache</b> ) 4 x 1.2TB HDD ( <b>Capacity</b> )
Storage All-Flash	2 x 400GB SSD ( <b>Cache</b> ) 4 x 1.92TB SSD ( <b>Capacity</b> )
Network	4 x 10Gb SFP+
iDRAC	iDRAC8 Ent w/ vFlash, 8GB SD
Power	2 x 750W PSUs

### 3.2.2.3 vSRN R630-C7 Configuration

There are two diskgroups in this configuration, diskgroup one (1) and diskgroup two (2) as depicted below, which consists of 1 x Cache SSD and 2 x Capacity HDD/SSD disks per diskgroup.



This would be classed as a vSRN HY-8 (Hybrid) or AF-8 (All-Flash) configuration on the vSRN HCL.

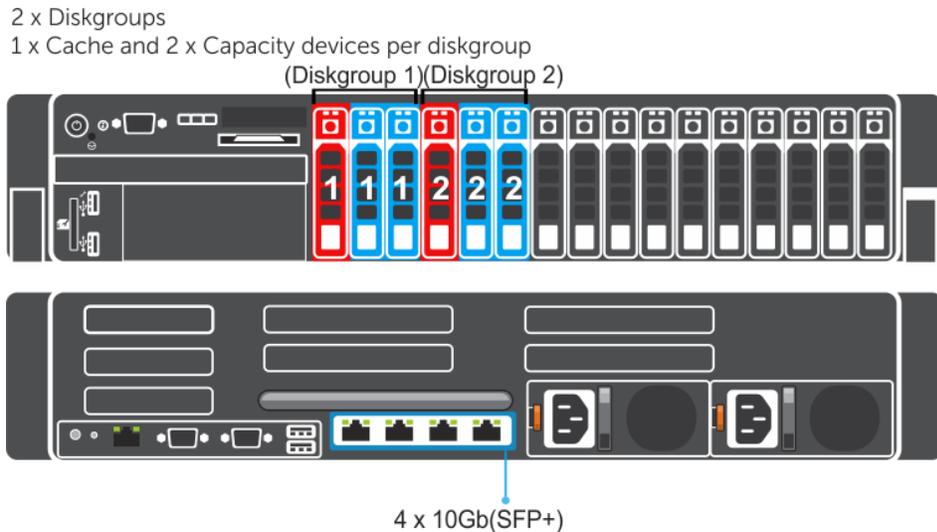
R630	C7
CPU	2 x E5-2698v4 (20C, 2.2GHz)
Memory	16 x 32GB 2400MT/s RDIMMs Effective speed: 2400MT/s @ 512GB
Storage Ctrlrs	PERC H730/H330 – no RAID
Boot Device	64GB Satadom or 2 X 16GB SD Module
Storage Hybrid	2 x 800GB SSD ( <b>Cache</b> ) 6 x 1.2TB HDD ( <b>Capacity</b> )
Storage All-Flash	2 x 800GB SSD ( <b>Cache</b> ) 6 x 1.92TB SSD ( <b>Capacity</b> )
Network	4 x 10Gb SFP+
iDRAC	iDRAC8 Ent w/ vFlash, 8GB SD
Power	2 x 750W PSUs

### 3.2.3 vSRN R730

This dual socket CPU platform runs the fastest Intel Xeon E5-2600 v4 family of processors, can host up to 24 DIMMs of DDR4 RAM, supports up to 16 x 2.5" SAS disks and can be outfitted with 2 double-wide GPU accelerators from AMD or NVIDIA.

#### 3.2.3.1 vSRN R730-A3 Configuration

The Dell PowerEdge R730 offers uncompromising performance and scalability in a 2U form factor. There are two diskgroups in this configuration, diskgroup one (1) and diskgroup two (2) as depicted below, which consists of 1 x Cache SSD and 2 x Capacity HDD/SSD disks per diskgroup.

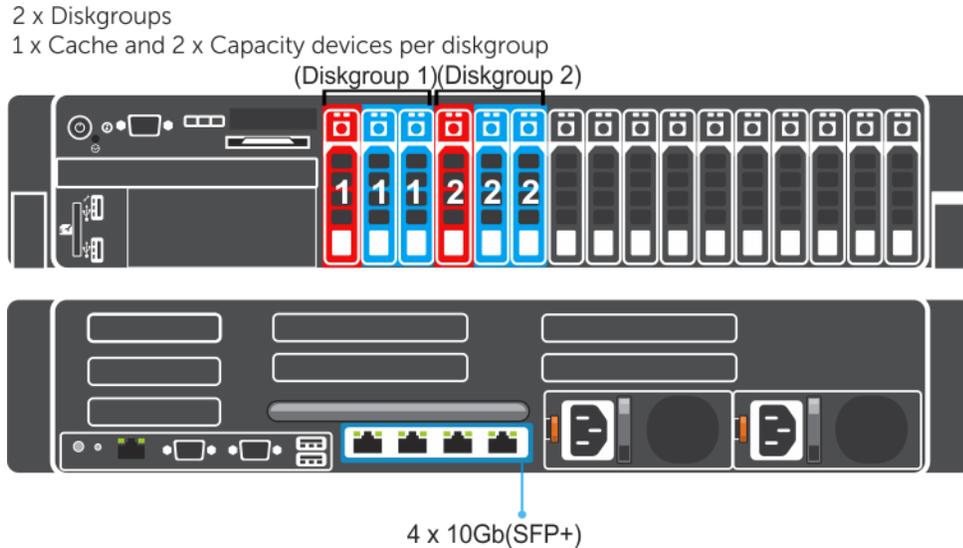


This would be classed as a vSRN HY-6 (Hybrid) or AF-4 (All-Flash) configuration on the vSRN HCL.

R730	A3
CPU	2 x E5-2630v4 (10C, 2.2GHz)
Memory	16 x 16GB 2400MT/s RDIMMs Effective speed: 2400MT/s @ 256GB
Storage Ctrlrs	PERC H730/H330 – no RAID
Boot Device	64GB Satadom or 2 X 16GB SD Module
Storage Hybrid	2 x 400GB SSD (Cache) 4 x 1.2TB HDD (Capacity)
Storage All-Flash	2 x 400GB SSD (Cache) 2 x 1.92TB SSD (Capacity)
Network	4 x 10Gb SFP+
iDRAC	iDRAC8 Ent w/ vFlash, 8GB SD
Power	2 x 750W PSUs

### 3.2.3.2 vSRN R730-B5 Configuration

There are two diskgroups in this configuration, diskgroup one (1) and diskgroup two (2) as depicted below, which consists of 1 x Cache SSD and 2 x Capacity HDD/SSD disks per diskgroup.

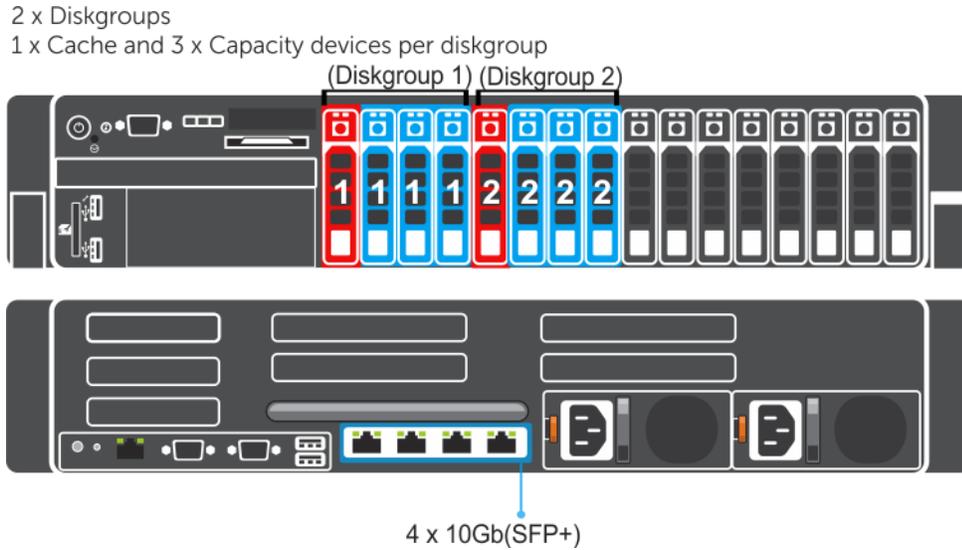


This would be classed as a vSRN HY-6 (Hybrid) or AF-4 (All-Flash) configuration on the vSRN HCL.

R730	B5
CPU	2 x E5-2660v4 (14C, 2.0GHz)
Memory	24 x 16GB 2400MT/s RDIMMs Effective speed: 2133MT/s @ 384GB
Storage Ctrlrs	PERC H730/H330 – no RAID
Boot Device	64GB Satadom or 2 X 16GB SD Module
Storage Hybrid	2 x 400GB SSD ( <b>Cache</b> ) 4 x 1.2TB HDD ( <b>Capacity</b> )
Storage All-Flash	2 x 400GB SSD ( <b>Cache</b> ) 4 x 1.92TB SSD ( <b>Capacity</b> )
Network	4 x 10Gb SFP+
iDRAC	iDRAC8 Ent w/ vFlash, 8GB SD
Power	2 x 750/1100W PSUs

### 3.2.3.3 vSRN R730-C7 Configuration

There are two diskgroups in this configuration, diskgroup one (1) and diskgroup two (2) as depicted below, which consists of 1 x Cache SSD and 3 x Capacity HDD/SSD disks per diskgroup.



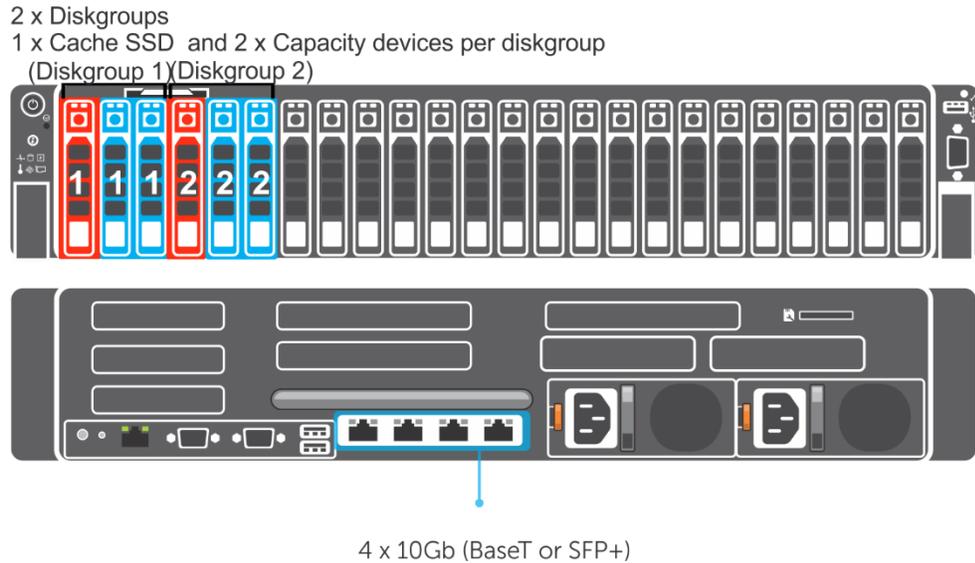
This would be classed as a vSRN HY-8 (Hybrid) or AF-8 (All-Flash) configuration on the vSRN HCL.

R730	C7
CPU	2 x E5-2698v4 (20C, 2.2GHz)
Memory	16 x 32GB 2400MT/s RDIMMs Effective speed: 2400MT/s @ 512GB
Storage Ctrls	PERC H730/H330 – no RAID
Boot Device	64GB Satadom or 2 X 16GB SD Module
Storage Hybrid	2 x 800GB SSD ( <b>Cache</b> ) 6 x 1.2TB HDD ( <b>Capacity</b> )
Storage All-Flash	2 x 800GB SSD ( <b>Cache</b> ) 6 x 1.92TB SSD ( <b>Capacity</b> )
Network	4 x 10Gb SFP+
iDRAC	iDRAC8 Ent w/ vFlash, 8GB SD
Power	2 x 750/1100W PSUs

### 3.2.4 vSRN R730XD

#### 3.2.4.1 vSRN R730XD-A3 Configuration

The R730XD configuration consists of both SSD for the cache and capacity tier. The cache tier uses write intensive SSD and mixed use SSDs for the capacity tier. There are different options for the boot device for ESXi, it can be on the SD cards, 2 x Flex bay drives or on a SATADOM configuration. There are two diskgroups in this configuration, diskgroup one (1) and diskgroup two (2) as depicted below.

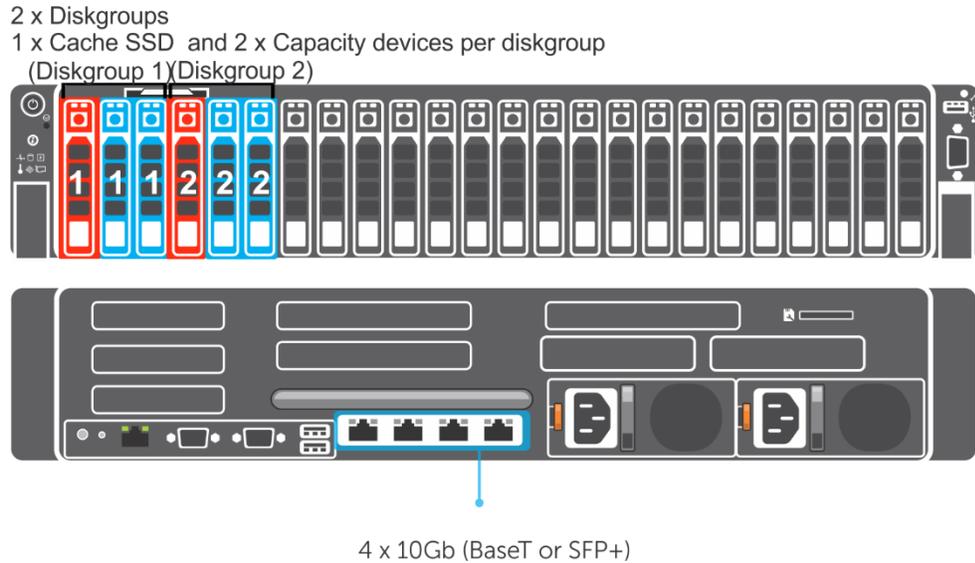


This would be classed as a vSRN HY-6 (Hybrid) or AF-4 (All-Flash) configuration on the vSRN HCL.

R730XD	A3
CPU	2 x E5-2630v4 (10C, 2.2GHz)
Memory	16 x 16GB 2400MT/s RDIMMs Effective speed: 2400MT/s @ 256GB
Storage Ctrls	PERC H730/H330 – no RAID
Boot Device	64GB Satadom or 2 X 16GB SD Module
Storage Hybrid	2 x 400GB SSD (Cache) 4 x 1.2TB HDD (Capacity)
Storage All-Flash	2 x 400GB SSD (Cache) 2 x 1.92TB SSD (Capacity)
Network	4 x 10Gb SFP+
iDRAC	iDRAC8 Ent w/ vFlash, 8GB SD
Power	2 x 750W PSUs

### 3.2.4.2 vSRN R730XD-B5 Configuration

There are two diskgroups in this configuration, diskgroup one (1) and diskgroup two (2) as depicted below, which consists of 1 x Cache SSD and 2 x Capacity HDD/SSD disks per diskgroup.

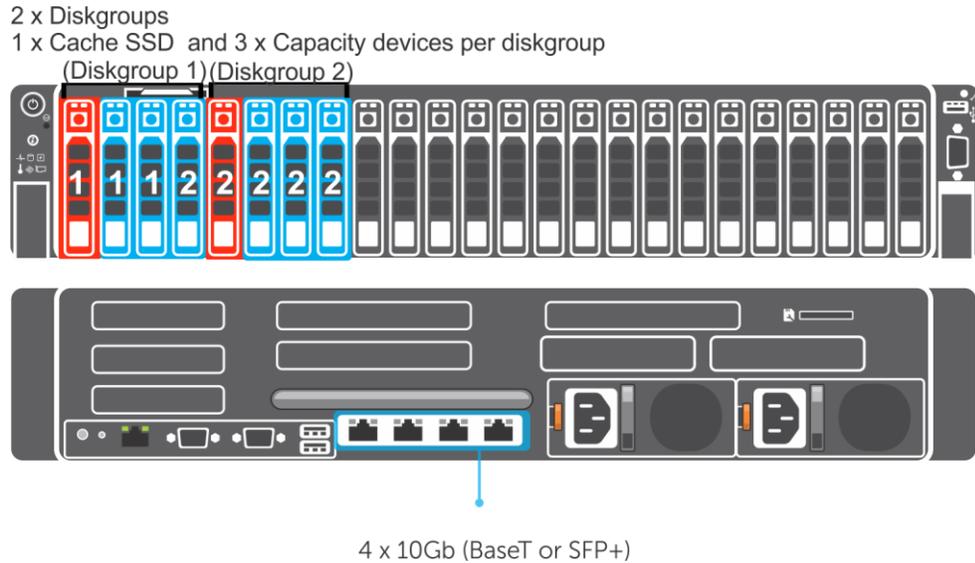


This would be classed as a vSRN HY-6 (Hybrid) or AF-4 (All-Flash) configuration on the vSRN HCL.

R730	B5
CPU	2 x E5-2660v4 (14C, 2.0GHz)
Memory	24 x 16GB 2400MT/s RDIMMs Effective speed: 2133MT/s @ 384GB
Storage Ctrls	PERC H730/H330 – no RAID
Boot Device	64GB Satadom or 2 X 16GB SD Module
Storage Hybrid	2 x 400GB SSD ( <b>Cache</b> ) 4 x 1.2TB HDD ( <b>Capacity</b> )
Storage All-Flash	2 x 400GB SSD ( <b>Cache</b> ) 4 x 1.92TB SSD ( <b>Capacity</b> )
Network	4 x 10Gb SFP+
iDRAC	iDRAC8 Ent w/ vFlash, 8GB SD
Power	2 x 750/1100W PSUs

### 3.2.4.3 vSRN R730XD-C7 Configuration

There are two diskgroups in this configuration, diskgroup one (1) and diskgroup two (2) as depicted below, which consists of 1 x Cache SSD and 3 x Capacity HDD/SSD disks per diskgroup.



This would be classed as a vSRN HY-8 (Hybrid) or AF-8 (All-Flash) configuration on the vSRN HCL.

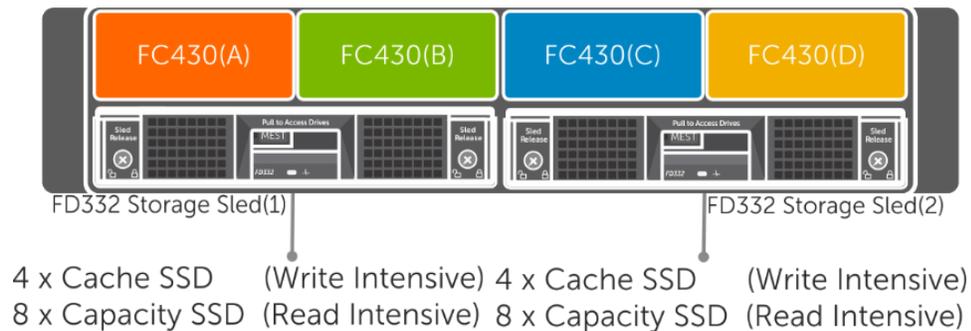
R730XD	C7
CPU	2 x E5-2698v4 (20C, 2.2GHz)
Memory	16 x 32GB 2400MT/s RDIMMs Effective speed: 2400MT/s @ 512GB
Storage Ctrls	PERC H730/H330 – no RAID
Boot Device	64GB Satadom or 2 X 16GB SD Module
Storage Hybrid	2 x 800GB SSD (Cache) 6 x 1.2TB HDD (Capacity)
Storage All-Flash	2 x 800GB SSD (Cache) 6 x 1.92TB SSD (Capacity)
Network	4 x 10Gb SFP+
iDRAC	iDRAC8 Ent w/ vFlash, 8GB SD
Power	2 x 750/1100W PSUs

### 3.2.5 vSRN FC430

The Dell PowerEdge FX2 is a 2U enclosure that can support up to four half-width compute sleds, up to eight quarter-width compute sleds, up to two full-width compute sleds, or a mix of compute sled types. The PowerEdge FX2s enclosure also supports half-width storage sleds mapped to the compute sleds.

#### 3.2.5.1 vSRN FC430-B5 Configuration

The Dell PowerEdge FC430 is a quarter-width sled that supports up to two Intel Xeon E5-2600v4 processors and eight DIMMs each. The FC430 FX configuration in this RA also includes 2 x PowerEdge FD332 storage sleds providing 4 x 400 GB cache SSD and 12 x 1.92TB capacity SSD per storage Sled. The SSDs in Slots 0-7 of Sled 1 are assigned to FC430 (A) and SSDs 8-15 are assigned to FC430 (B). The SSDs in Sled 2 have the same disk assignments, 0-7 on Sled 2 are assigned to FC430 (C) and 8-15 assigned to FC430 (D).



This would be classed as a vSRN AF-6 (All-Flash) configuration on the vSRN HCL.

FC430	B5
CPU	2 x E5-2660v4 (14C, 2.0GHz)
Memory	8 x 64GB 2400MT/s LRDIMMs Effective speed: 2400MT/s @ 512GB
Storage Ctrlrs	PERC H730 – no RAID
Boot Device	2 X 16GB SD Module
Storage	2 x 400GB SSD 2.5" (Cache) 4 x 1.92TB SSD 2.5" (Capacity)
Network	10Gb Ethernet
iDRAC	iDRAC8 Ent w/ vFlash, 8GB SD

### 3.2.6 vSAN FC630

These configurations at the time of release of this document are not a certified vSRN configuration but in fact a vSAN configuration built from certified components listed on the vSAN HCL and the vSRN HCL certification is pending.

#### 3.2.6.1 FC630-B5 Configuration

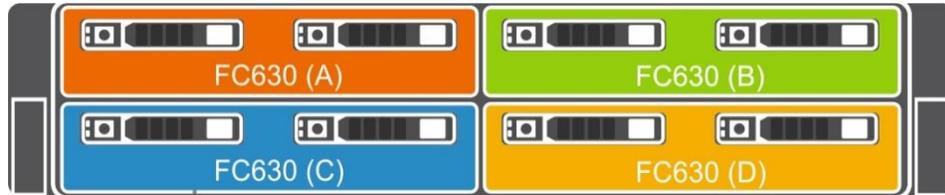
The Dell PowerEdge FC630 is a half-width sled that supports up to two Intel Xeon E5-2600v4 processors and 24 DIMMs per server. This configuration uses local disk for vSAN, there are two internal 2.5" SSDs per server, one disk for the cache device and the other is used as the capacity device.



FC630	B5
CPU	2 x E5-2660v4 (14C, 2.0GHz)
Memory	24 x 16GB 2400MT/s RDIMMs Effective speed: 2133MT/s @ 384GB
Storage Ctrls	PERC H730 – no RAID
Boot Device	2 X 16GB SD Module
Storage	1 x 400GB SSD 2.5" (Cache) 1 x 1.92TB SSD 2.5" (Capacity)
Network	10Gb Ethernet
iDRAC	iDRAC8 Ent w/ vFlash, 8GB SD

### 3.2.6.2 FC630-C7 Configuration

This FC630 configuration has a larger capacity device and also uses the C7 configuration of dual 2698v4 CPUs and 512 GB of Memory. This configuration is using local internal disk and that is the limiting factor in the configuration which determines the amount of VMs we can have per host.



1 x Cache SSD  
1 x Capacity SSD  
(per node)

FC630	C7
CPU	2 x E5-2698v4 (20C, 2.2GHz)
Memory	16 x 32GB 2400MT/s RDIMMs Effective speed: 2400MT/s @ 512GB
Storage Ctrls	PERC H730 – no RAID
Boot Device	2 X 16GB SD Module
Storage	1 x 400GB SSD 2.5" (Cache) 1 x 3.84TB SSD 2.5" (Capacity)
Network	10Gb Ethernet
iDRAC	iDRAC8 Ent w/ vFlash, 8GB SD

### 3.2.7 Summary comparison

The summary of all the VDI Platforms in this Reference Architecture are included in the below table. This table shows the total usable VM space which has been calculated on the RAW storage value. Also of note is that this the default storage policy has been applied which has an FTT=1. The All-Flash configuration can leverage the dedup & compression capabilities of VSAN and for examples on the potential space savings of this visit Section 4.1.3.1

Platform	Raw (per Host)	Slack Space (30%)	Usable VM Space (FTT=1)	Memory	CPU	Controller
A3 Hybrid	1.92TB	720GB	1.7TB	256GB	10 Core	H730
B5 Hybrid	4.8TB	720GB	1.7TB	384GB	14 Core	H730
C7 Hybrid	7.2TB	1.1TB	2.5TB	512GB	20 Core	H730
A3 All-Flash	3.84TB	576GB	1.3TB	256GB	10 Core	HBA330
B5 All-Flash	7.68TB	1.1TB	2.7TB	384GB	14 Core	HBA330
C7 All-Flash	11.52TB	1.72TB	4TB	512GB	20 Core	HBA330
FC430-B5	7.68TB	1.1TB	2.7TB	512GB	14 Core	H730
FC630-B5 All-Flash	1.92TB	288GB	672GB	384GB	14 Core	H730
FC630-C7 All-Flash	3.84TB	576GB	1.4TB	512GB	20 Core	H730

## 3.3 GPUs

### 3.3.1 NVIDIA Tesla GPUs

Accelerate your most demanding enterprise data center workloads with NVIDIA® Tesla® GPU accelerators. Scientists can now crunch through petabytes of data up to 10x faster than with CPUs in applications ranging from energy exploration to deep learning. In addition, Tesla accelerators deliver the horsepower needed to run bigger simulations faster than ever before. For enterprises deploying VDI, Tesla accelerators are perfect for accelerating virtual desktops.

#### 3.3.1.1 NVIDIA Tesla M10

The NVIDIA® Tesla® M10 is a dual-slot 10.5 inch PCI Express Gen3 graphics card featuring four mid-range NVIDIA Maxwell™ GPUs and a total of 32GB GDDR5 memory per card (8GB per GPU). The Tesla® M10 doubles the number of H.264 encoders over the NVIDIA® Kepler™ GPUs and improves encoding quality, which enables richer colors, preserves more details after video encoding, and results in a high-quality user experience.



The NVIDIA® Tesla® M10 GPU accelerator works with NVIDIA GRID™ software to deliver the industry's highest user density for virtualized desktops and applications. It supports up to 64 desktops per GPU card (up to 128 desktops per server) and gives businesses the power to deliver great graphics experiences to all of their employees at an affordable cost.

Specs	Tesla M10
Number of GPUs	4 x NVIDIA Maxwell™ GPUs
Total CUDA cores	2560 (640 per GPU)
GPU Clock	Idle: 405MHz / Base: 1033MHz
Total memory size	32GB GDDR5 (8GB per GPU)
Max power	225W
Form Factors	Dual slot (4.4" x 10.5")
Aux power	8-pin connector
PCIe	x16 (Gen3)
Cooling solution	Passive

### 3.3.1.2 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.



Accelerate your most demanding enterprise data center workloads with NVIDIA® Tesla® GPU accelerators. Scientists can now crunch through petabytes of data up to 10x faster than with CPUs in applications ranging from energy exploration to deep learning. Plus, Tesla accelerators deliver the horsepower needed to run bigger simulations faster than ever before. For enterprises deploying VDI, Tesla accelerators are perfect for accelerating virtual desktops.

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

## 3.4 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.

### 3.4.1 Wyse 3040 Thin Client (ThinOS, ThinLinux)

The Wyse 3040 is the industry's first entry-level Intel x86 quad-core thin client, powered by a quad-core Intel Atom 1.44GHz processor, delivering robust connectivity options with a choice of Wyse ThinOS or ThinLinux operating systems. The Wyse 3040 is Dell's lightest, smallest and most power-efficient thin client – it consumes 3.3 Watts in idle state – and offers superb performance and manageability for task and basic productivity users. Despite its small size, the 3040 includes all typical interfaces such as four USB ports including USB 3.1, two DisplayPort interfaces and wired and wireless options. It is highly manageable as it can be monitored, maintained, and serviced remotely via Wyse Device Manager (WDM) or Wyse Management Suite. For more information, please visit: [Link](#)



### 3.4.2 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](#)



### 3.4.3 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](#)

### 3.4.4 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](#)

### 3.4.5 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](#)

### 3.4.6 Wyse 5060 Thin Client (ThinOS, ThinLinux, WES7P, WIE10)



The Wyse 5060 offers high performance, reliability and flexible OS options, featuring all the security and management benefits of Dell thin clients. It come with flexible OS options: ThinOS (with or without PCoIP), ThinLinux, Windows Embedded Standard 7P (WES7P) or Windows 10 IoT Enterprise (WIE10). 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. It is powered by a quad-core AMD 2.4GHz processor, 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 Management Suite or Microsoft SCCM (5060 with Windows versions). Customers choosing WIE10 licenses can save about \$50/device/year as WIE10 qualifies under Microsoft Software Insurance, without the need to have more expensive VDA licenses to connect to a Windows virtual desktop. For more information, please visit: [Link](#)

### 3.4.7 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](#).

### 3.4.8 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](#)



### 3.4.9 Latitude 3480 and 5280 Mobile Thin Clients (Win 10 IoT)

Designed to securely deliver virtual desktops and applications to mobile users who want to connect a broad range of peripherals, the Latitude 3480 and 5280 mobile thin clients run **Windows 10 IoT Enterprise**. They support a wide variety of connection brokers including Citrix XenDesktop/XenApp, Microsoft RDS and VMware Horizon right out of the box, and are an ideal alternative to much less secure Chromebooks.

The Latitude 3480 features an Intel dual core processor with integrated graphics for a rich multimedia experience, and delivers great value with a 14" Full-HD display and robust connectivity with plenty of ports.



The Latitude 5280 delivers excellent performance with 12.5-inch, Full HD display. It offers the ability to support a 4K monitor via an optional docking station, and it supports a broad mix of peripheral attachments and network connections. They are easily manageable through Wyse Device Manager (WDM), Wyse Management Suite and Microsoft's System Center Configuration Manager (SCCM). For enhanced security, optional advanced threat protection in the form of Dell Threat Defense offers proactive malware protection. For more information, please visit the following pages for: [Latitude 3480](#) , [Latitude 5280](#)



## 4 Software Components

### 4.1 VMware

#### 4.1.1 VMware vSphere 6.x

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 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 used in the solution is a vCenter Server Appliance (VCSA) residing on a host in the management Tier. Horizon Composer is installed on a standalone Windows Server 2012 R2 or later VM when using the VCSA.

VMware vSphere® is the next-generation infrastructure for next-generation applications. It provides a powerful, flexible, and secure foundation for business agility that accelerates the digital transformation to cloud computing and promotes success in the digital economy.

#### **Improved Appliance Management**

vCenter Server Appliance also exclusively provides improved appliance management capabilities. The vCenter Server Appliance Management interface continues its evolution and exposes additional configuration data. In addition to CPU and memory statistics, it now shows network and database statistics, disk space usage and health data. This reduces reliance on a command-line interface for simple monitoring and operational tasks.

#### **VMware vCenter High Availability**

vCenter Server has a new native high availability solution that is available exclusively for vCenter Server Appliance. This solution consists of active, passive, and witness nodes that are cloned from the existing vCenter Server instance. The VMware vCenter® High Availability (vCenter HA) cluster can be enabled, disabled, or destroyed at any time. There is also a maintenance mode that prevents planned maintenance from causing an unwanted failover. vCenter HA uses two types of replication between the active and passive nodes: Native PostgreSQL synchronous replication is used for the vCenter Server database; a separate asynchronous file system replication mechanism is used for key data outside of the database.

Failover can occur when an entire node is lost—host failure, for example—or when certain key services fail. For the initial release of vCenter HA, a recovery time objective (RTO) of about 5 minutes is expected, but this can vary slightly depending on the load, size, and capabilities of the underlying hardware.

## Backup and Restore

New in vCenter Server is native backup and restore for the vCenter Server Appliance. This new, out-of-the-box functionality enables users to back up vCenter Server and Platform Services Controller appliances directly from the VAMI or API. The backup consists of a set of files that is streamed to a storage device of the user's choosing using SCP, HTTP(S), or FTP(S) protocols. This backup fully supports VCSA instances with both embedded and external Platform Services Controller instances.

## vSphere HA Support for NVIDIA GRID vGPU Configured VMs

vSphere HA now protects VMs with the NVIDIA GRID vGPU shared pass-through device. In the event of a failure, vSphere HA attempts to restart the VMs on another host that has an identical NVIDIA GRID vGPU profile. If there is no available healthy host that meets this criterion, the VM fails to power on. For more information on HA Support for NVIDIA GRID vGPU please visit the Blog article located [Here](#).

For more information on VMware vSphere and what's new in this release, visit [link](#).

### 4.1.2 VMware vSAN

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

**Deduplication and compression:** 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 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:** 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 vSAN IOPS limits are now available. Quality of service for vSAN is a Storage Policy Based Management (SPBM) rule. Because quality of service is applied to 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.

**iSCSI Access:** vSAN iSCSI Access enables vSAN to support physical workloads using iSCSI initiators and allows you to eliminate cost and complexity of a separate physical array. All core functionality is available and managed through vCenter

**Full-Featured PowerCLI:** vSAN includes a complete set of PowerCLI cmdlets providing the scalability and ease of enterprise-class automation.

**Next-Generation Hardware:** vSAN continues to support next-generation hardware including large capacity drives with 512e support.

For more information on what's new in vSAN, visit this [Link](#).

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 vSAN.

### 4.1.3 vSAN Best Practices

When determining the amount of capacity required for a vSAN Design we need to pay close attention to the number of Failures to Tolerate (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).

We also need to factor in how much free capacity or "Slack Space" needs to be preserved when designing the capacity requirement for the vSAN Cluster. The recommendation by VMware is that this should be 30%. The reasoning for this slack space size that the 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 vSAN6.X Cluster [here](#).

#### 4.1.3.1 All-Flash versus Hybrid

The most significant new features in this latest version of vSAN are Deduplication & Compression and erasure coding. These features are only supported in an All-Flash 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 it 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 and 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 vSAN 6.5, RAID-1 (Mirroring) was used as the failure tolerance method. vSAN 6.5 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.

RAID-5 as described in the above table 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 vSAN Cluster.

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

**vSAN Hybrid Read:** For an object placed on a 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 vSAN does not consume more read cache than is necessary, this avoids caching the data in multiple locations.

**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 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.

**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.

**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, the other 70% is assigned to the read cache.

### 4.1.3.2 VM Storage Policies for 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 vSAN layer when a VM is created. The VM virtual disk is distributed across the vSAN datastore per policy definition to meet the requirements.

VMware Horizon 7 has a built in storage policy for 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).

Name	Description	VC
Virtual SAN Default Storage ...	Storage policy used as default f...	vsm-vc.osprey.com
VVol No Requirements Policy	Allow the datastore to determine...	vsm-vc.osprey.com
VM_HOME_e8971e95-d947...	View Auto Created	vsm-vc.osprey.com
OS_DISK_FLOATING_e897...	View Auto Created	vsm-vc.osprey.com
REPLICA_DISK_e8971e95-...	View Auto Created	vsm-vc.osprey.com
PERSISTENT_DISK_e8971...	View Auto Created	vsm-vc.osprey.com

Each policy can be edited but it is recommended to refer to design and sizing guide for vSAN 6.x located [here](#) before making any change to the policy.

**OS\_DISK\_FLOATING\_e8971e95-d947-4fd2-9cee-8455f3e95a83: Edit VM Storage Policy**

**Name and description**  
Rule-Set 1  
Storage compatibility

**Rule-Set 1**  
Select rules specific for a datastore type. Rules can be based on data services provided by datastore or based on tags. The VM storage policy will match datastores that satisfy all the rules in at least one of the rule-sets.

**Rules based on data services** vSAN

- Number of disk stripes per object: 1
- Flash read cache reservation (%): 0.0000
- Number of failures to tolerate: 1
- Object space reservation (%): 0

<Add rule>

**Rules based on tags**  
Add tag-based rule...

**Storage Consumption Model**  
A virtual disk with size 100 GB would consume:  
Storage space: 200.00 GB  
Initially reserved storage space: 0.00 B  
Reserved flash space: 0.00 B

Add another rule set Remove this rule set  
OK Cancel

#### 4.1.4 VMware Horizon

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

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

virtual desktop infrastructure. For the complete set of details, please see the Horizon resources page at <http://www.vmware.com/products/horizon-view/resources.html>.

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.

#### 4.1.4.1 What's new in this release of Horizon

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

##### 4.1.4.1.1 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.

#### 4.1.4.1.2 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.

#### 4.1.4.1.3 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.

#### 4.1.4.1.4 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.

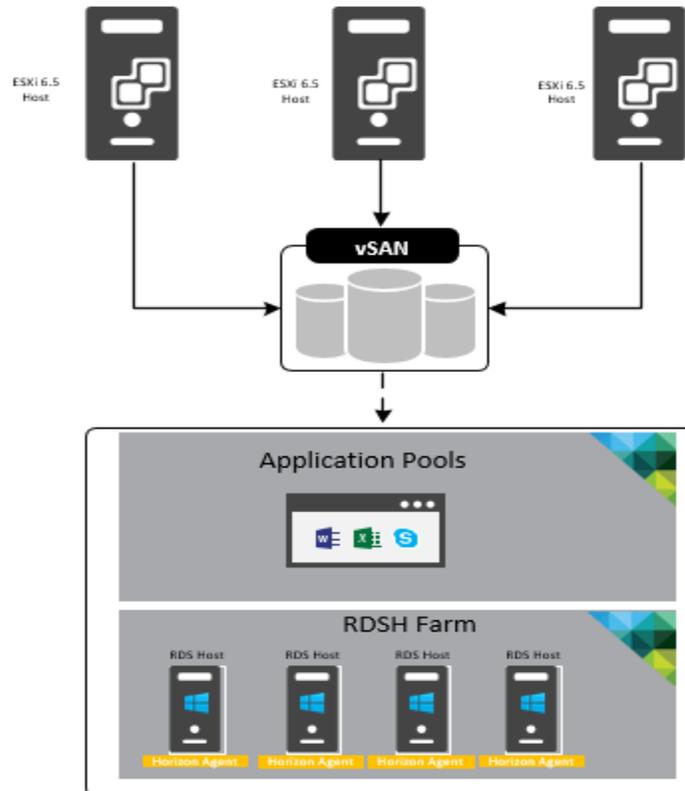
**Policy-Managed Client Features**, which enables IT to use policy to define which specific security-impacting features, are accessible upon login. These include clipboard redirection, USB, printing, and client-drives. All of these can be enforced contextually, based on role, evaluated at logon/logoff, 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.

## 4.1.5 VMware Horizon Apps

Horizon Apps (Deliver Virtual Applications to Any Device, Anywhere!!)

The ability to support published application with Horizon has been available with VMware Horizon 7 Enterprise but now we have the standalone options of VMware Horizon Apps Standard & Advance. Horizon provides a platform to deliver an enterprise-class application publishing solution as well as a virtual desktops. VMware Horizon leverages Microsoft Remote Desktop Session Host (RDSH) to deliver published applications as well as published desktops running on the Microsoft RDSH Server.



VMware Horizon features and components such as the Blast Extreme display protocol, instant-clone provisioning, App Volumes application delivery, and User Environment Manager are heavily integrated into RDSH to provide a seamless user experience and an easy-to-manage, scalable solution.

Next Generation Application and Delivery Platform via Just-in-Time Management Platform (JMP) are available via Horizon Apps Advanced Edition or Horizon 7 Enterprise. MP apps offer simple image management, quick scaling and zero downtime while providing simple and powerful user and group policy controls at the push of a button.

Horizon JMP Apps are composed of:

**VMware Instant Clones Technology:** RDSH farms can be provisioned rapidly via instant cloned Microsoft RDSH Servers.

**VMware App Volumes:** real time application delivery via Appstacks mapped to the RDSH Servers. App Volumes allows you to separate the Windows OS image from the application images. Groups of applications can be installed into virtual disks called AppStacks. The appropriate AppStack can then be assigned to the RDSH farm to personalize the applications delivered.

**VMware User Environment Manager (UEM):** VMware UEM simplifies end-user profile management by offering personalization and dynamic policy configuration to the RDSH Server you can configure fine-grained policies for folder redirection, mapping the user's home drive, configuring location-based printers, and application blocking—all based on user accounts. You can use the Horizon 7 Smart Policies feature to enable or disable client features based on user device, location, and other defined conditions.

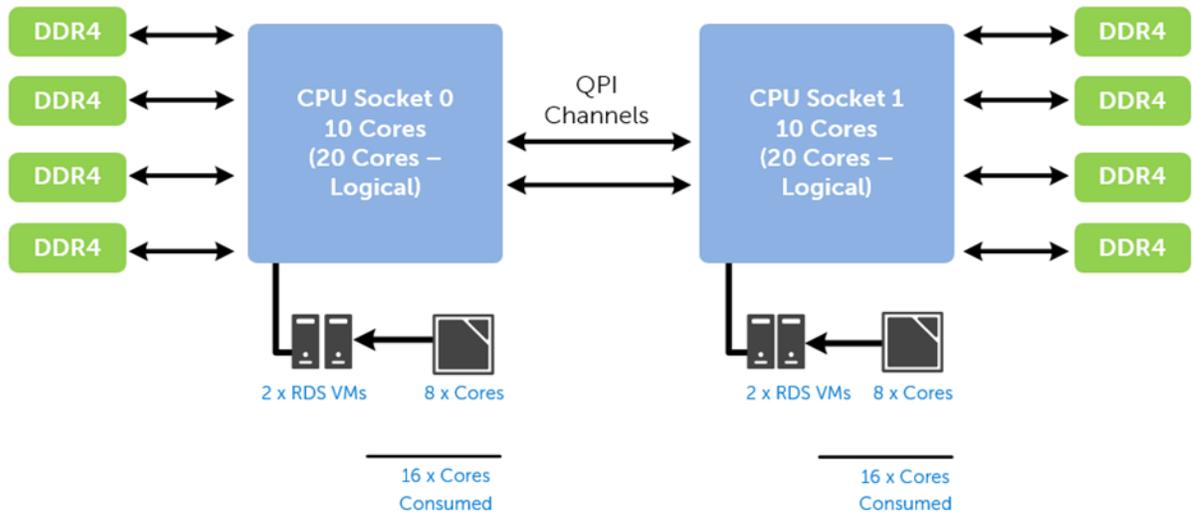
## 4.1.6 NUMA architecture considerations

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 VSRN is as follows:

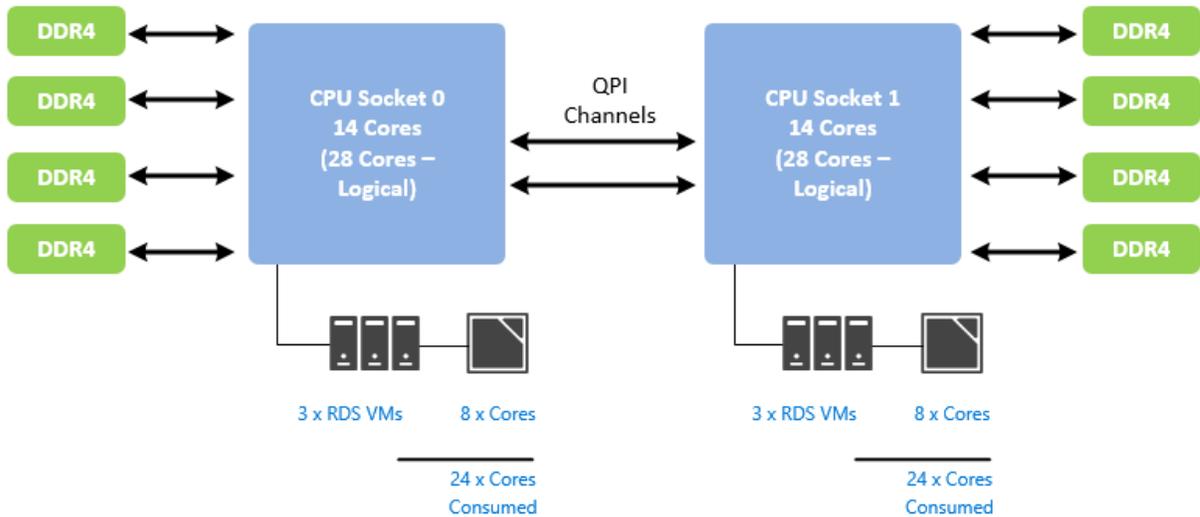
### 4.1.6.1 A3 Configuration NUMA Alignment

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



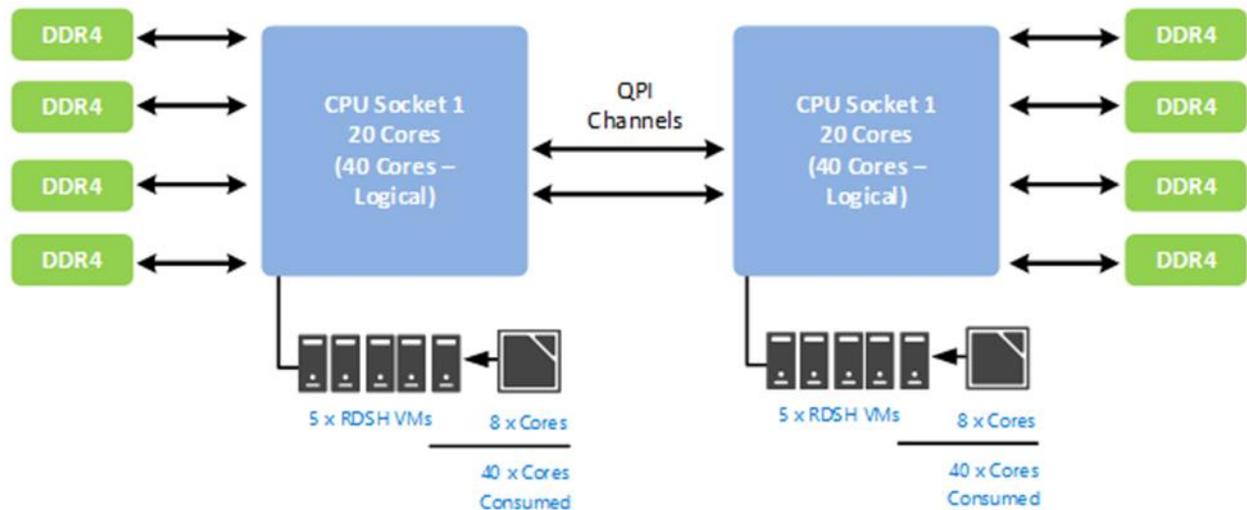
#### 4.1.6.2 B5 Configuration NUMA Alignment

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



#### 4.1.6.3 C7 Configuration NUMA Alignment

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

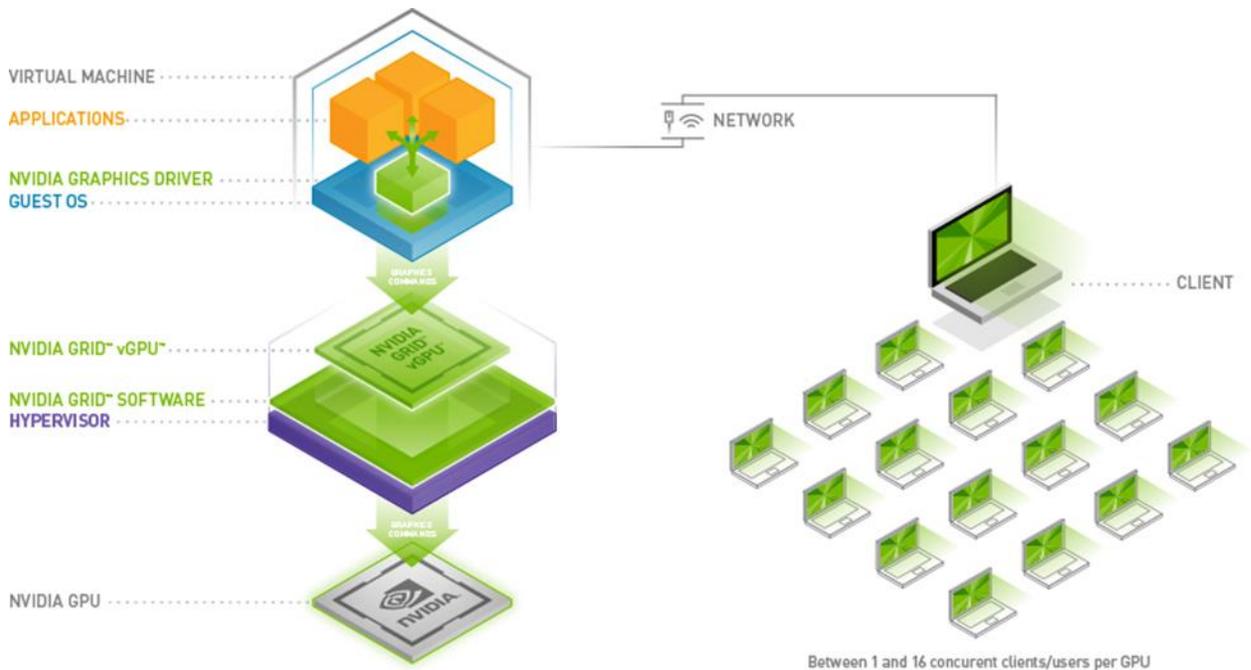


## 4.2 NVIDIA GRID vGPU

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 by NVIDIA Corporation. Copyright NVIDIA Corporation)

#### 4.2.1 vGPU Profiles

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

The 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: [LINK](#)

The number of users per server is determined by the number of GPU cards in the system (max 2), vGPU profiles used for each GPU in a card (2 GPUs per card), and GRID license type. The same profile must be used on a single GPU but profiles can differ across GPUs in a single card.

**NVIDIA® Tesla® M10 GRID vGPU Profiles:**

Card	vGPU Profile	Graphics Memory (Frame Buffer)	Virtual Display Heads	Maximum Resolution	Maximum Graphics-Enabled VMs		
					Per GPU	Per Card	Per Server (2 cards)
Tesla M10	M10-8Q	8GB	4	4096x2160	1	4	8
	M10-4Q	4GB	4	4096x2160	2	8	16
	M10-2Q	2GB	4	4096x2160	4	16	32
	M10-1Q	1GB	2	4096x2160	8	32	64
	M10-0Q	512MB	2	2560x1600	16	64	128
	M10-1B	1GB	4	2560x1600	8	32	64
	M10-0B	512MB	2	2560x1600	16	64	128
	M10-8A	8GB	1	1280x1024	1	4	8
	M10-4A	4GB			2	8	16
	M10-2A	2GB			4	16	32
	M10-1A	1GB			8	32	64

Card	vGPU Profile	Guest VM OS Supported*		GRID License Required
		Win	64bit Linux	
Tesla M10	M10-8Q	●	●	GRID Virtual Workstation
	M10-4Q	●	●	
	M10-2Q	●	●	
	M10-1Q	●	●	
	M10-0Q	●	●	
	M10-1B	●		GRID Virtual PC
	M10-0B	●		
	M10-8A	●		GRID Virtual Application
	M10-4A	●		
	M10-2A	●		
	M10-1A	●		

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	

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

**NVIDIA® Tesla® M60 GRID vGPU Profiles:**

Card	vGPU Profile	Graphics Memory (Frame Buffer)	Virtual Display Heads	Maximum Resolution	Maximum Graphics-		
					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	

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

### 4.2.1.1 GRID vGPU Licensing and Architecture

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

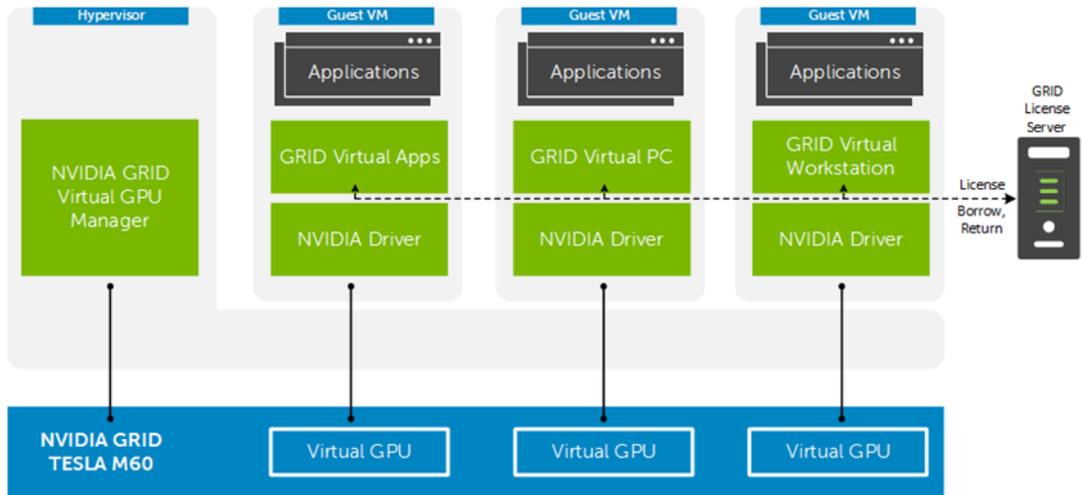


NVIDIA GRID Virtual Applications	NVIDIA GRID Virtual PC	NVIDIA GRID Virtual Workstation
For organizations deploying RDSH solutions. Designed to deliver Windows applications at full performance.	For users who need a virtual desktop, but also need a great user experience leveraging PC applications, browsers, 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.



(Image provided by NVIDIA Corporation. Copyright NVIDIA Corporation)

## 5 Solution architecture for vSRN with Horizon

### 5.1 Management Server Infrastructure

Role	vCPU	RAM (GB)	NIC	OS + Data vDisk (GB)	Tier 2 Volume (GB)
VMware vCenter Server Appliance	2	8	1	290	-
Horizon Connection Server	2	8	1	60	-
SQL Server	5	8	1	60	210 (VMDK)
File Server	1	4	1	60	2048 (VMDK)
<b>Total</b>	<b>7</b>	<b>28</b>	<b>4</b>	<b>470</b>	<b>2358</b>

#### 5.1.1 RDSH VM Configuration

The recommended number of RDSH VMs and their configurations on ESXi are summarized below and take into account proper NUMA balancing assuming the CPU. The amount of RDSH vms per Server depend on the CPU configuration and for more information on NUMA please refer to the [NUMA Architecture Considerations](#) section.

RDSH VM configuration on ESXi

Role	vCPU	RAM (GB)	NIC	OS vDisk (GB)	Tier 2 Volume (GB)
RDSH VM	8	32	1	80	-

## 5.1.2 NVIDIA GRID License Server Requirements

When using NVIDIA Tesla M60 cards, graphics enabled VMs must obtain a license from a GRID License server on your network to be entitled for vGPU. To configure, a virtual machine with the following specifications must be added to a management host in addition to the management role VMs.

Role	vC PU	RAM (GB)	NIC	OS + Data vDisk (GB)	Tier 2 Volume (GB)
NVIDIA GRID License Srv	2	4	1	40 + 5	-

GRID License server software can be installed on a system running the following operating systems:

- Windows 7 (x32/x64)
- Windows 8.x (x32/x64)
- Windows 10 x64
- Windows Server 2008 R2
- Windows Server 2012 R2
- Red Hat Enterprise 7.1 x64
- CentOS 7.1 x64

Additional license server requirements:

- A fixed (unchanging) IP address. The IP address may be assigned dynamically via DHCP or statically configured, but must be constant.
- At least one unchanging Ethernet MAC address, to be used as a unique identifier when registering the server and generating licenses in NVIDIA's licensing portal.
- The date/time must be set accurately (all hosts on the same network should be time synchronized).

### 5.1.3 SQL Databases

The VMware databases will be hosted by a single dedicated SQL 2016 (or higher) 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:

- Horizon Connection Server
- vCenter
- Horizon 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).

### 5.1.4 DNS

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

#### 5.1.4.1 DNS for SQL

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

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

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

## 5.2 Storage Architecture Overview

All Dell VSRN Hybrid appliances come with two tiers of local storage by default, SSD for performance and HDD for capacity. These diskgroups need a minimum of 1 x SSD and 1 x HDD per diskgroup and with the Dell VSRN appliances we offer the option of one or two diskgroup configurations. These local storage diskgroups are configured into one Software Defined Storage pool via VSAN which are shared across all hosts in the VSRN Cluster.

### 5.2.1 vSAN

To enable vSAN, simply select the Datacenter in vSphere, go to menu Manage, Settings and General. Click Edit button and select Turn ON vSAN. There are two modes when configuring vSAN: automatic and manual. If you setup vSAN in automatic mode then all empty local disks will be used in the creation of the shared data store. If you configure it in manual mode then you can manually select which disks to use in the creation of the data store.



## 5.3 Virtual Networking

### 5.3.1 VSRN Network Configuration

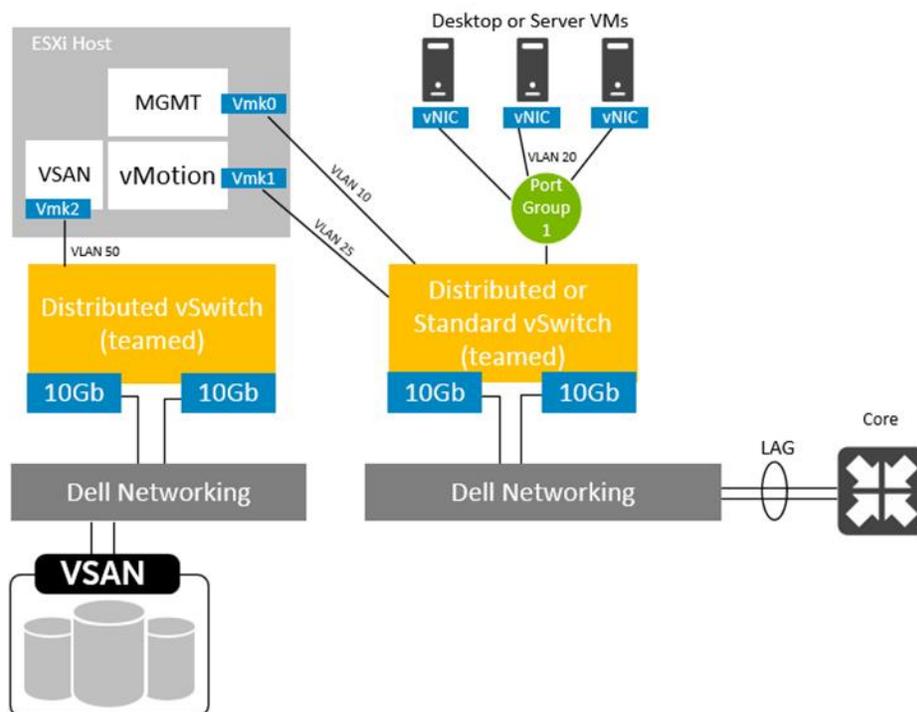
The network configuration for the Dell vSRN appliances utilizes a 10Gb converged infrastructure model. The management, vMotion, VDI traffic and vSAN are configured across 2 x 10Gb NICs configured in an active/active team and the traffics is separated out via VLAN. When deploying larger cluster configurations it may be optimal to split out the management and vSAN traffic. The Server configurations are equipped with 4 x 10GB connections to facilitate this

The network configuration in this model is the same for the compute and management layer. They both share the local storage VSAN configuration taking advantage of HA including Live Migrations. The following outlines the VLAN requirements for the Compute and Management hosts in this solution model:

- VSRN 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
  - vSAN VLAN: Configured for vSAN traffic – L2 switched only via ToR switch
  - vMotion VLAN: Configured for Live Migration traffic – L2 switched only, trunked from Core (HA only)
- A VLAN for iDRAC is configured for all hardware management traffic – L3 routed via core switch

This traffic is combined within a single switch in smaller stacks to minimize the initial investment; however, VLANs are required for each traffic type to enable traffic separation. Each VSAN host will have a quad port NDC which includes 4 x 10Gb interfaces and the 2 free 10Gb interfaces can be used to separate out the vSAN traffic for larger clusters, configure the LAN traffic from the server to the ToR switch as a LAG.

The following screenshot shows the VMKernel adapter for the management network (vmk0) and vMotion network (vmk1) on a standard or distributed switch and 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.

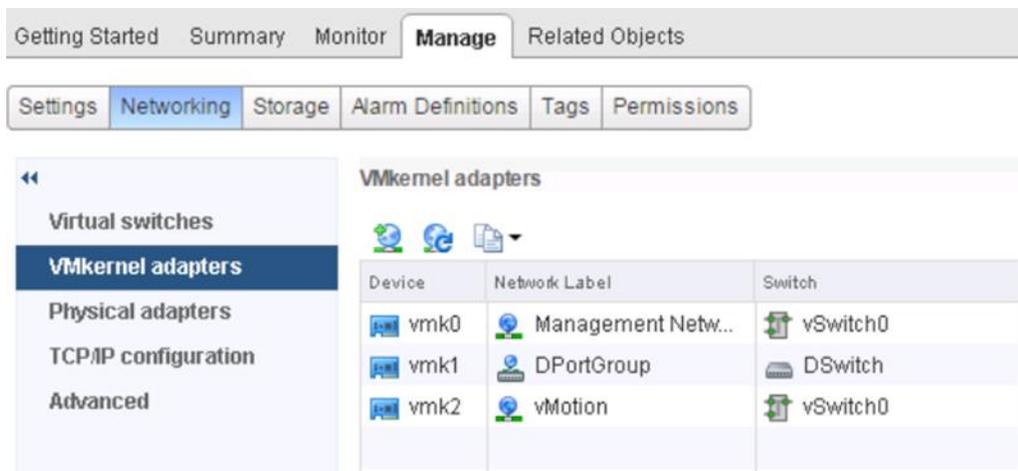
dvSwitches should be used as desired for VM traffic especially in larger deployments to ease the management burden across numerous hosts. In the VSRN rack model both the mgmt. hosts connect to

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.

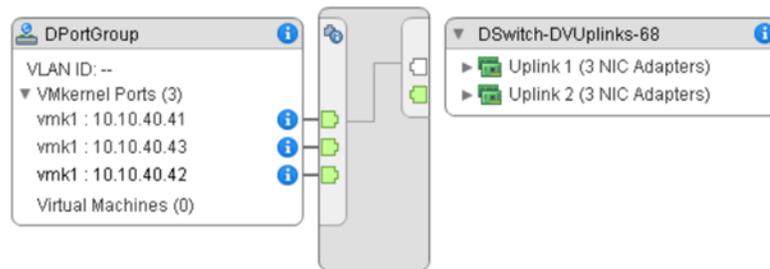
vSAN cluster networking includes at least two VMKernel ports, one for management traffic and one for vSAN traffic. If vMotion, Storage vMotion or High Availability functionality is required in addition, a third VMKernel port is to be configured for this.

vSAN traffic can be used on 1Gb networks as well as 10Gb networks for Hybrid configuration but 10Gb 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 10Gb network with standard switch. It is recommended that the network configuration for the vSAN storage is a 10Gb network with distributed switch configuration.



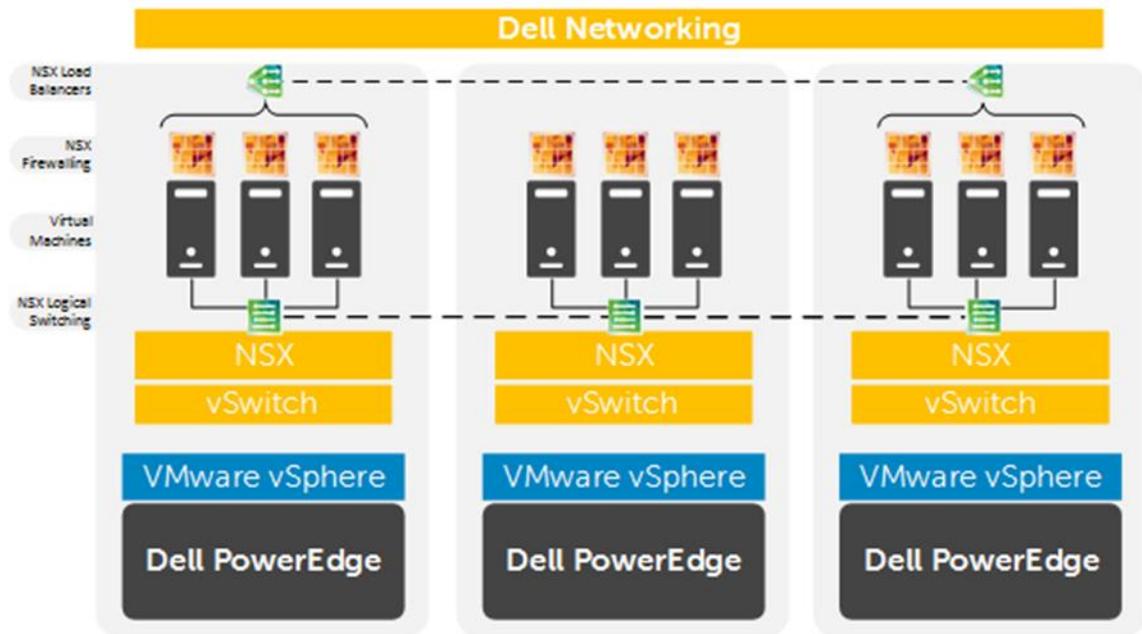
The distributed switch configuration is the same on all VSRN storage hosts. It is recommended to have at least two uplinks for each host to provide load balancing and fail back redundancy.



### 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 providing 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 complements Dell Networking components deployed ToR, leaf/spine or at the core.

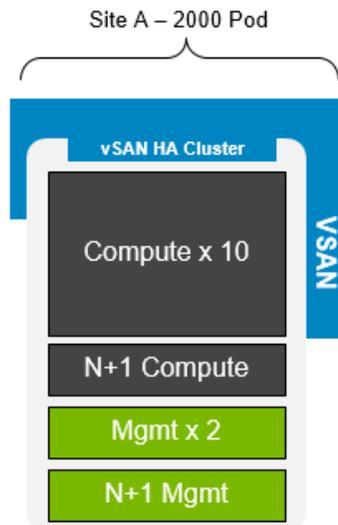


<b>Key Features of Dell Open Networking and VMware NSX</b>	
Power of Choice	Choose from best-of-breed open networking platforms, operating systems and applications.
Accelerated Innovation	Take advantage of open networking with open source standards-based tools and expertise to help accelerate innovation.
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	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.

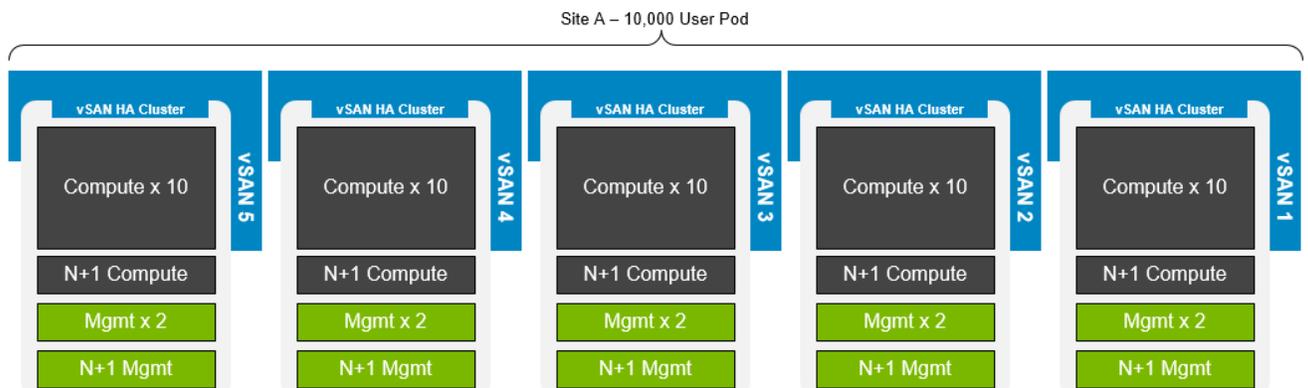
For more information on VMware NSX and integrated offers from Dell Networking please see the Dell Networking [Solution Brief](#) and the [Reference Architecture](#).

## 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 vSAN SDS pool in a modular fashion. The scaling limit for vSAN 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 10 compute nodes for with 200 Task User VMs 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 VCSA is protected by native HA which is a new feature with vSphere 6 and this is covered in the VMware vSphere section.



- 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. The below table shows the scalability options for each component.

Component	Metric	Horizontal Scalability	Vertical Scalability
<b>Virtual Desktop Host/Compute Servers</b>	VMs per physical host	Additional hosts and clusters added as necessary	Additional RAM or CPU compute power
<b>Horizon Composer</b>	Desktops per instance	Additional physical servers added to the Management cluster to deal with additional management VMs.	Additional RAM or CPU compute power
<b>Horizon Connection Servers</b>	Desktops per instance	Additional physical servers added to the Management cluster to deal with additional management VMs.	Additional VCS Management VMs.
<b>VMware vCenter</b>	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.
<b>Database Services</b>	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
<b>File Services</b>	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, 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.

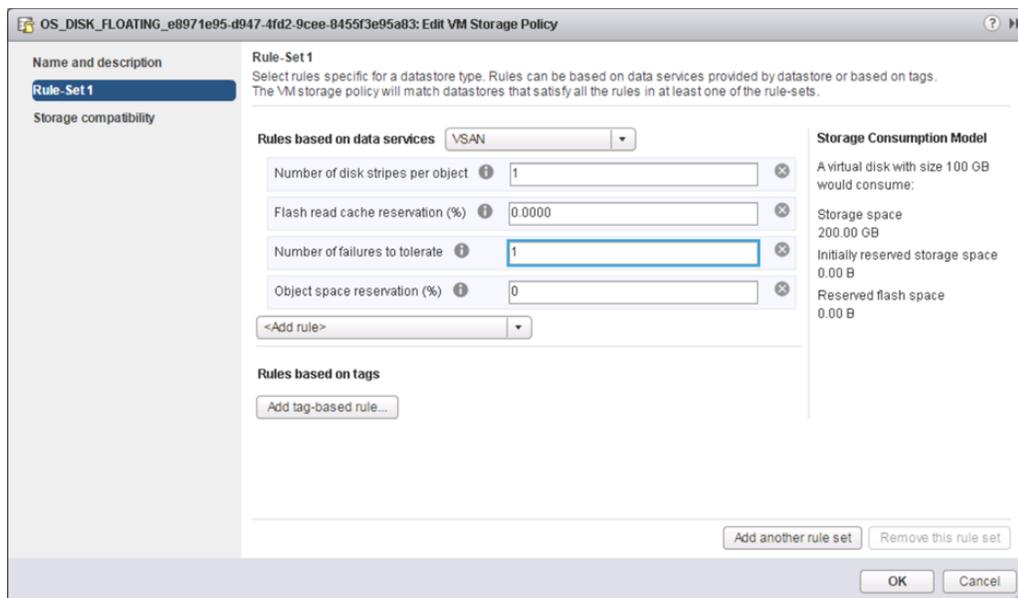
- Additional switches added to the existing thereby equally spreading each host's network connections across multiple switches.

- 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 vSAN HA/ FTT configuration

The minimum configuration required for vSAN 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, and 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 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 vSAN storage so can make full use of vSphere HA and vSRN. Compute 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 vSAN 6.x limits dictated by VMware (6,400 VMs per 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.

<b>vSAN Limits</b>	<b>Minimum</b>	<b>Maximum</b>
Number of supported ESXi hosts per vSAN cluster	3	64
Number of supported VMs per host	1	200*
Number of supported VMs per vSAN Cluster	1	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 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.

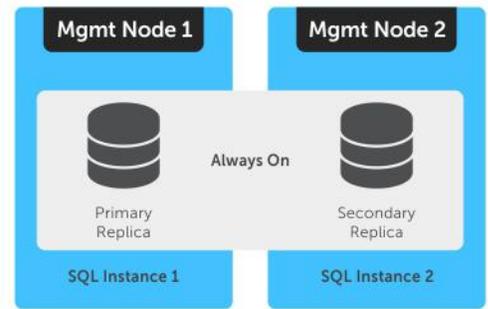
### 5.5.4 Horizon HCS High Availability

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.

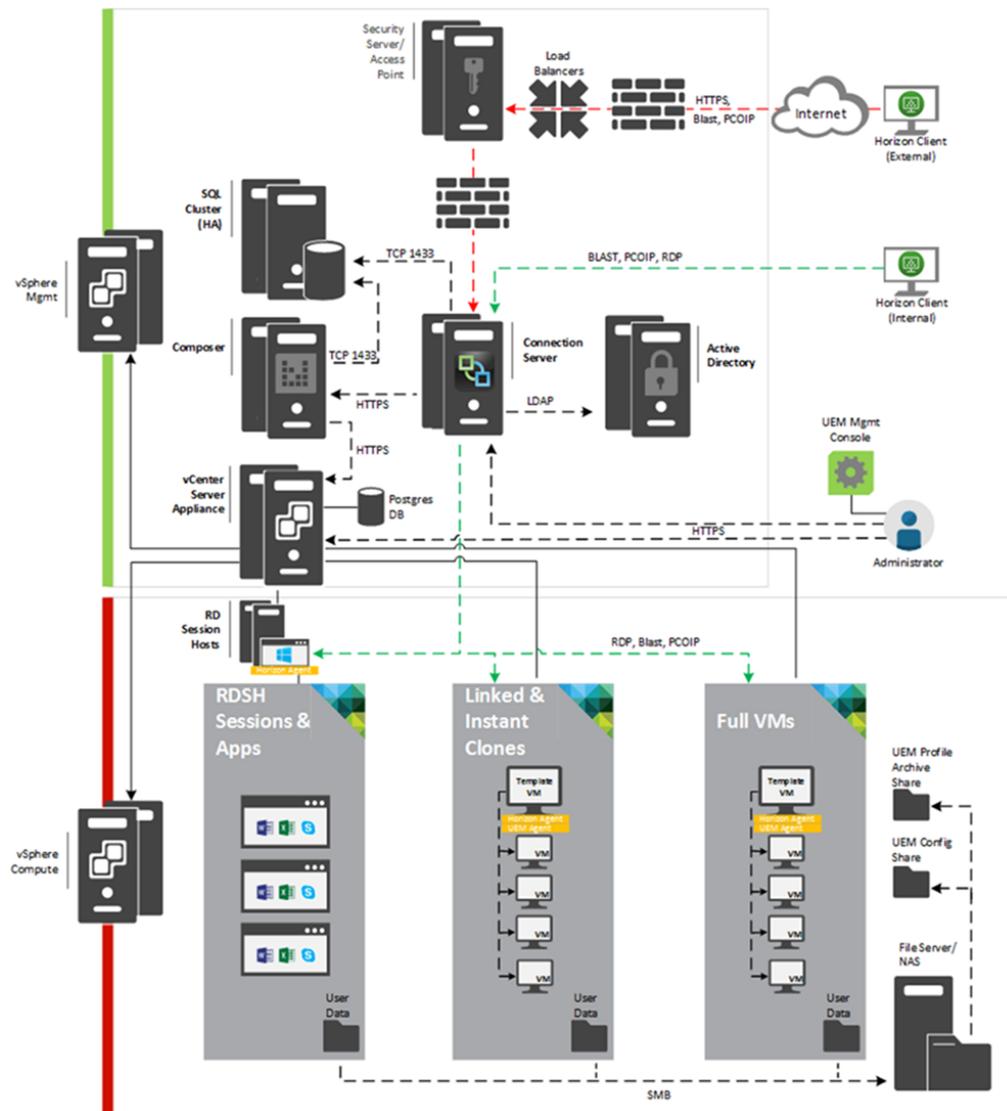
### 5.5.5 SQL Server High Availability

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 see these links for more information: [LINK1](#) and [LINK2](#)

### 5.6 VMware Horizon Communication Flow



## 6 Solution Performance and Testing

At the time of publication, these are the available density recommendations per appliance/node. Please refer to section 3.2. section for hardware specifications. 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 amount of 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.

User density summary

Hypervisor	Provisioning	Profile	Template OS	Config	User Density
ESXi 6.5	Linked Clone	Task	Windows 10	B5	150
ESXi 6.5	Linked Clone	Knowledge	Windows 10	B5	130
ESXi 6.5	Linked Clone	Power	Windows 10	B5	105
ESXi 6.5	Horizon 7 Apps/ RDSH	Task	W2K12R2	B5	200
ESXi 6.5	Linked Clone	Task	Windows 10	C7	250*
ESXi 6.5	Linked Clone	Knowledge	Windows 10	C7	170
ESXi 6.5	Linked Clone	Power	Windows 10	C7	140
ESXi 6.5	Horizon 7 Apps/ RDSH	Task	W2K12R2	C7	350
ESXi 6.5	Linked Clone/vGPU	Power	Windows 10	C7	64
ESXi 6.5	Linked Clone	Task	Windows 10	FX2/FC430 B5 (512GB Mem)	170
ESXi 6.5	Linked Clone	Knowledge	Windows 10	FX2/FC430 B5 (512GB Mem)	130
ESXi 6.5	Linked Clone	Power	Windows 10	FX2/FC430 B5 (512GB Mem)	105
ESXi 6.5	Linked Clone	Task	Windows 10	FX2/FC430 B5 (256GB Mem)	140
ESXi 6.5	Linked Clone	Knowledge	Windows 10	FX2/FC430 B5 (256GB Mem)	100
ESXi 6.5	Linked Clone	Power	Windows 10	FX2/FC430 B5 (256GB Mem)	80
ESXi 6.5	Linked Clone	Power	Windows 10	FX2/FC630 B5	70

The detailed results for the testing are in the following sections.

## 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

\*Turbo mode is enabled; therefore, the CPU threshold is increased as it will be reported as over 100% utilization when running with turbo.

## 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)	VSRN R630/730 AF/HY & VSRN FC430 AF
Hypervisor(s)	ESXi 6.5
Broker technology	Horizon 7.2
Broker database	Microsoft SQL 2016
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	1920 X 1080	Windows 10 Enterprise 64-bit
Graphics LVSI Custom – Density	2	4 GB	4GB	1920 X 1080	Windows 10 Enterprise 64-bit
Graphics LVSI Custom - Performance	4	8GB	8GB	1920 X 1080	Windows 10 Enterprise 64-bit

## 6.2.2 Platform Configuration

The hardware configuration details are listed in section 3.2.

## 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
B5	ESXi 6.0 U2	Horizon 7 & Linked Clone	Task Worker	150	68%	337GB	211GB	6
B5	ESXi 6.0 U2	Horizon 7 & Linked Clone	Knowledge Worker	130	88%	329GB	312GB	6.5
B5	ESXi 6.0 U2	Horizon 7 & Linked Clone	Power Worker	100	82%	372GB	296GB	7
B5	ESXi 6.5	Horizon 7 Apps/ RDSH	Task Worker	200	95%	155GB	69GB	2
C7	ESXi 6.5	Horizon 7 & Linked Clone	Task Worker	250	92%	505 GB	209 GB	7.1
C7	ESXi 6.5	Horizon 7 & Linked Clone	Knowledge Worker	170	83 %	495 GB	181 GB	6.6
C7	ESXi 6.5	Horizon 7 & Linked Clone	Power Worker	145	93%	498 GB	213 GB	7.2
C7	ESXi 6.5	Horizon 7 Apps/ RDSH	Task Worker	350	92%	211GB	162GB	5.7
C7	ESXi 6.5	Horizon 7 vGPU-M10-1B	Power Worker	64	48%	285GB	269GB	12
FC430 B5	ESXi 6.5	Horizon 7 & Linked Clone	Task Worker	150*	75%	335GB	130GB	11.6
FC430 B5	ESXi 6.5	Horizon 7 & Linked Clone	Knowledge Worker	130	95%	425GB	150GB	12.6
FC430 B5	ESXi 6.5	Horizon 7 & Linked Clone	Power Worker	105	95%	450GB	150GB	16.7
FC630 B5	ESXi 6.5	Horizon 7 & Linked Clone	Power Worker	70	49%	284GB	108GB	12.5

Note: \* We have increased this number to 170 for the FX2/FC430 B5 to bring this in line with the other results.

**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.

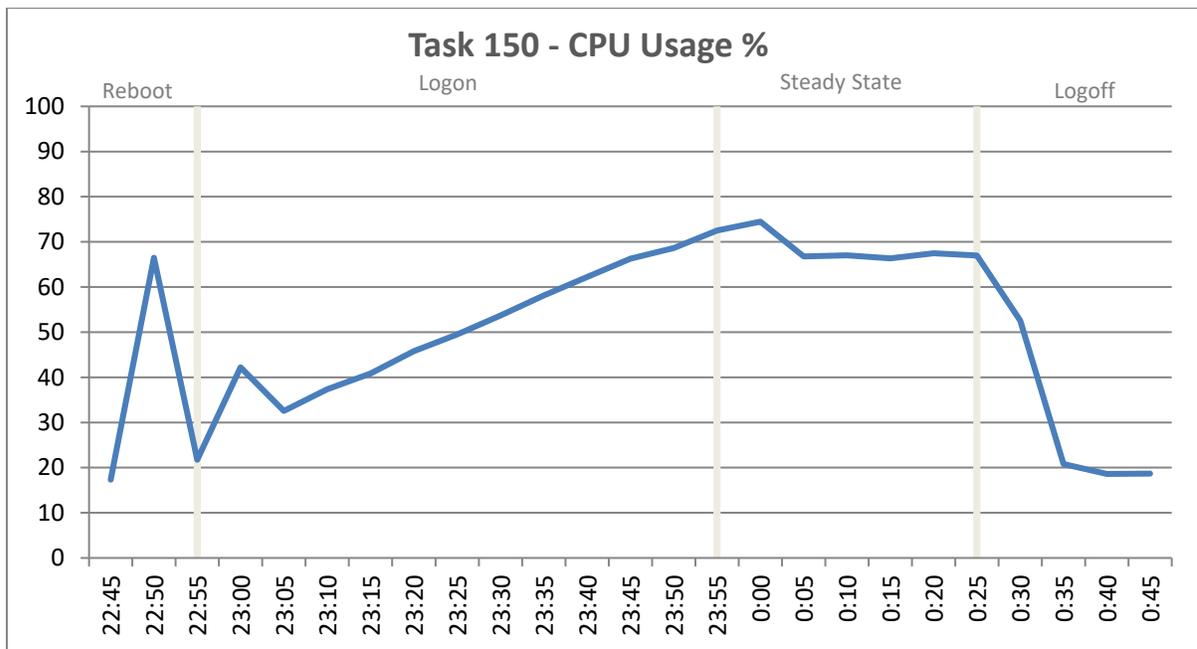
**Avg Net Mbps/User:** Amount of network usage over the steady state period divided by the number of users. For clusters, this is the combined average of all compute hosts over the steady state period divided by the number of users on a host.

### 6.3.1 vSRN-B5

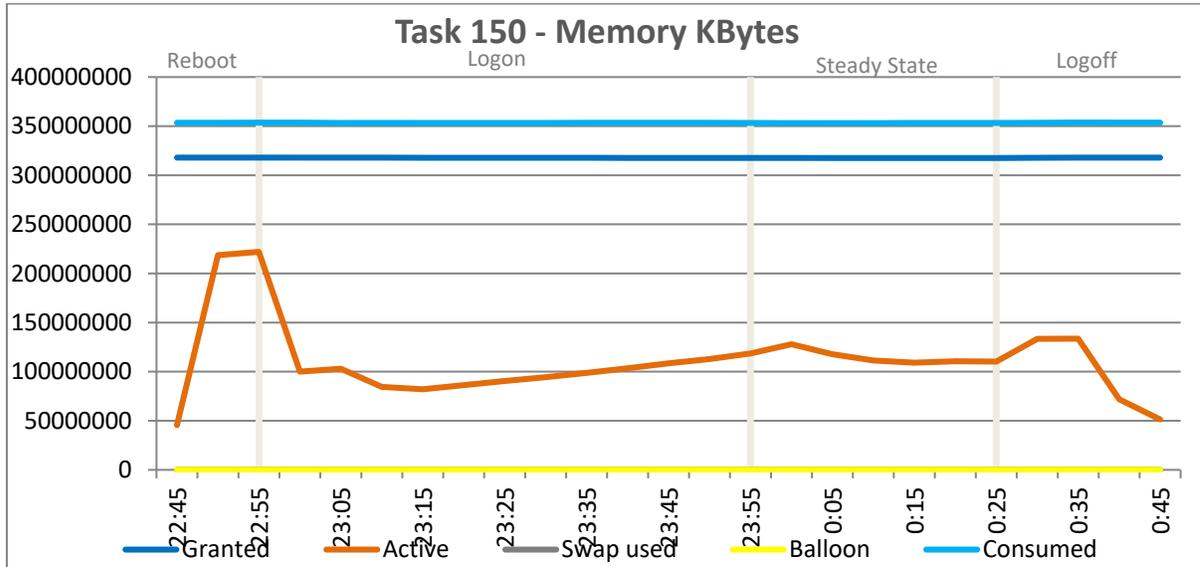
Refer to Section 3.2 for hardware configuration details .This would be classed as a vSRN HY-6 (Hybrid) or AF-4 (All-Flash) configuration on the vSRN HCL

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

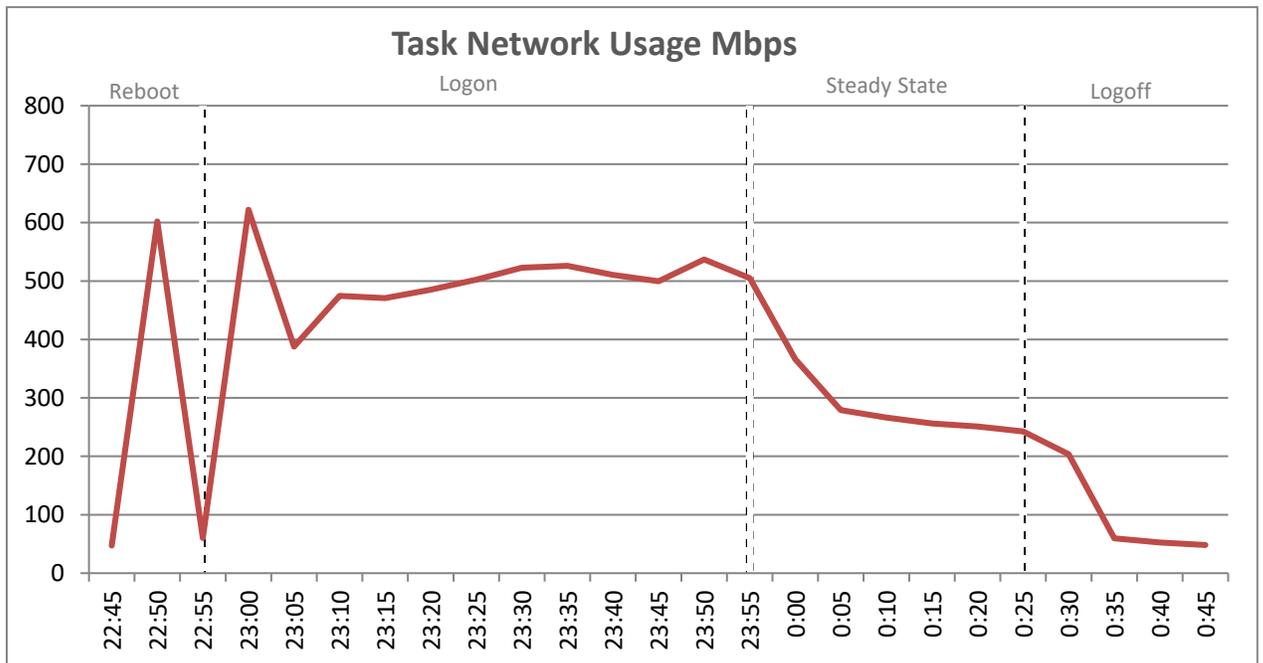
Maximum CPU utilization was in the region of the 75% threshold indicating the number of users tested were less than the 95% threshold so there is more room for user for this config/workload and will be revisited in a future revision.



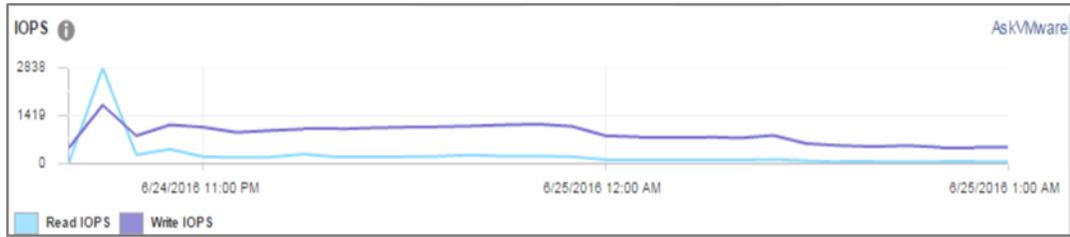
Memory usage is monitored on ESXi host, memory usage monitored are consumed, active, balloon and swap used, as swap and ballooning usage would indicate host memory reached saturation point and the VMperformance may start deteriorating.



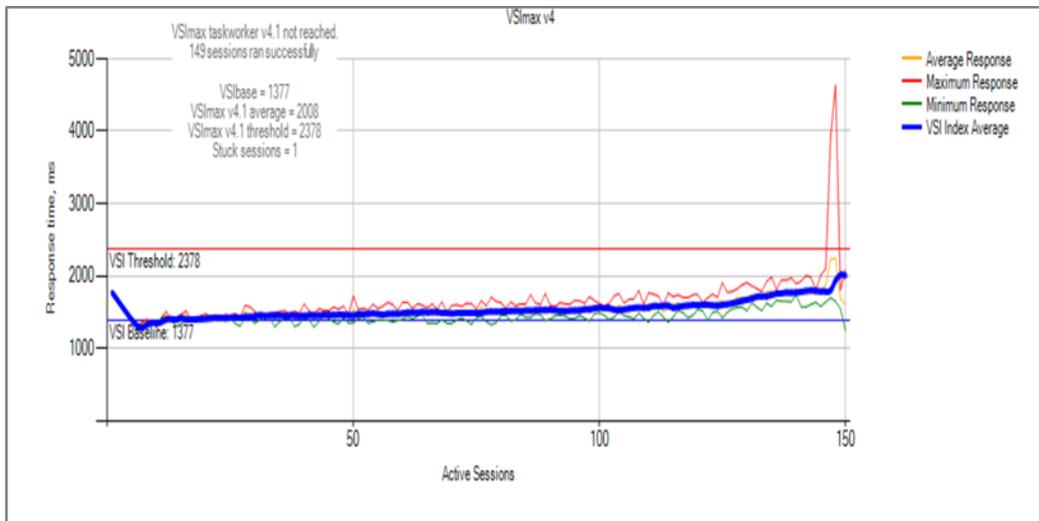
There were no issues with network usage on any of the test runs. All tests showed that the reboot of the VM pool before beginning Login VSI testing produced the highest spike in network activity.



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.

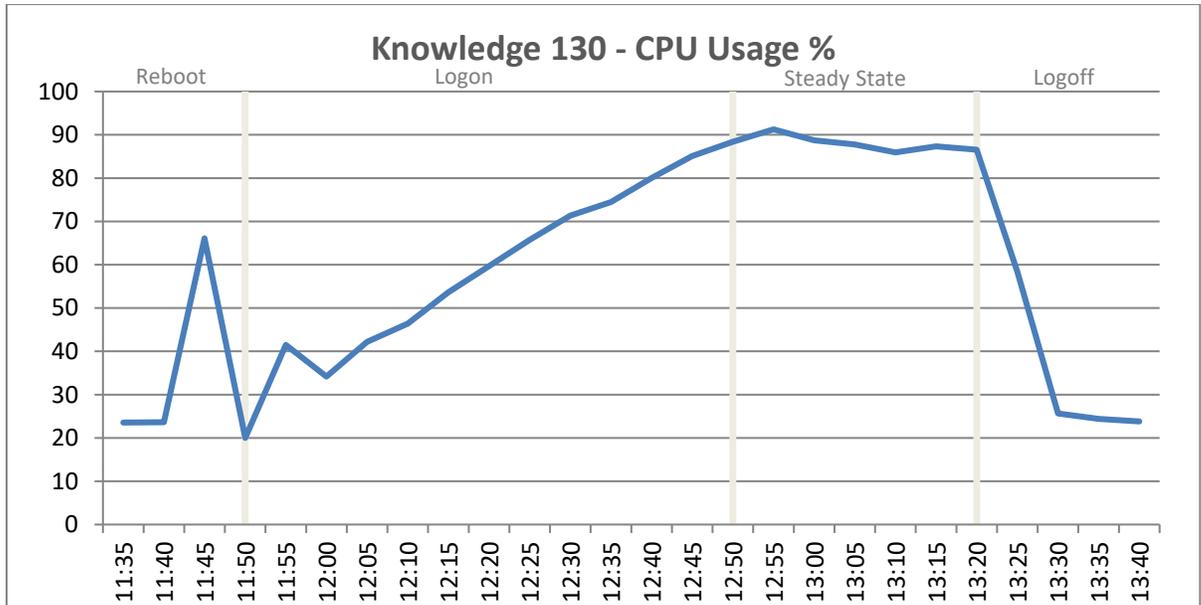


The VSI Max was not reached as shown below.

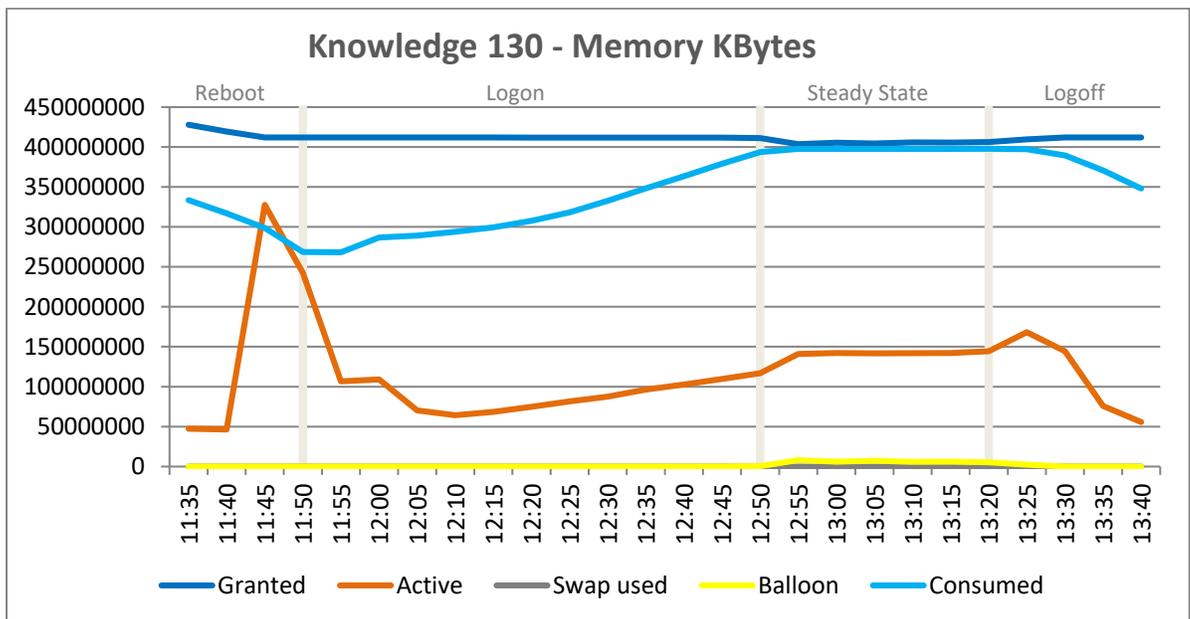


### 6.3.1.2 Knowledge Worker, ESXi 6.0U2, Horizon 7 linked clones

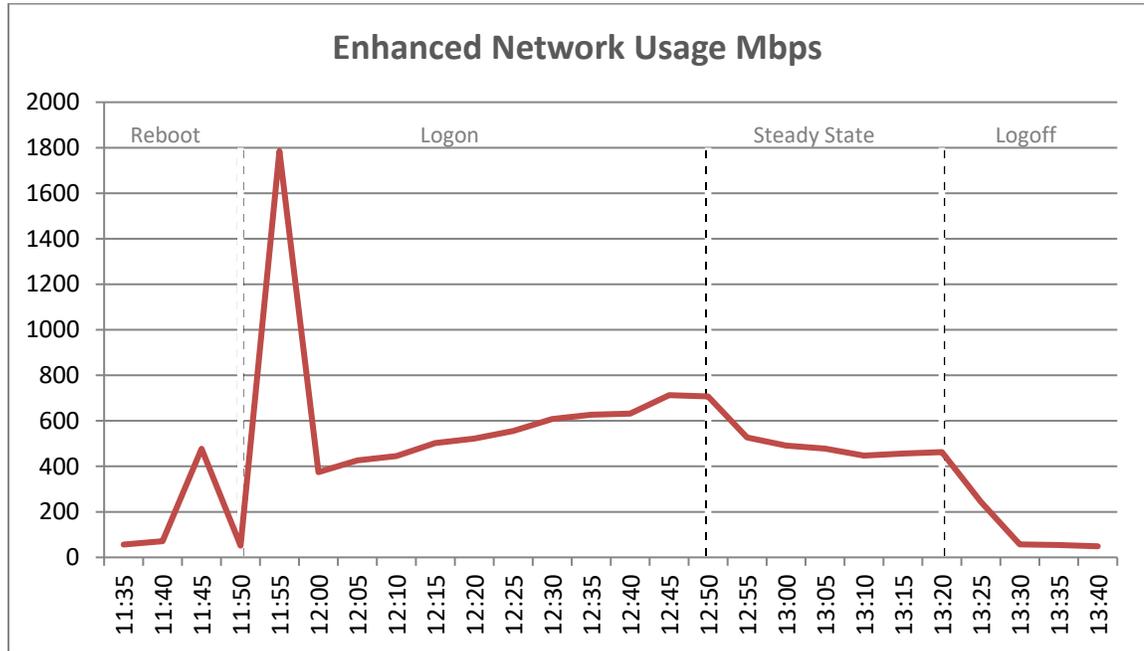
Maximum CPU utilization was in the region of the 85% threshold indicating the number of users tested was within the appropriate PAAC resource threshold.



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.

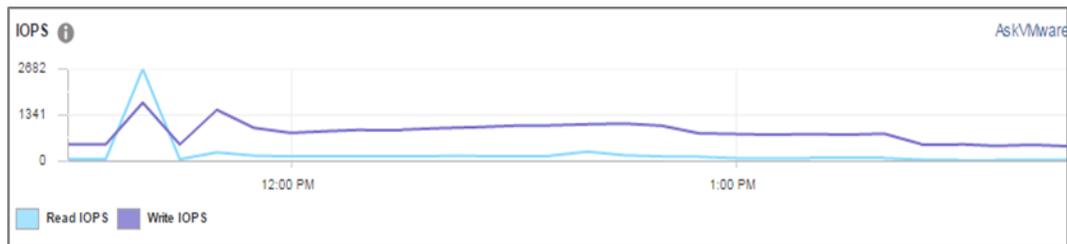


There were no issues with network usage on any of the test runs. All tests showed that the reboot of the VM pool before beginning Login VSI testing produced the highest spike in network activity. There is a significant reduction in activity once the steady state phase is reached after all machines have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

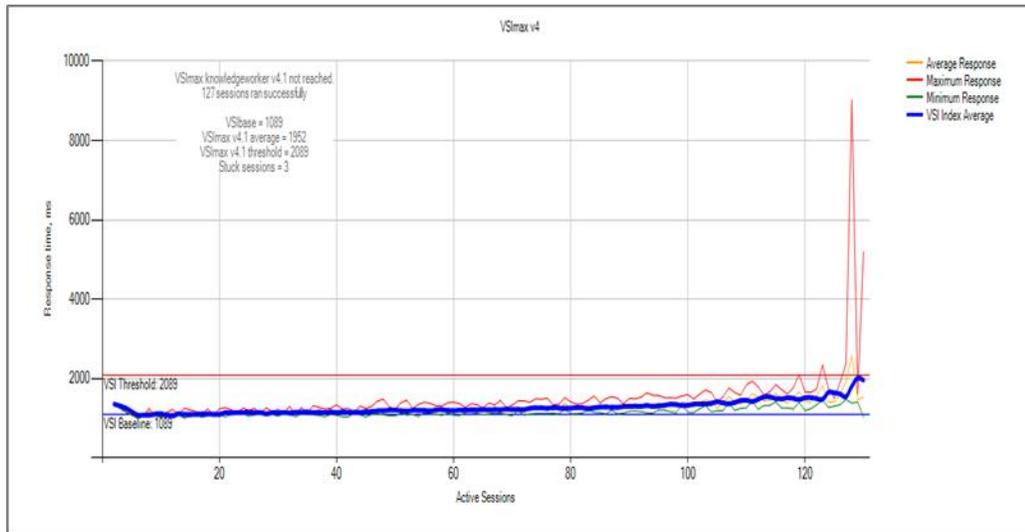


For the logon and steady state phase of each test, the latency remained well below the 20ms threshold reaching a max of 2-3 ms during the test run. The IOPS peaked during the boot phase and for each profile test and then settled thereafter during the login phase and reduced once steady state was reached.

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.

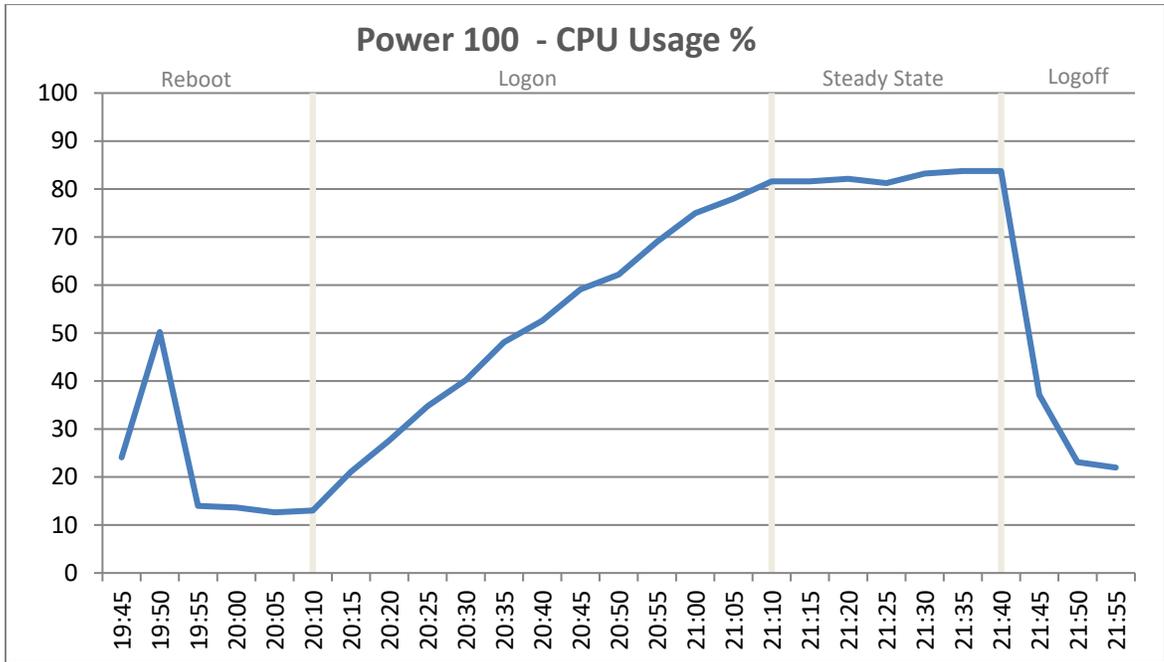


The VSI Max was not reached as shown below.

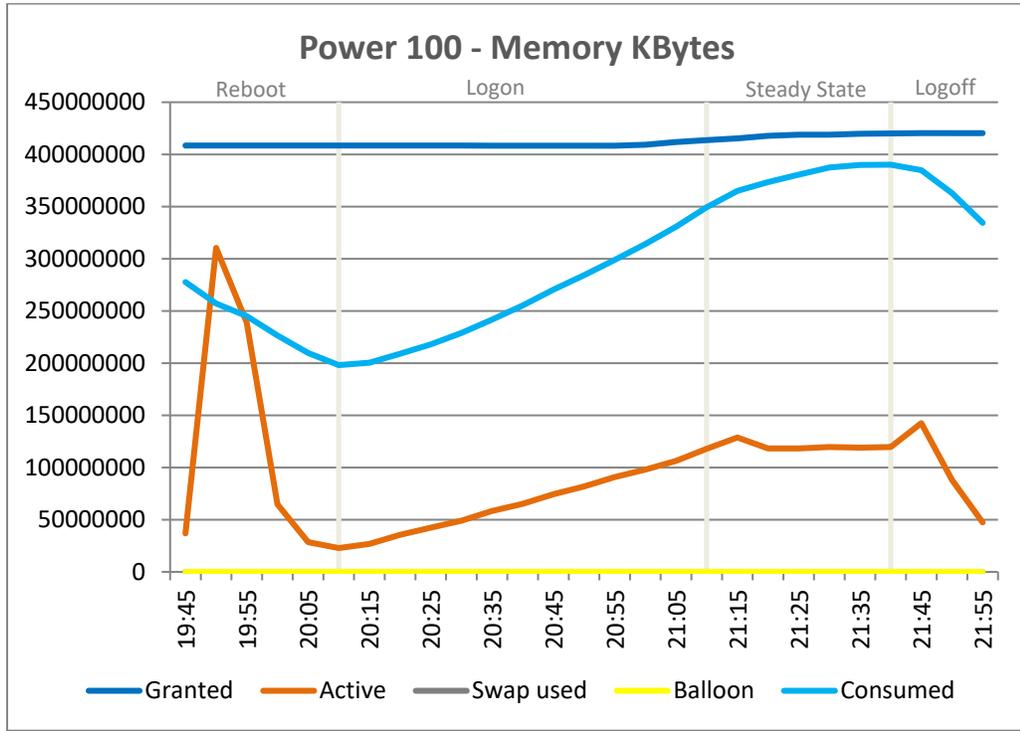


### 6.3.1.3 Power Worker, 100 users, ESXi 6.0U2, Horizon 7 linked clones

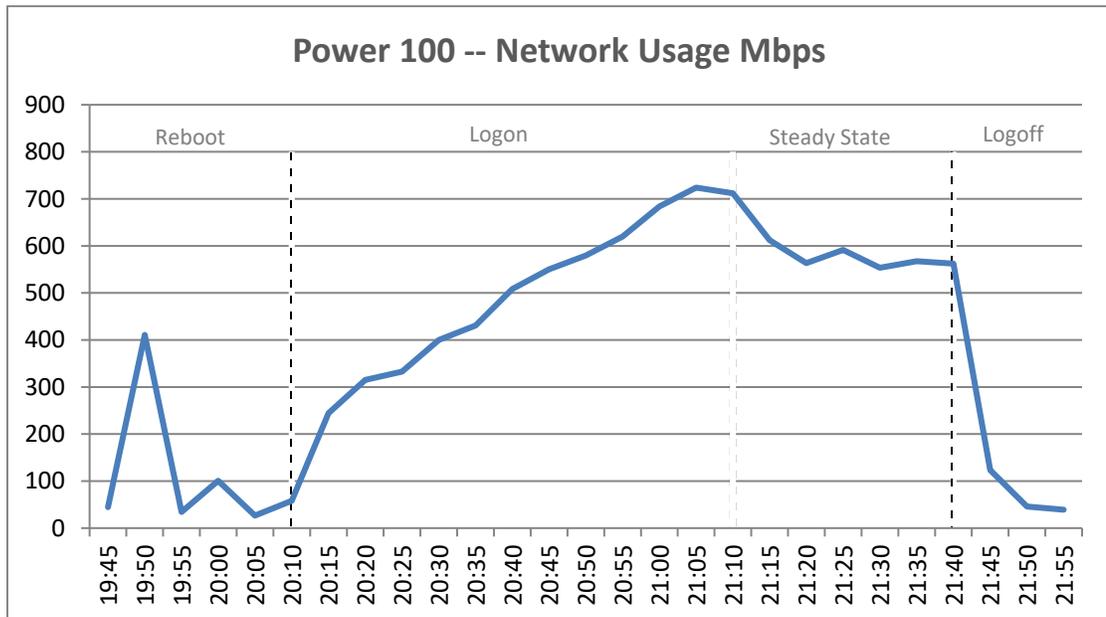
Maximum CPU utilization was in the region of the 85% threshold indicating the number of users tested was within the appropriate PAAC resource threshold.



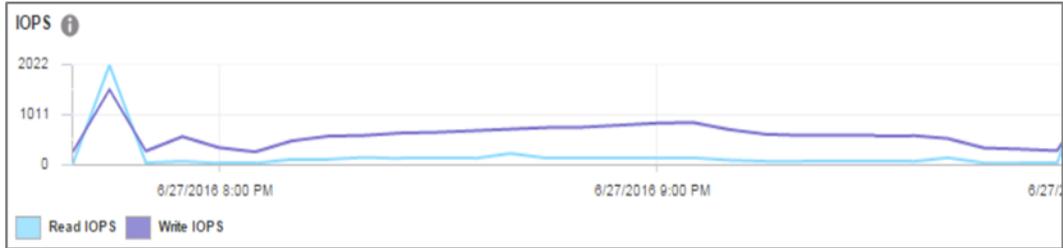
Memory usage is monitored on ESXi host, memory usage monitored are consumed, active, balloon and swap used, as swap and ballooning usage would indicate host memory reached saturation point and the VM performance may start deteriorating



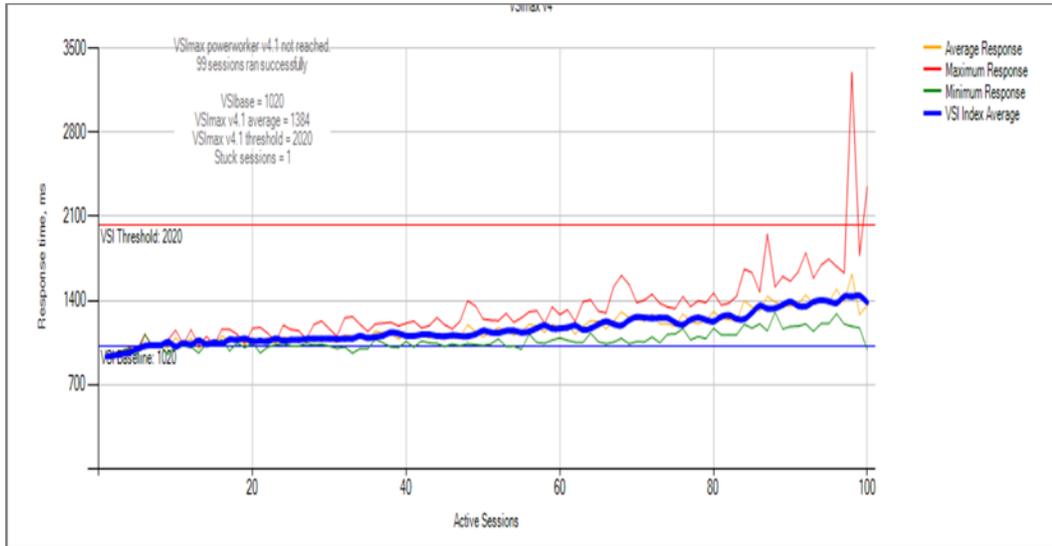
There were no issues with network usage on any of the test runs. All tests showed that the reboot of the VM pool before beginning Login VSI testing produced the highest spike in network activity



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.

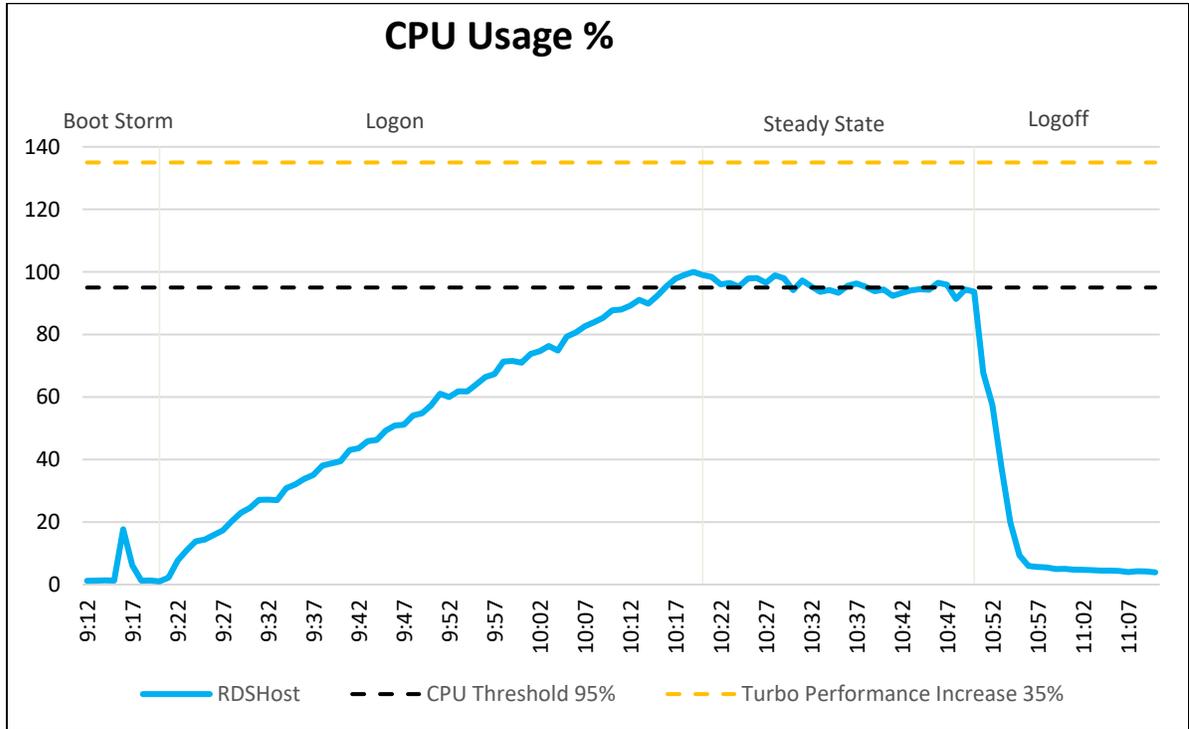


The VSI Max was not reached as shown below.

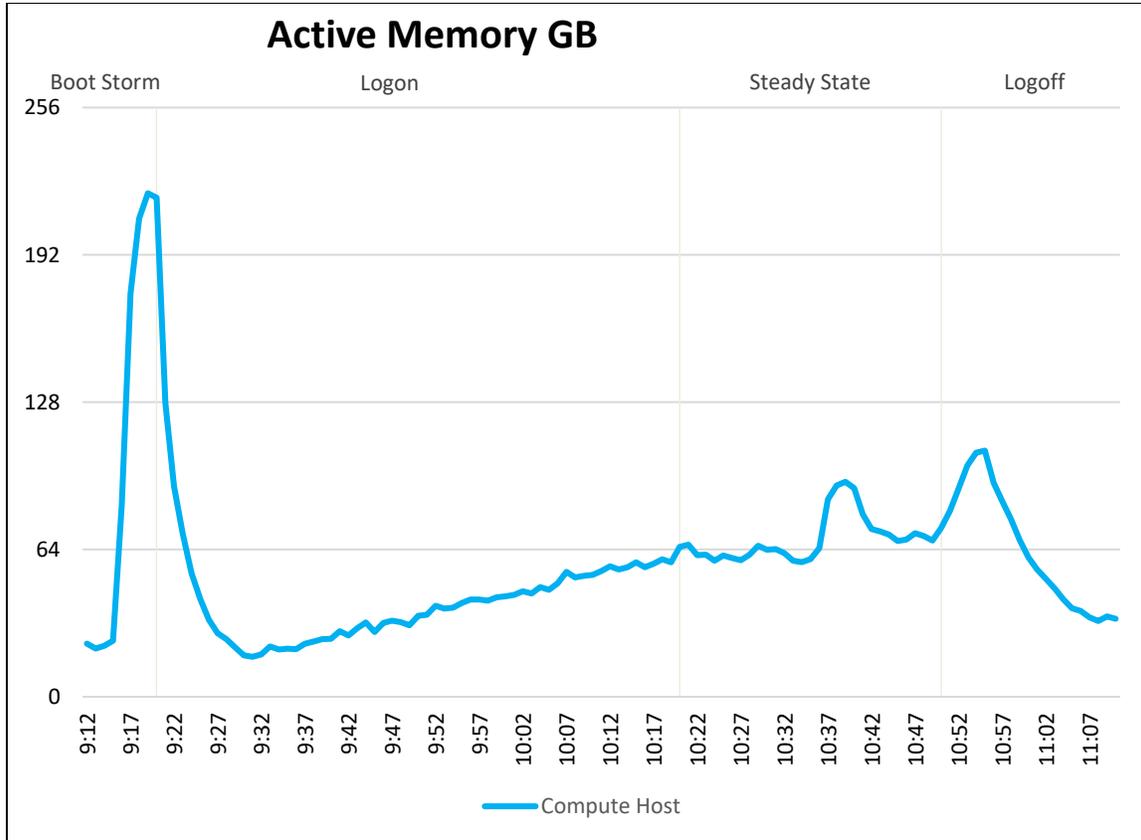


### 6.3.1.4 Task Worker, 200 users, ESXi 6.5 Horizon Apps/RDSH

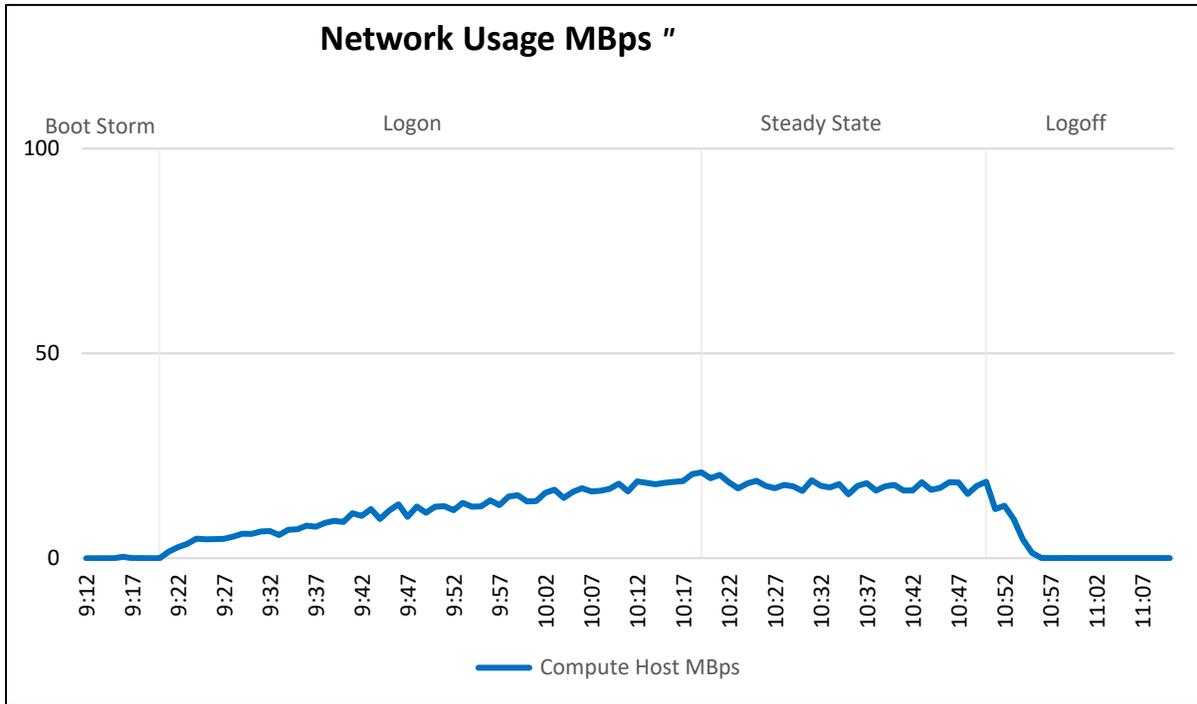
The below graph shows the performance data of the physical compute hosts for 205 user sessions distributed over 7 vRDS hosts. The CPU reaches a steady state average of 95.4% on the compute host during the test cycle when 205 sessions are logged on to 8 vRDS hosts.



In regards to memory consumption for the compute host, out of a total of 384 GB available memory never reached maximum. The compute host reached a maximum memory consumption of 168 GB with active memory usage reaching a max of 218 GB. Average memory consumption of 155 GB and active memory usage 68 GB during steady state period. Zero ballooning occurred, and no memory swapping on compute host.

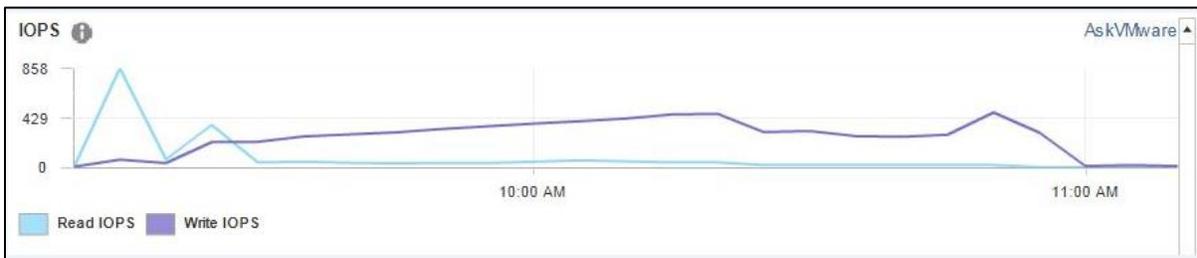


Network bandwidth is not an issue on this test run with a steady state average of approximately 18 MBps and a Max of 21 MBps over the duration of the test.

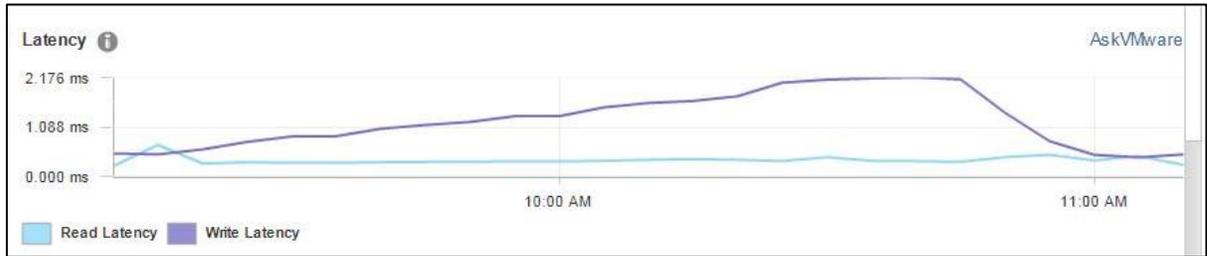


The IOPS graphs and IOPS numbers are taken from the vCenter web console and the graphs clearly show the initial boot storm during a reboot of all the desktops, followed by a settle period, logon, steady state and logoff phases. The graph displays the Disk IOPS figure for the Compute Host.

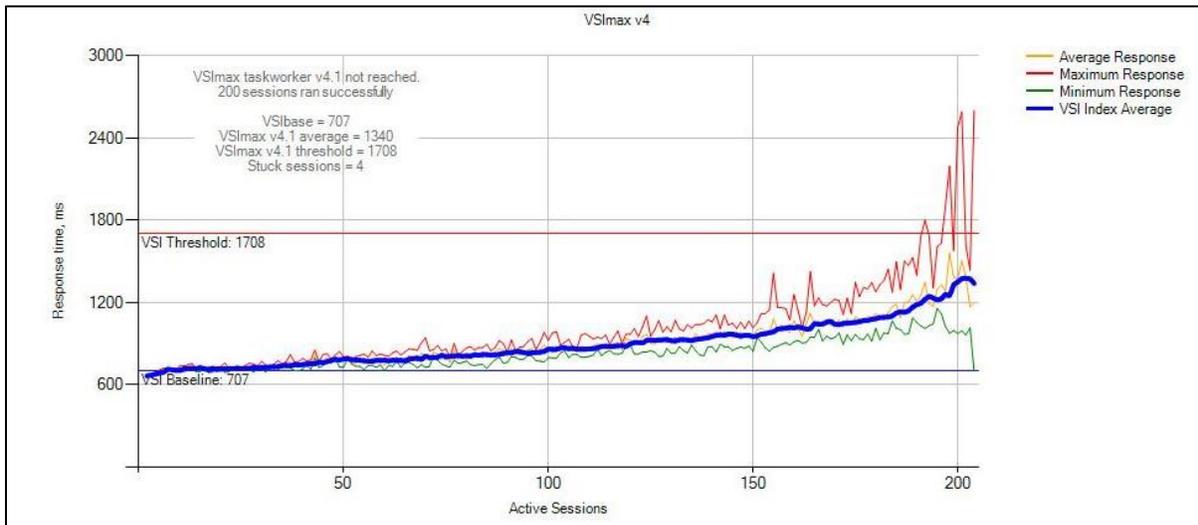
The Host reached a maximum of 926 Disk IOPS during the boot storm before test start and 514 IOPS at the start of steady state.



Disk I/O Latency was not an issue. The maximum latency reached was approximately .7 ms during the boot storm and 1.8 ms at the beginning of steady state. This was well below the 20 ms threshold that is regarded as becoming potentially troublesome.



The Login VSI Max user experience score shown below for this test was not reached indicating there was little deterioration of user experience during testing and manually interacting with the test sessions confirmed this.



Notes:

- As indicated above, the CPU graphs do not take into account the extra 35% of CPU resources available through the E5-2660v4’s turbo feature.
- Main constraint during testing proved to be CPU.
  - CPU Usage averaged above 95% usage during steady state.

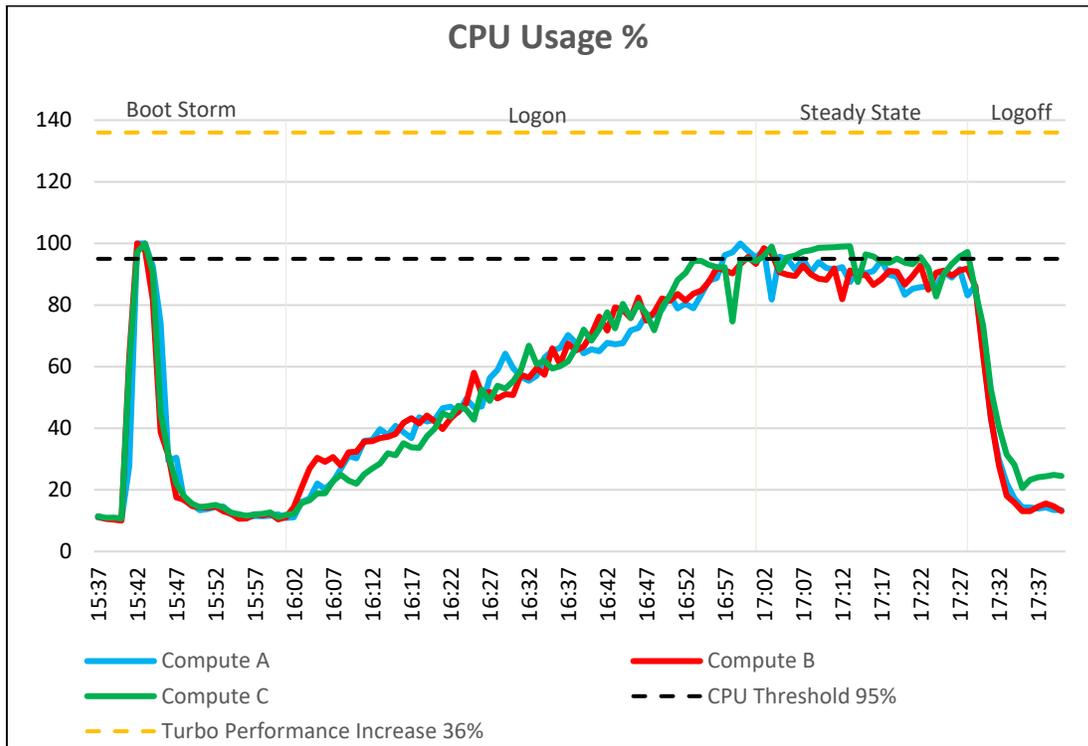
### 6.3.2 vSRN-C7

Refer to Section 3.2 for hardware configuration details. This would be classed as a vSRN HY-8 (Hybrid) or AF-8 (All-Flash) configuration on the vSRN HCL

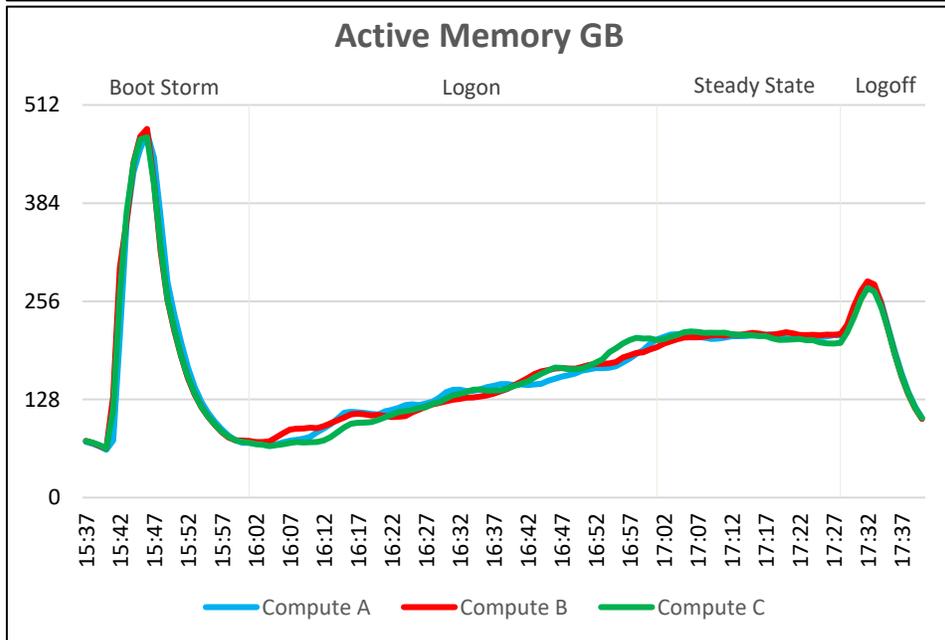
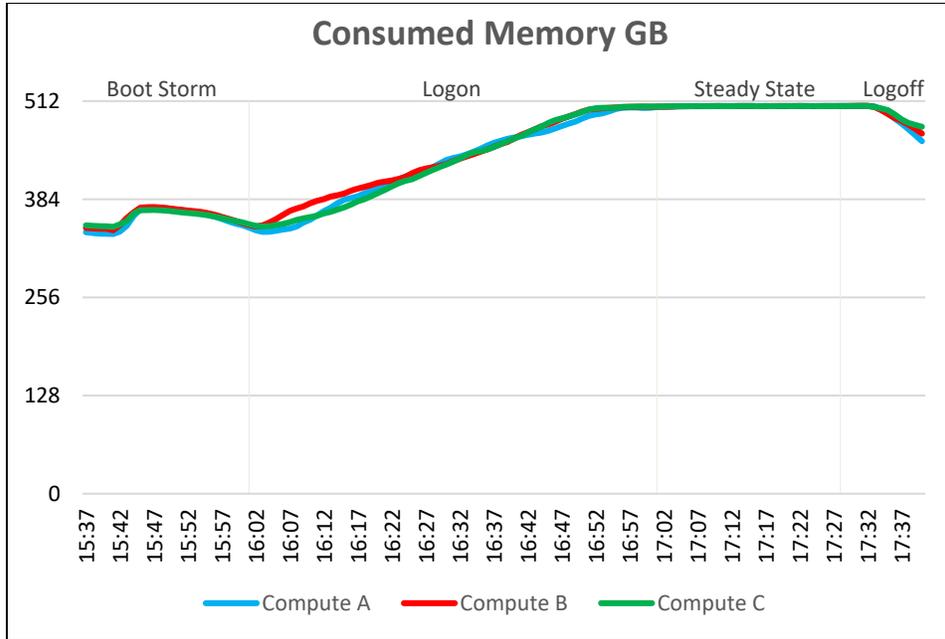
### 6.3.2.1 Task Worker, 250 users per host, ESXi 6.5, Horizon 7. linked clones

Each compute host was populated with 250 virtual machines per host. With all user virtual machines powered on and before starting test, the CPU usage was approximately 11%.

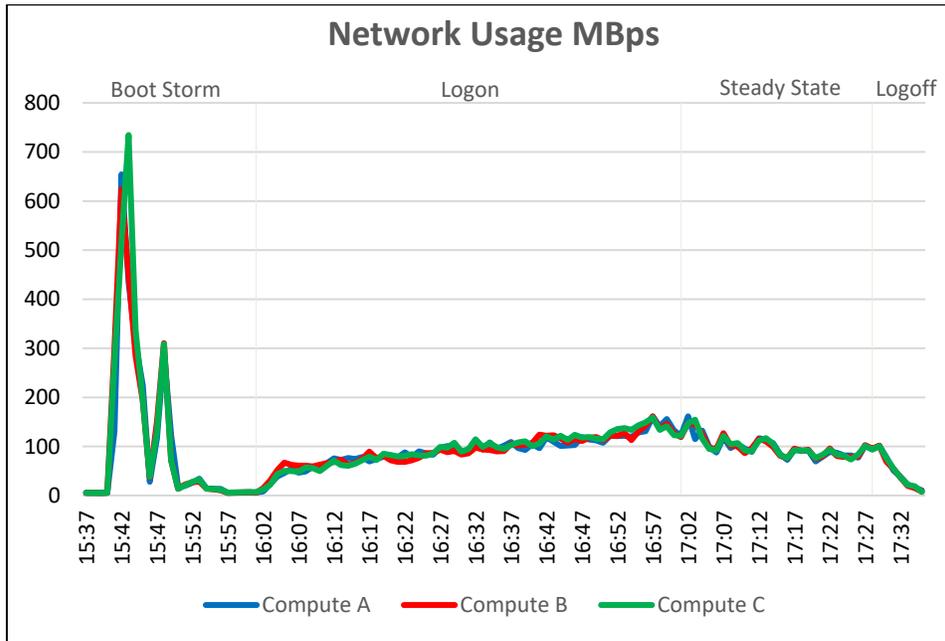
The below graph shows the performance data for 250 user sessions per host. The CPU reaches a steady state average of 92% across the three compute hosts during the test cycle when 250 users are logged on to each host.



In regards to memory consumption for the cluster, out of a total of 512 GB available memory per node, memory usage was pushed close to its maximum. The compute hosts reached a maximum memory consumption of 505 GB with active memory usage reaching a max of 480 GB. Some ballooning occurred on all hosts over a period of about 40 minutes, mostly during the steady state period. The maximum amount of memory ballooning was approximately 17.7GB. There was no memory swapping on any of the hosts.

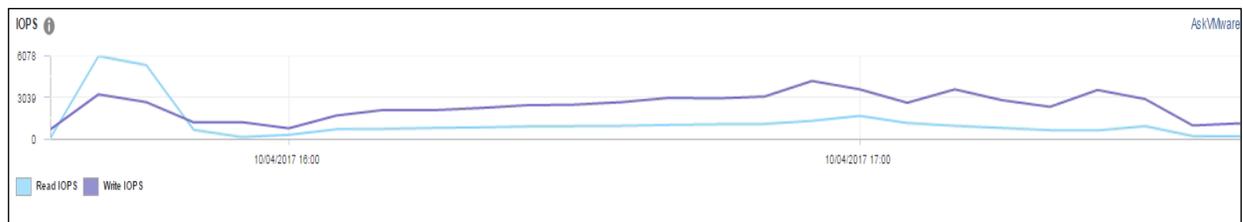


Network bandwidth is not an issue on this test run with a steady state peak of approximately 154 MBps. One of the hosts reached a peak of 734 MBps during the boot storm.

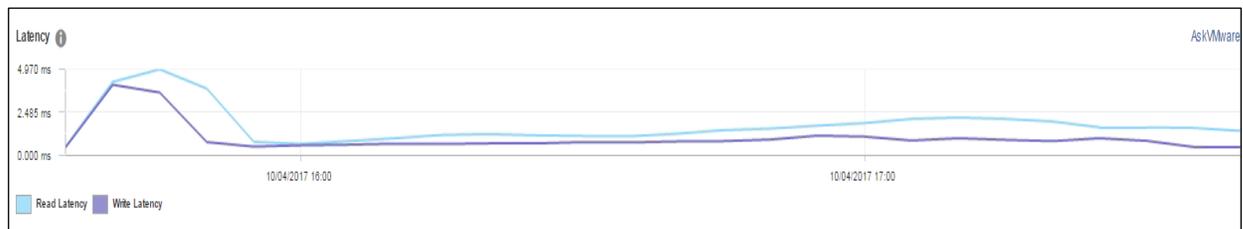


The IOPS graphs and IOPS numbers are taken from the vCenter web console and the graphs clearly show the initial boot storm during a reboot of all the desktops, followed by a settle period, logon, steady state and logoff phases. The graph displays the Disk IOPS figure for the VSAN cluster.

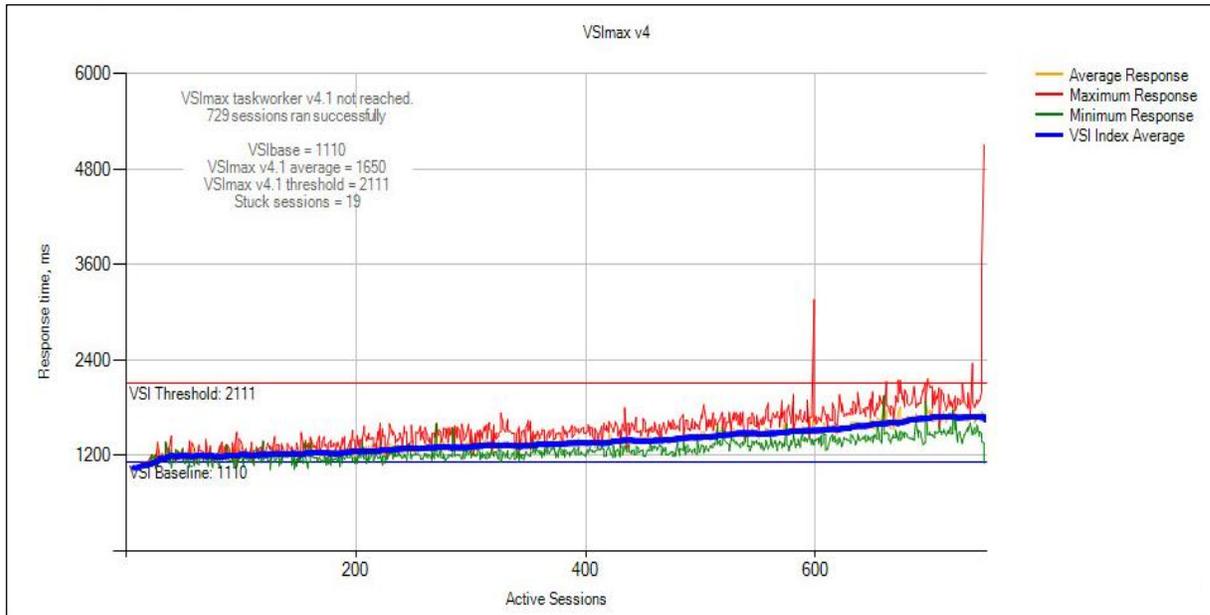
The cluster reached a maximum of 9,342 Disk IOPS during the boot storm before test start and 5,328 IOPS at the start of steady state.



Disk I/O Latency was not an issue. The maximum latency reached was approximately 8.4 ms during the boot storm and 2.9 ms at the beginning of steady state. This was well below the 20 ms threshold that is regarded as becoming potentially troublesome.



The Login VSI Max user experience score shown below for this test was not reached indicating there was little deterioration of user experience during testing and manually interacting with the test sessions confirmed this.



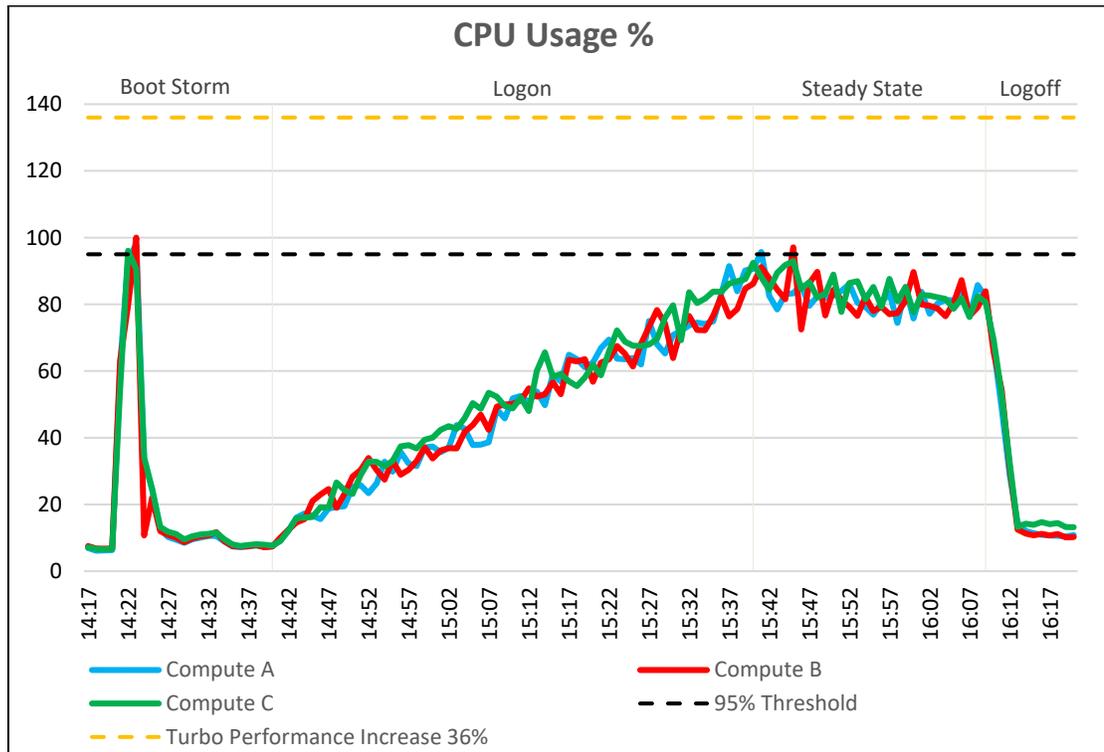
**Notes:**

- As indicated above, the CPU graphs do not take into account the extra 36% of CPU resources available through the E5-2698v4's turbo feature.
- 512 GB of memory installed on each node is just about sufficient for the number of desktops. With memory usage going close to maximum, no extra desktops could have been accommodated for this configuration.
- With extra memory installed there would have been capacity for some extra desktops as the CPU steady state average did not reach the 95% target threshold.

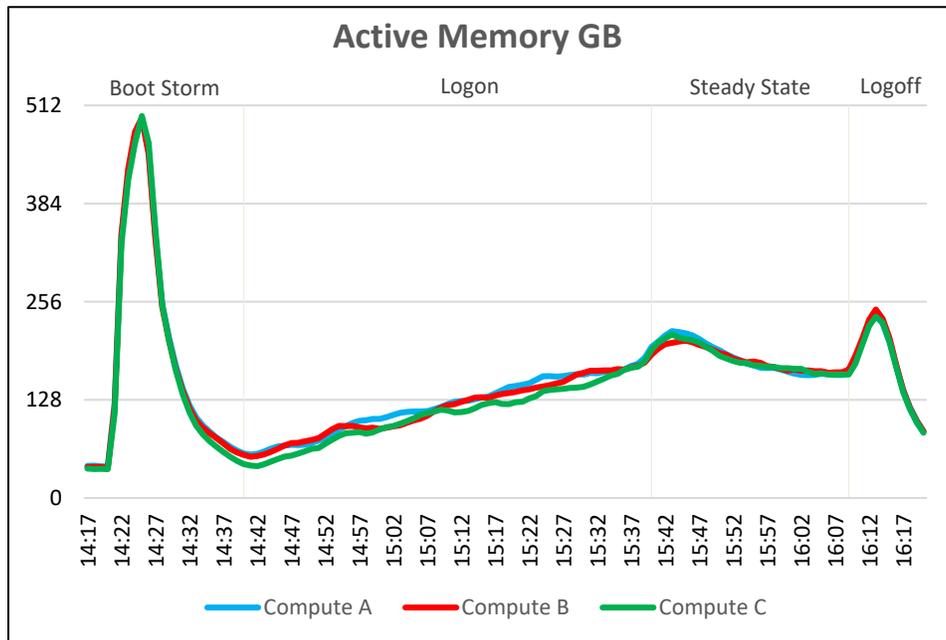
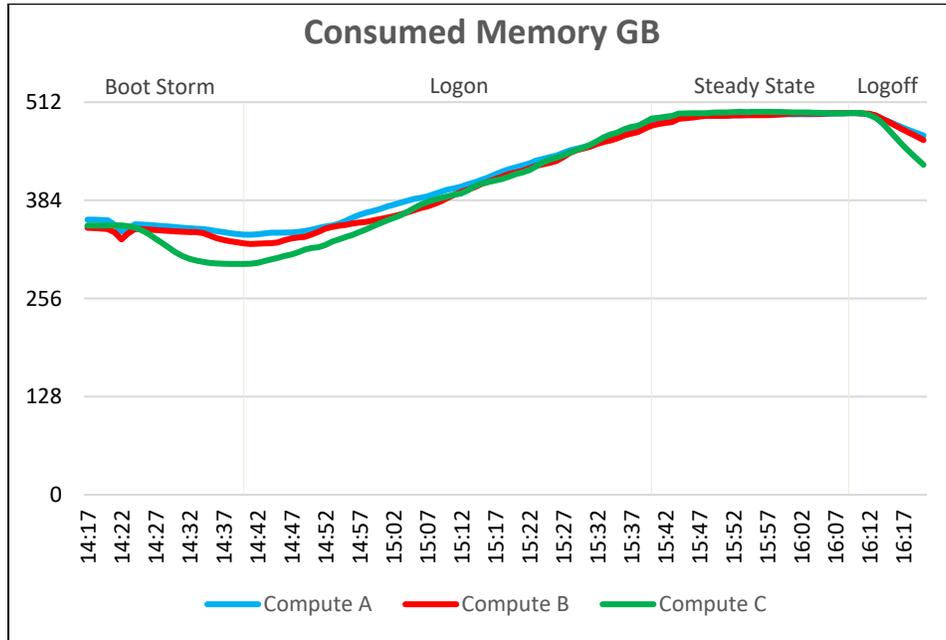
### 6.3.2.2 Knowledge Worker, 170 users per host, ESXi 6.5, Horizon 7 linked clones

Each compute host was populated with 170 virtual machines per host. With all user virtual machines powered on and before starting test, the CPU usage was approximately 7%.

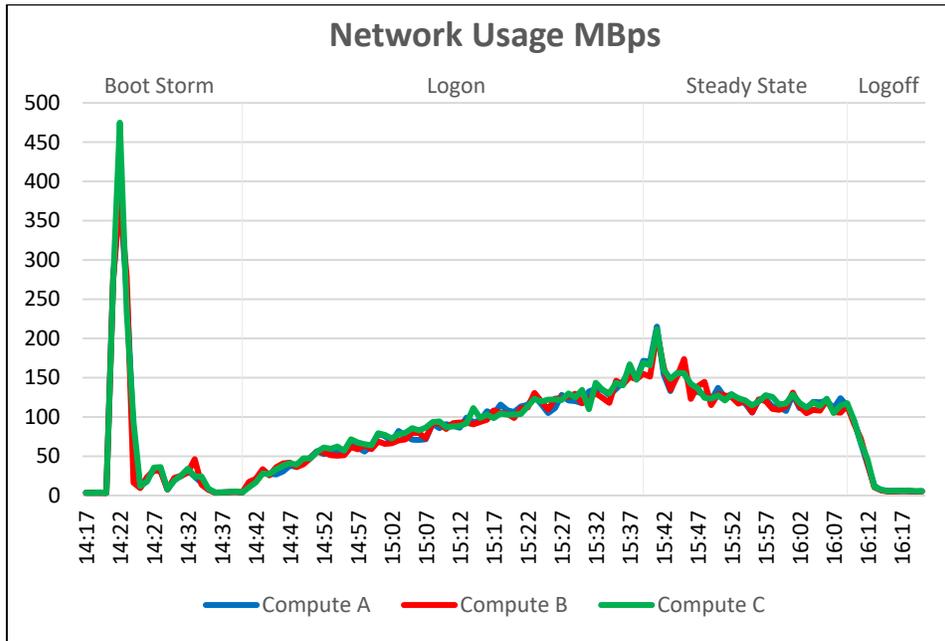
The below graph shows the performance data for 170 user sessions per host. The CPU reaches a steady state average of 82.7% across the three compute hosts during the test cycle when 170 users are logged on to each host.



In regards to memory consumption for the cluster, out of a total of 512 GB available memory per node, memory usage was pushed close to its maximum. The compute hosts reached a maximum memory consumption of 499 GB with active memory usage reaching a max of 498 GB during the boot storm. Some ballooning occurred on all hosts over a period of about 40 minutes, mostly during the steady state period. The maximum amount of memory ballooning was approximately 13 GB. There was no memory swapping on any of the hosts.

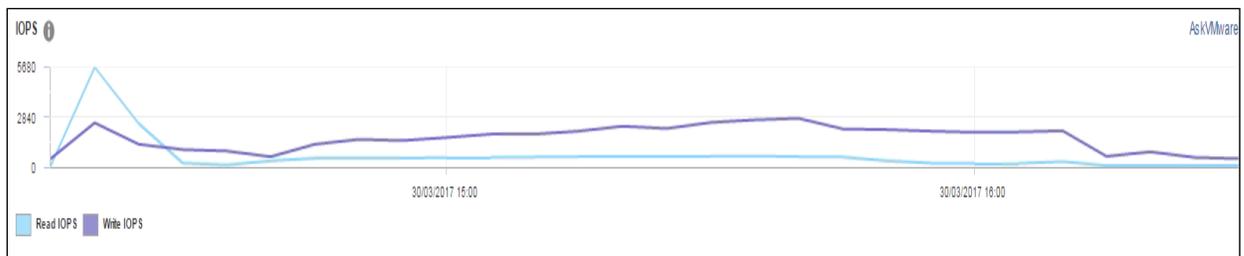


Network bandwidth is not an issue on this test run with a steady state peak of approximately 215 MBps. One of the hosts reached a peak of 474 MBps during the boot storm.

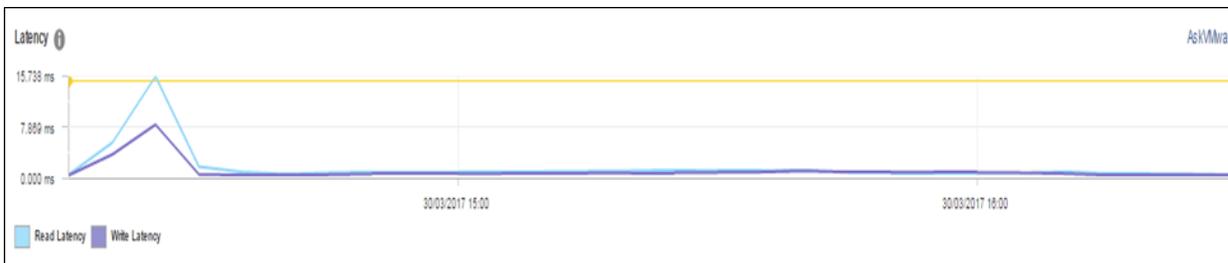


The IOPS graphs and IOPS numbers are taken from the vSphere web console and the graphs clearly show the initial boot storm during a reboot of all the desktops, followed by a settle period, logon, and steady state and logoff phases. The graph displays the Disk IOPS figure for the VSAN cluster.

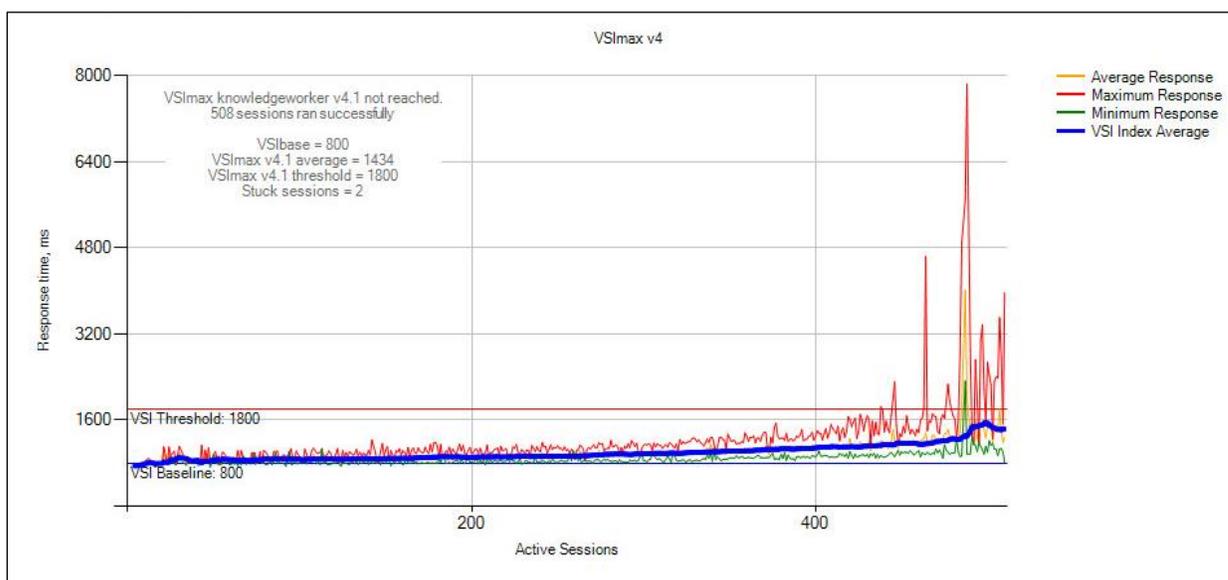
The cluster reached a maximum of 8,201 Disk IOPS during the boot storm before test start and 3,377 IOPS at the start of steady state.



Disk I/O Latency was not an issue. The maximum latency reached was approximately 2.3 ms during the steady state period of test and a brief spike to 24 ms during the boot storm. A spike list this, while over the 20 ms threshold should not be an issue as it occurred during a boot storm and latency returned to normal levels during testing.



The Login VSI Max user experience score shown below for this test was not reached indicating there was little deterioration of user experience during testing and manually interacting with the test sessions confirmed this.



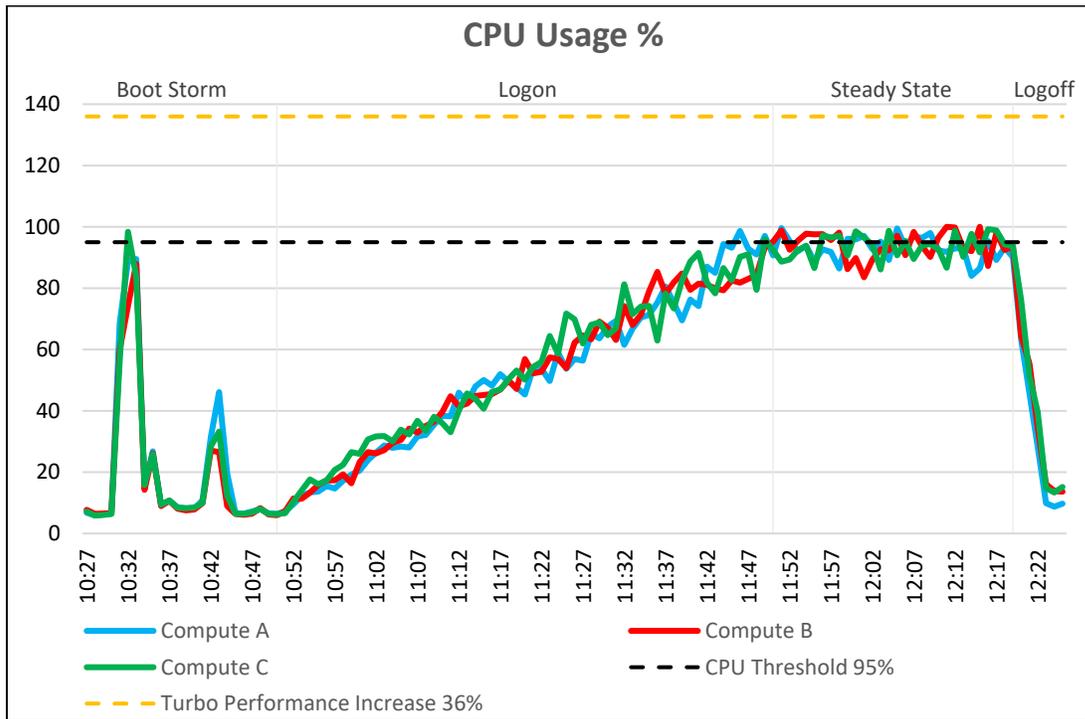
**Notes:**

- As indicated above, the CPU graphs do not take into account the extra 36% of CPU resources available through the E5-2698v4's turbo feature.
- 512 GB of memory installed on each node is just about sufficient for the number of desktops. With memory usage going close to maximum, no extra desktops could have been accommodated for this configuration.
- With extra memory installed there would have been capacity for some extra desktops as the CPU steady state average was well short of the 95% target threshold.

### 6.3.2.3 Power Worker, 145 users per host, ESXi 6.5, Horizon 7 linked clones

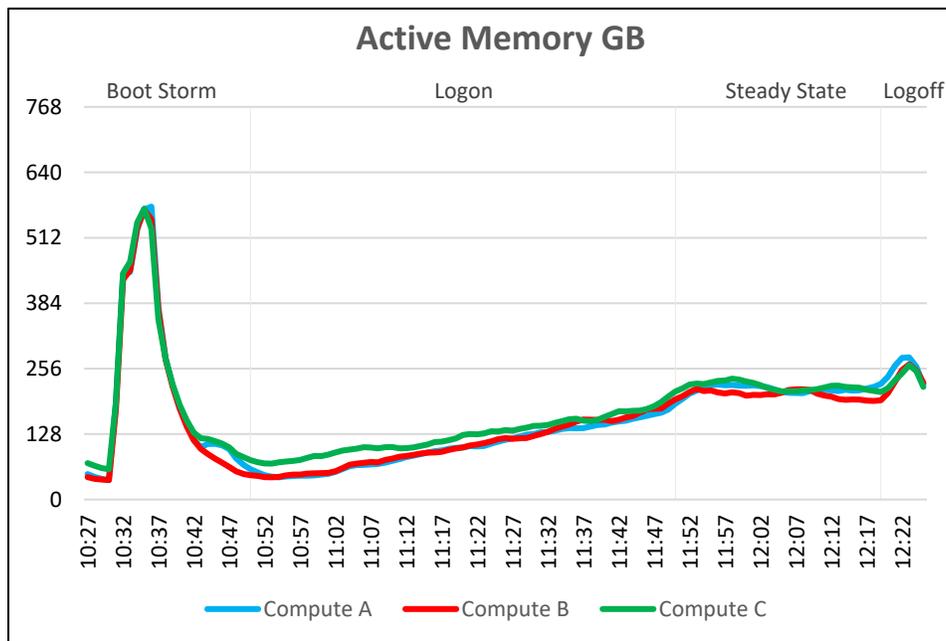
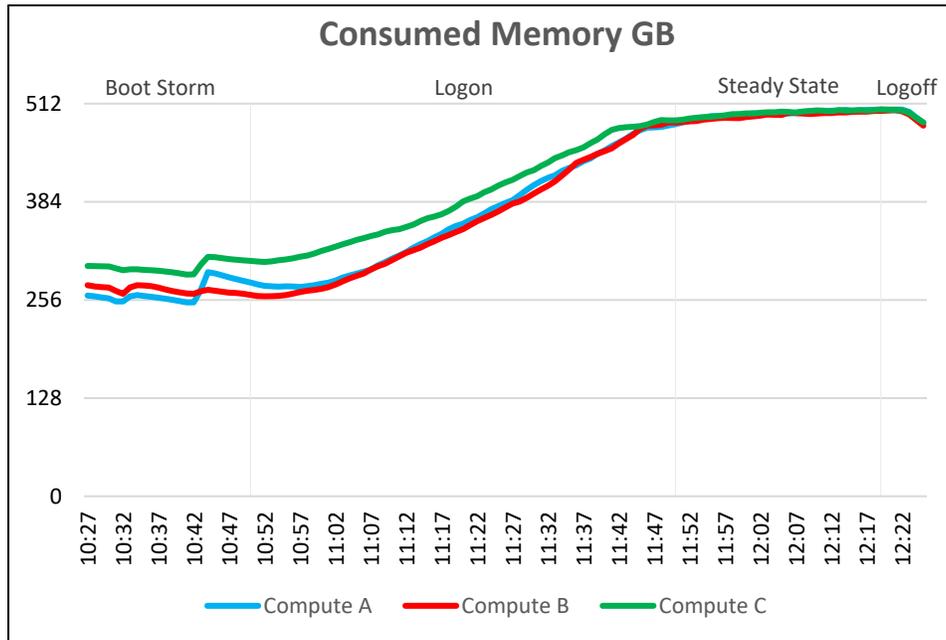
Each compute host was populated with 145 virtual machines per host. With all user virtual machines powered on and before starting test, the CPU usage was approximately 6%.

The below graph shows the performance data for 145 user sessions per host. The CPU reaches a steady state average of 93.6% across the three compute hosts during the test cycle when 145 users are logged on to each host.

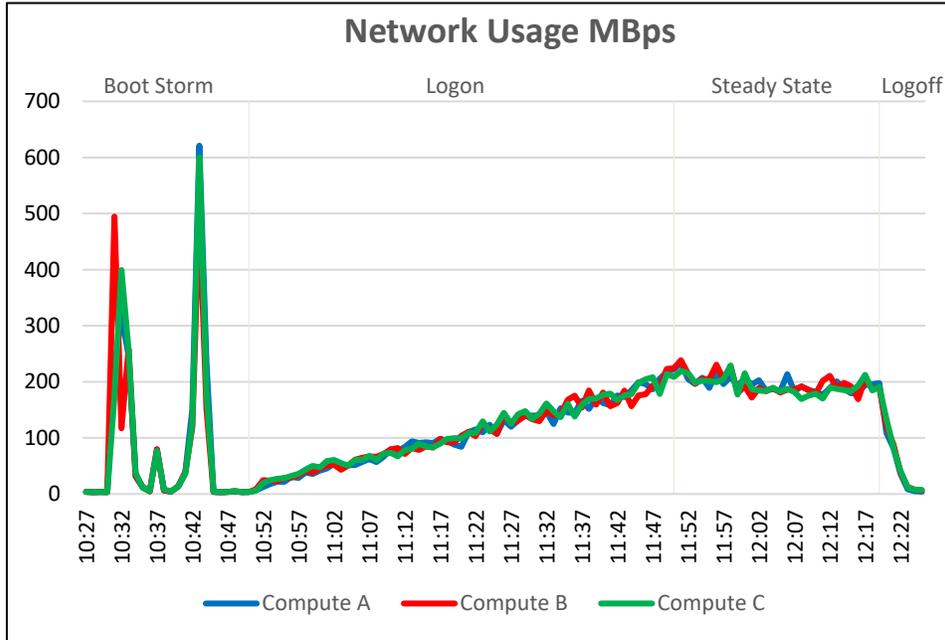


In regards to memory consumption for the cluster, out of a total of 512 GB available memory per node, memory usage was pushed close to its maximum. The compute hosts reached a maximum memory consumption of 503 GB with active memory usage reaching a max of 572 GB during the boot storm.

Some ballooning occurred on all hosts over a period of about 40 minutes, mostly during the steady state period. The maximum amount of memory ballooning was approximately 25 GB. There was a small amount of memory swapping on one of the hosts which reached a maximum of 0.00818 GB.

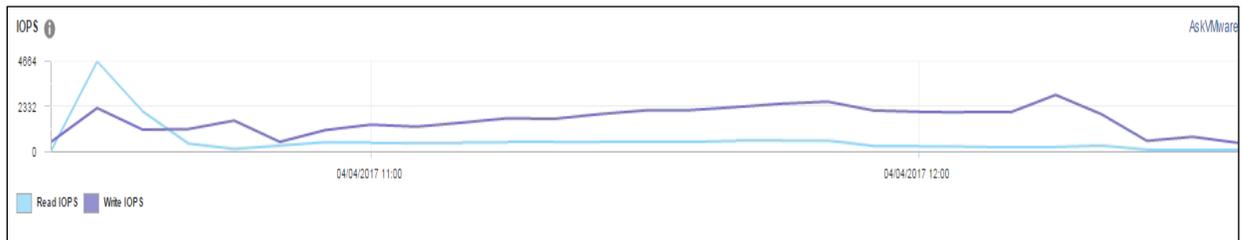


Network bandwidth is not an issue on this test run with a steady state peak of approximately 238 MBps. One of the hosts reached a peak of 621 MBps during the boot storm.

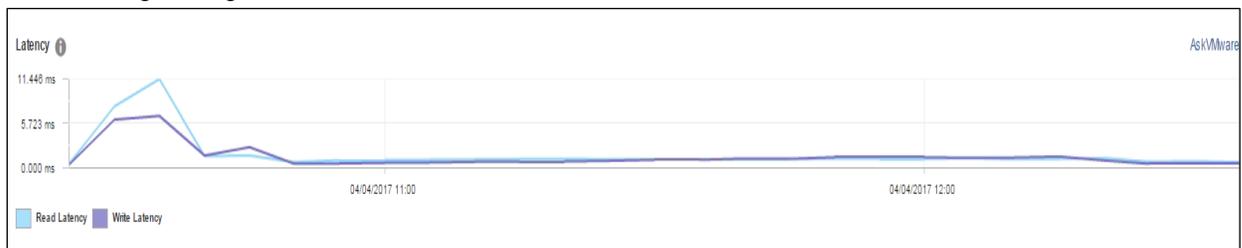


The IOPS graphs and IOPS numbers are taken from the vCenter web console and the graphs clearly show the initial boot storm during a reboot of all the desktops, followed by a settle period, logon, steady state and logoff phases. The graph displays the Disk IOPS figure for the VSAN cluster.

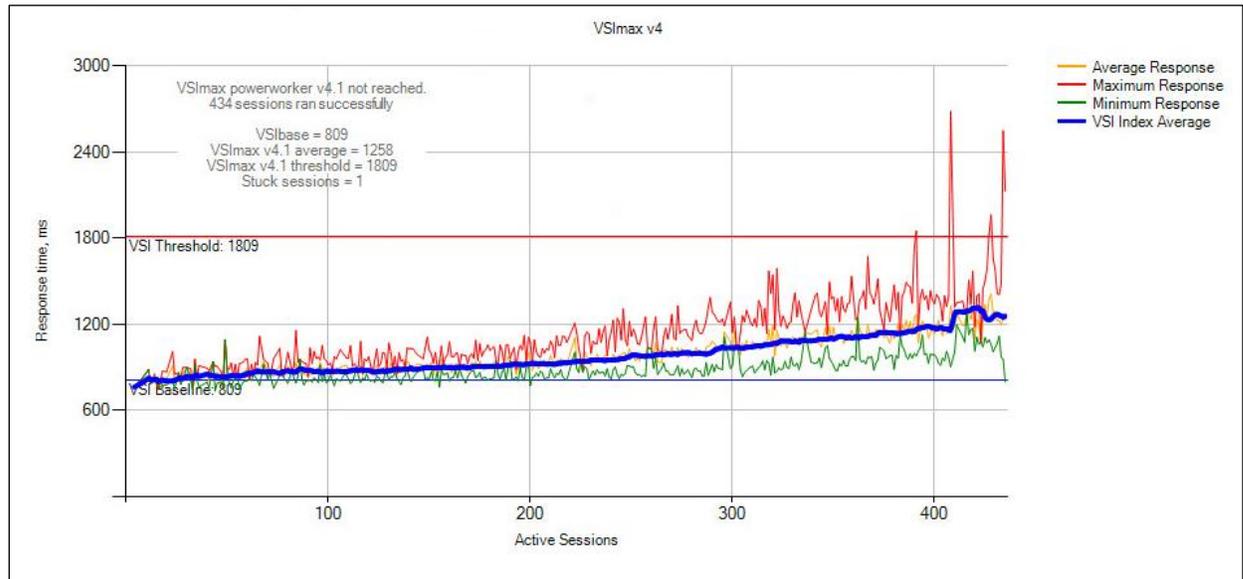
The cluster reached a maximum of 6,912 Disk IOPS during the boot storm before test start and 3,132 IOPS at the start of steady state.



Disk I/O Latency was not an issue. The maximum latency reached was approximately 2.5 ms during the steady state period of test and a brief spike to 17.4 ms during the boot storm. A spike list this, while close to the 20 ms threshold should not be an issue as it occurred during a boot storm and latency returned to normal levels during testing.



The Login VSI Max user experience score shown below for this test was not reached indicating there was little deterioration of user experience during testing and manually interacting with the test sessions confirmed this.



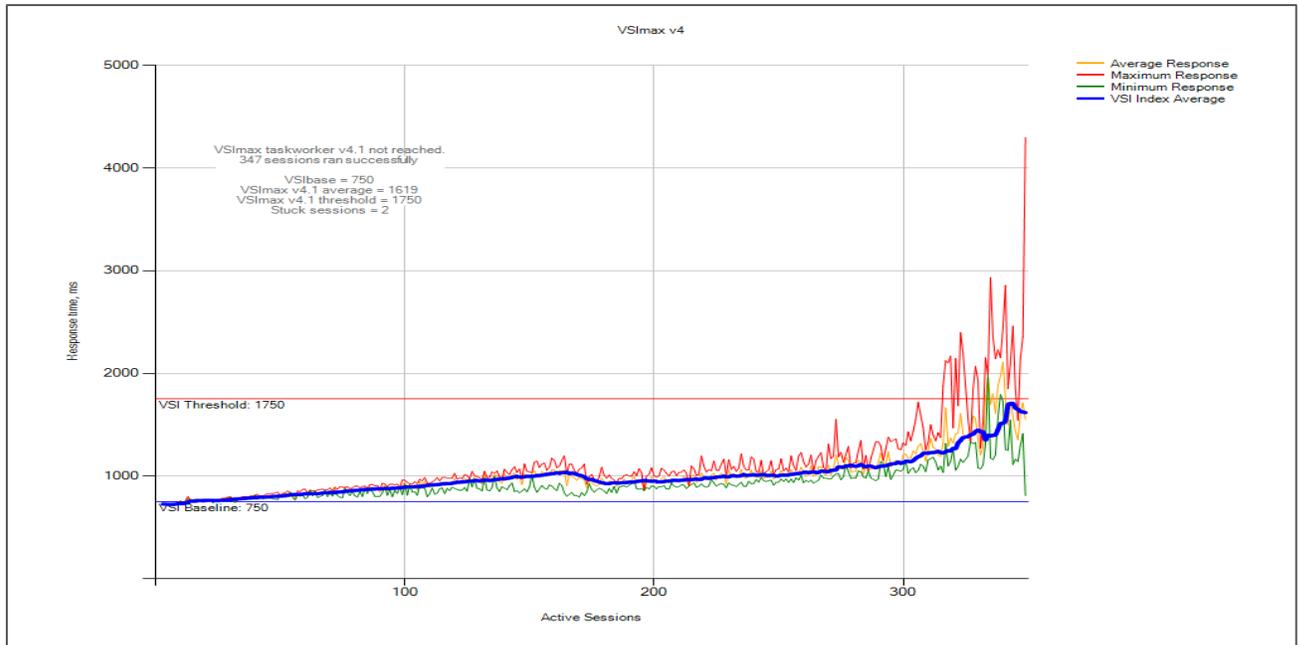
#### Notes:

- As indicated above, the CPU graphs do not take into account the extra 36% of CPU resources available through the E5-2698v4's turbo feature.
- 512 GB of memory installed on each node is just about sufficient for the number of desktops. With memory usage going close to maximum, no extra desktops could have been accommodated for this configuration.
- With extra memory installed there would have been capacity for a few extra desktops as the CPU steady state average did not reach the 95% target threshold.
- Note that during the boot storm the Active Memory metric exceeded the physical installed memory. This is possible because of the way vCenter statistically estimates the number of active memory pages touched by sampling rather than by actually measuring every memory page touched. Measuring every memory page touched would generate too much resources overheads for the hypervisor.

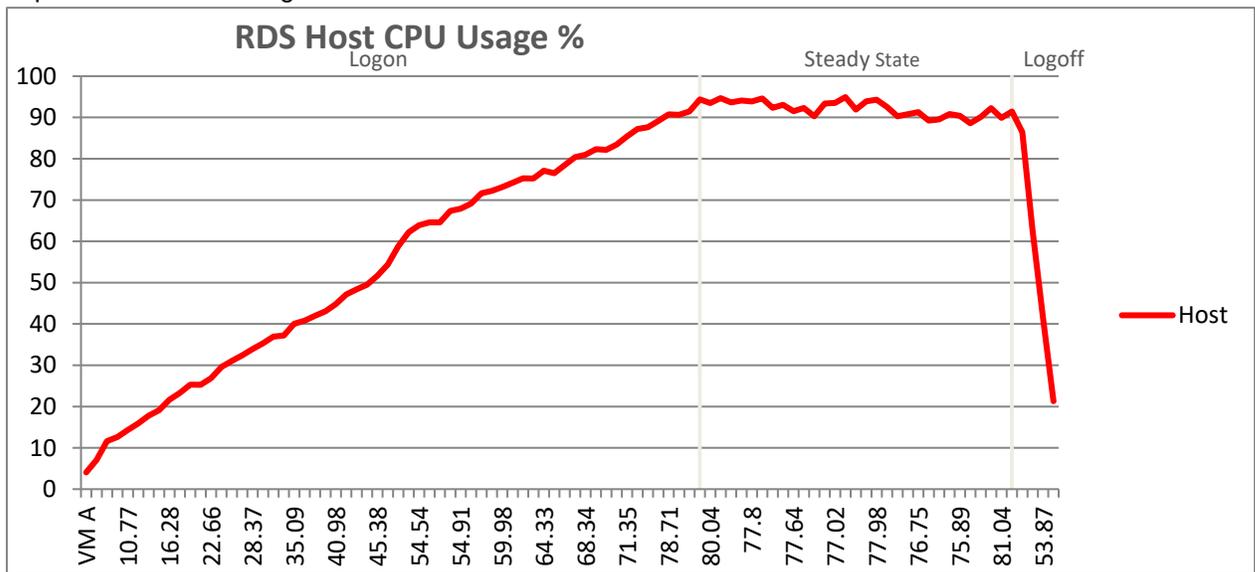
### 6.3.2.4 Task Worker 350 users: ESXi 6.5, Horizon Apps/RDSH

#### 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.

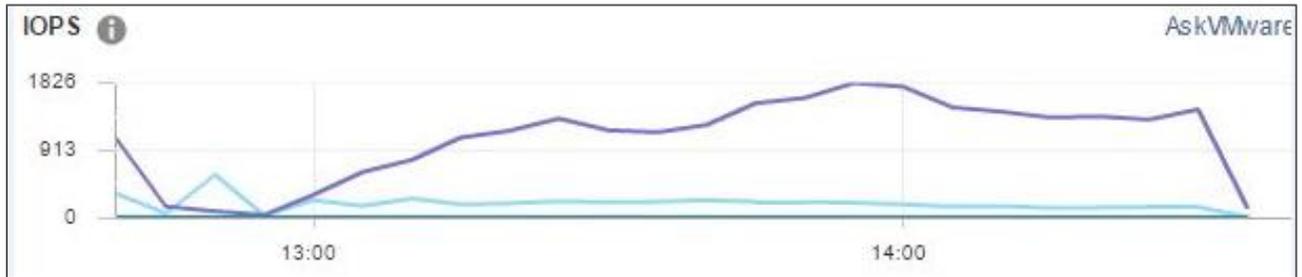


Latency on the datastore did not approach the 20ms threshold. There were no latency spikes during the logon or steady states phases of the test run with 4.5 ms the maximum value reached and 6.4 ms the maximum during the logoff phase.

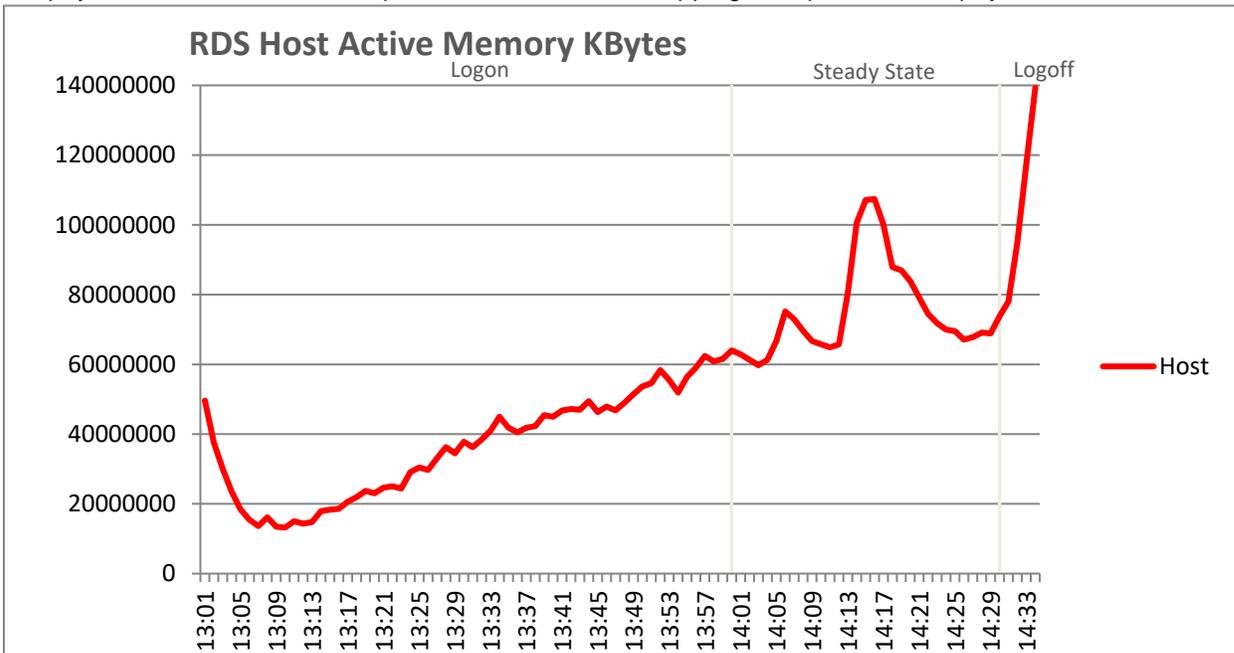
As there was no reboot spike during RDSH testing, IOPS increased steadily as more users logged on reaching a peak just as the steady state phase began. At peak, IOPS reached approximately 2000 resulting in approximately 5.7 IOPS per user.

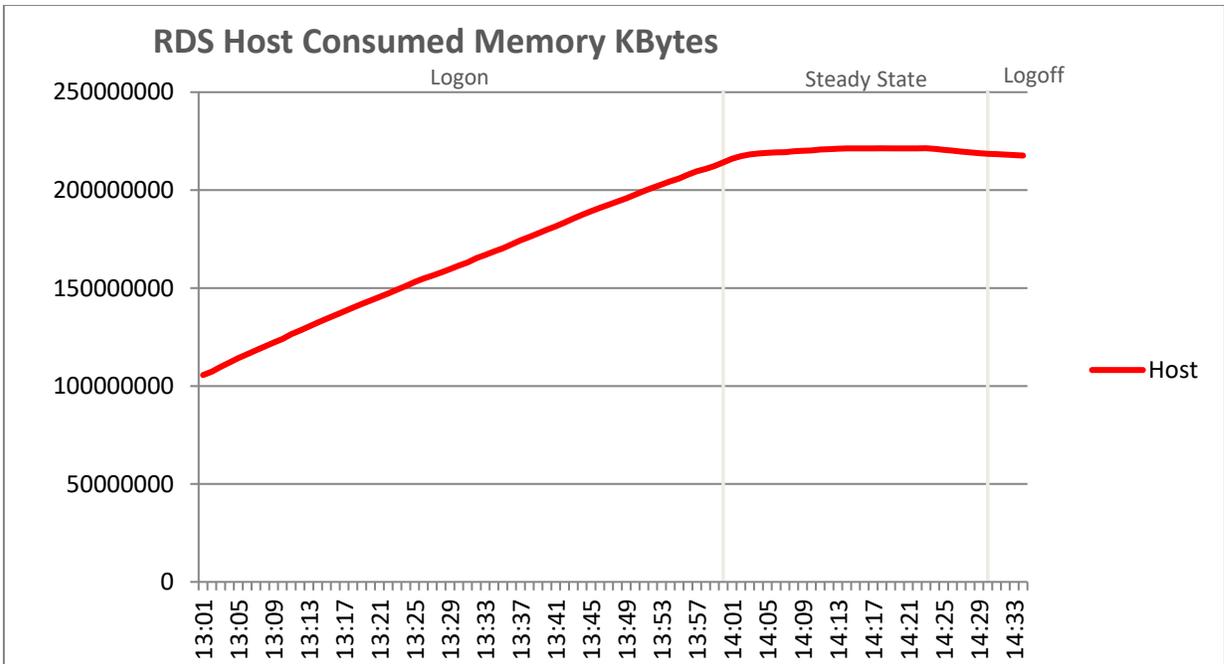
These charts are captured from within vSphere and were a feature released with VSAN6.2 so we do not need to use VSAN Observer as was previously the case with past VSAN validations.

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.

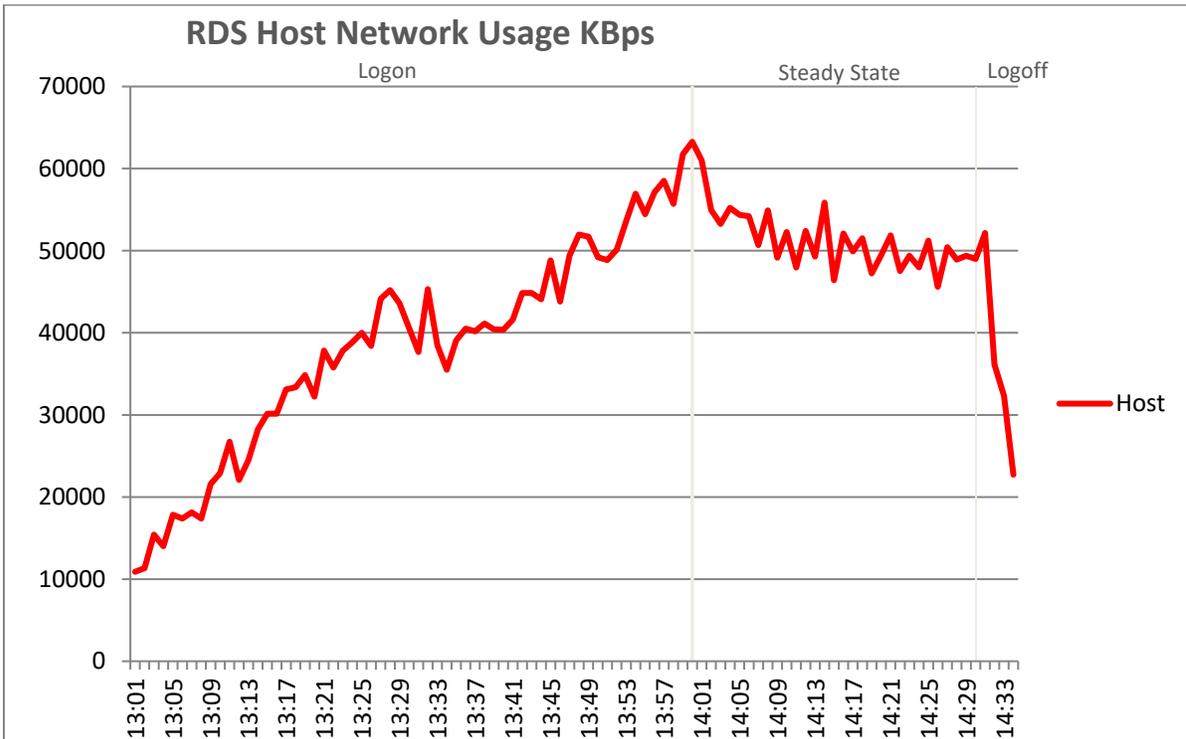


Memory usage is monitored on the ESXi host and the two sampled RDSH VMs. There was no ballooning on the physical host or the two samples RDSH VMs. No swapping took place on the physical host.



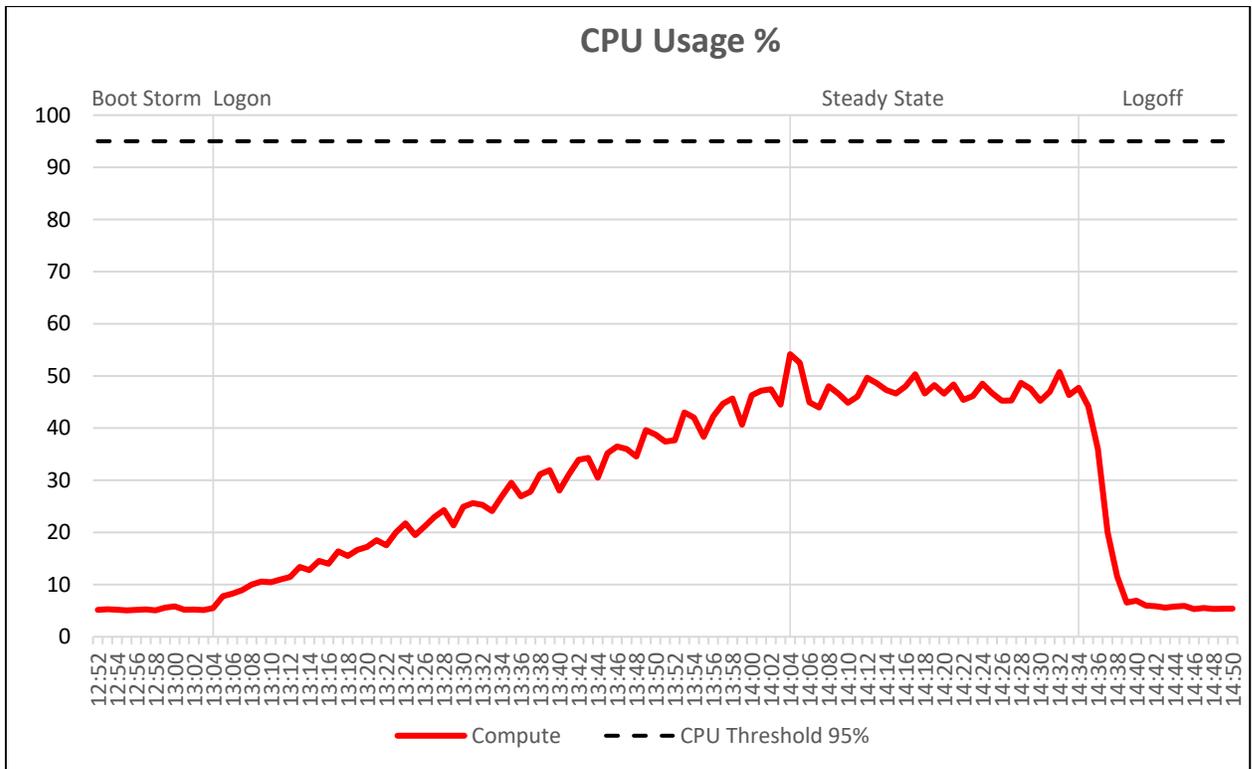


Network utilization was not an issue in this test with host usage reaching a maximum of approximately 65,000 KBps.



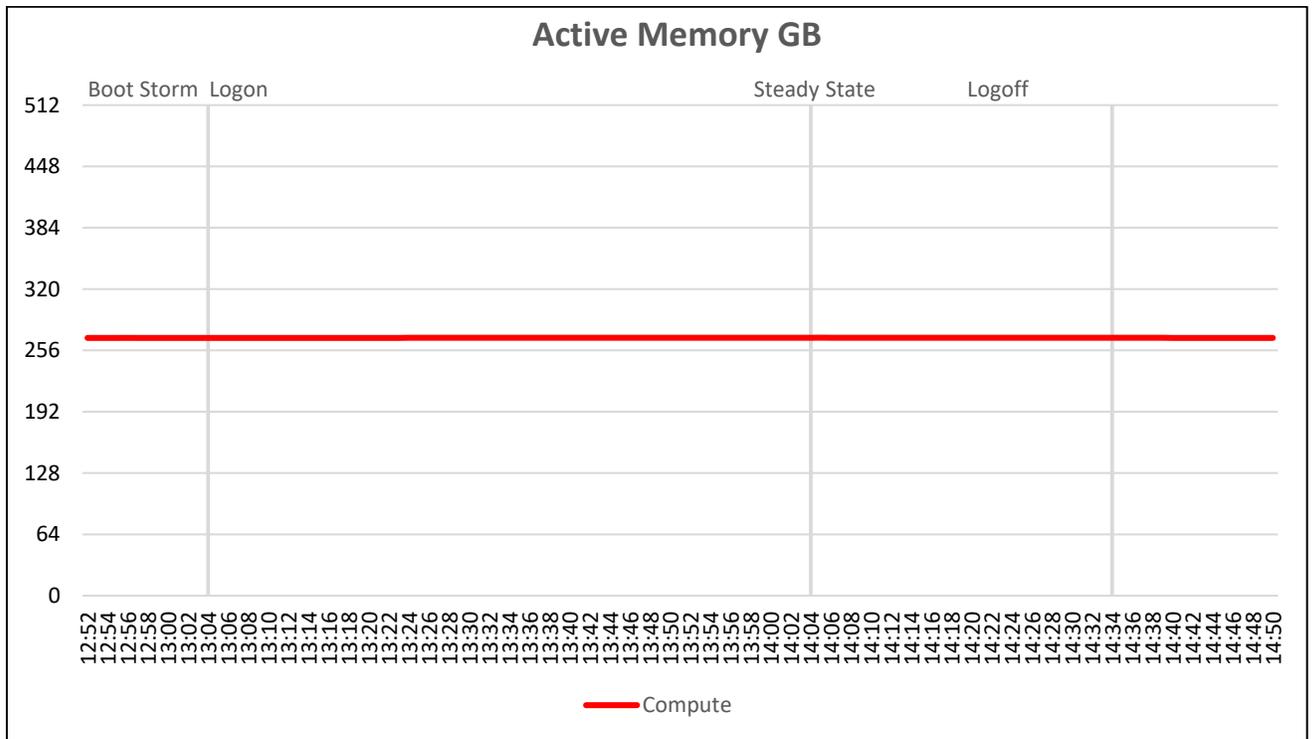
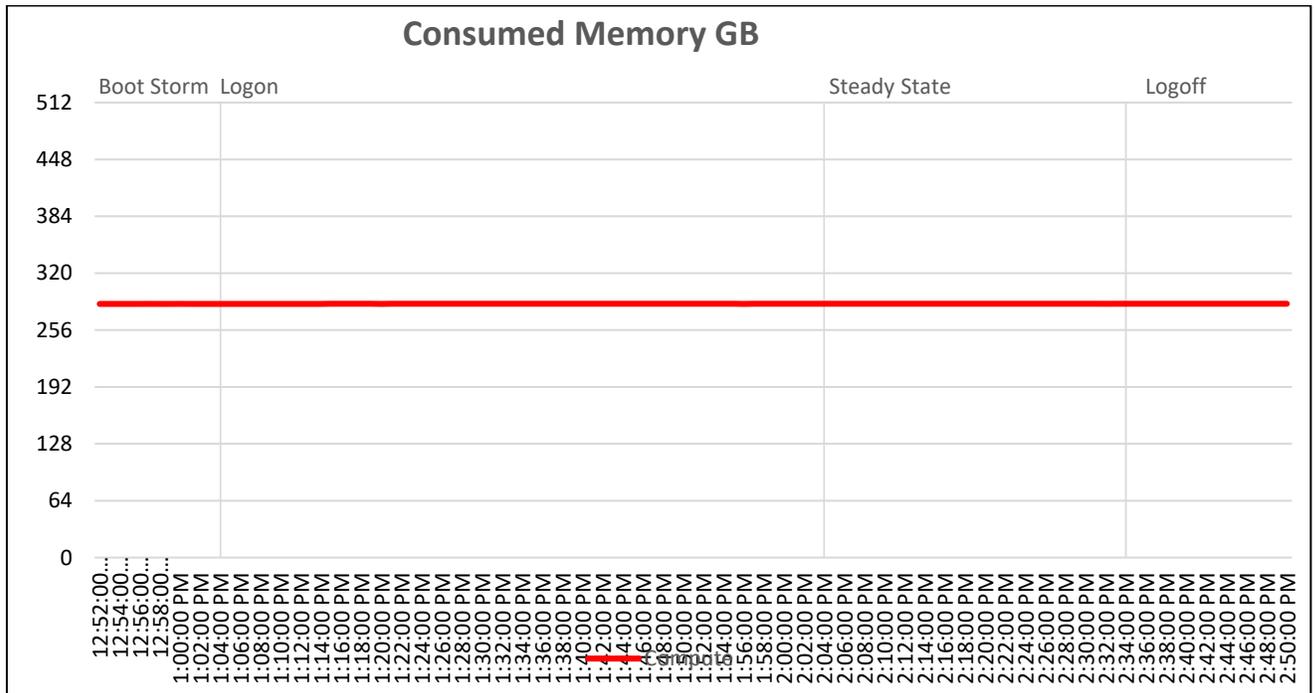
### 6.3.2.5 Power Worker (M10-1B), 64 vGPU users per host, ESXi 6.5, Horizon 7 linked clones

The compute host was populated with two NVIDIA M10 cards. The LVSI Power Worker workload was run for this test.

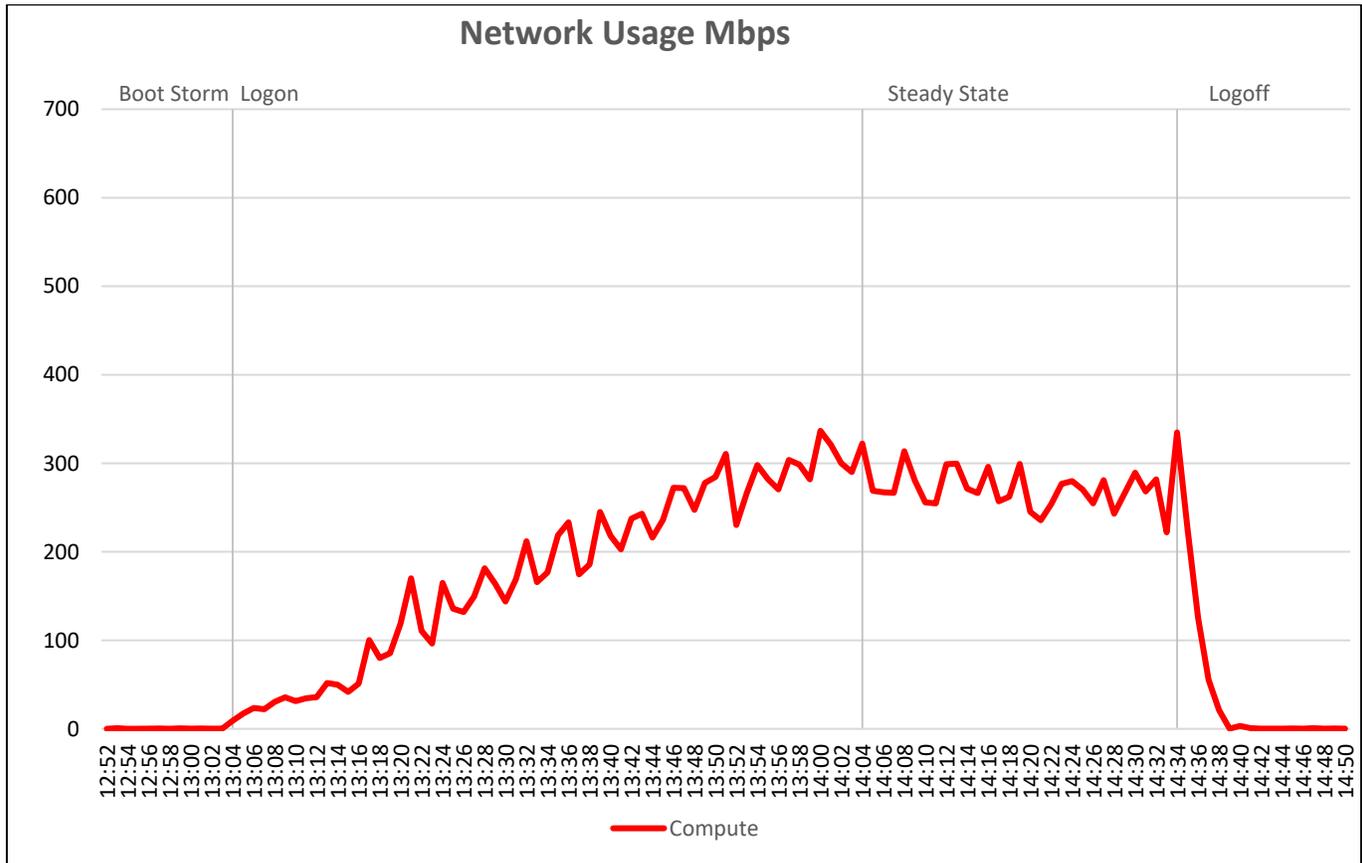


CPU usage was well below the 95% mark using 64 clients with M10-1b profile.

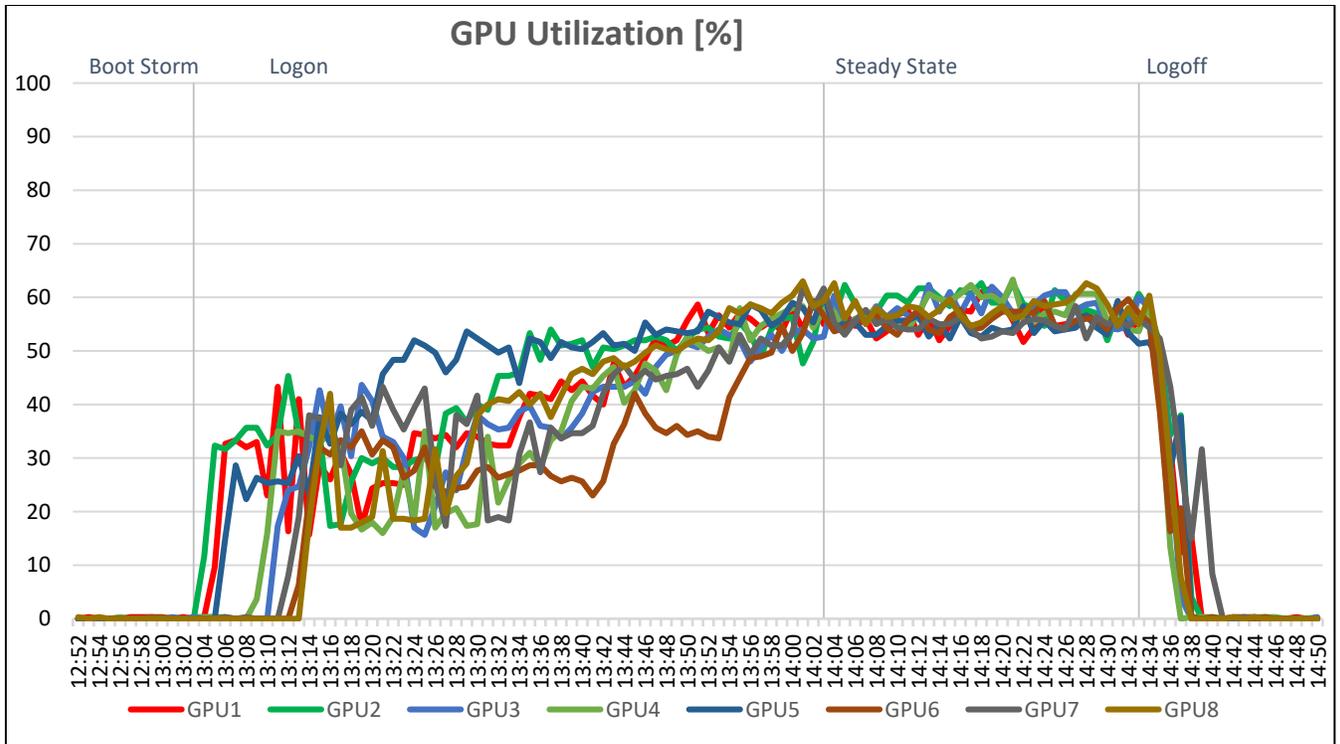
Since memory usage is reserved with vGPU applications, consumed and active memory will show a value slightly above the reserved value of: 64 clients \* 4GB per client. Installed memory for this test was 512 GB. During tests no memory ballooning or swapping occurred.



Network bandwidth usage showed a consistent usage during steady state of around 280 Mbits/sec. Network usage was not considered to be an issue during tests.



Average GPU processor usage was in the range of 60% during steady state. This utilization can be expected for the Power Worker load, since it is considered not to be a high load in GPU terms.



Compute – GPU Utilization

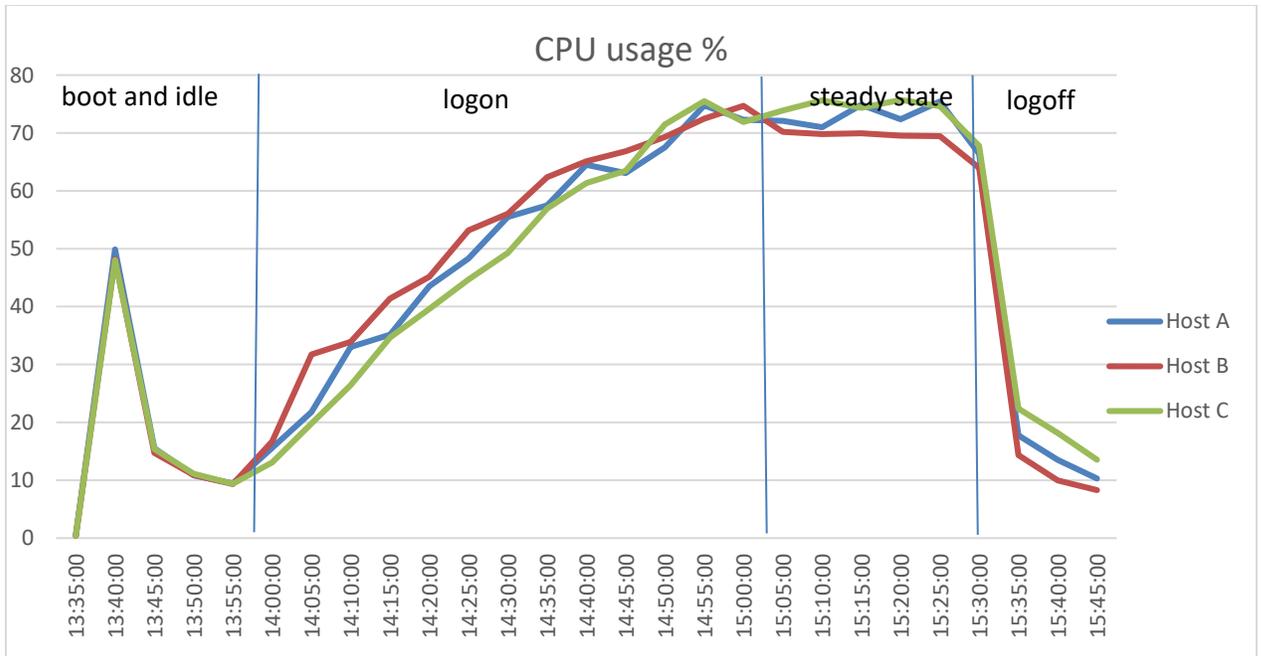
### 6.3.3 vSRN-FC430 B5

Refer to Section 3.2 for hardware configuration details. This would be classed as a vSRN HY-6 (Hybrid) or AF-4 (All-Flash) configuration on the vSRN HCL

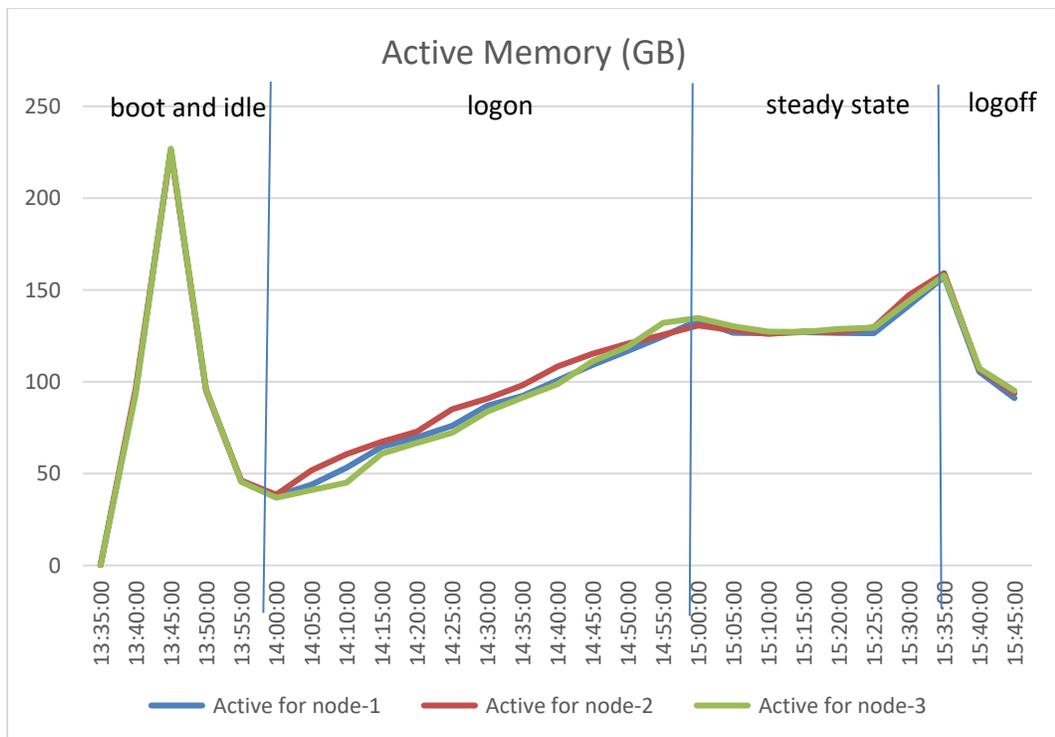
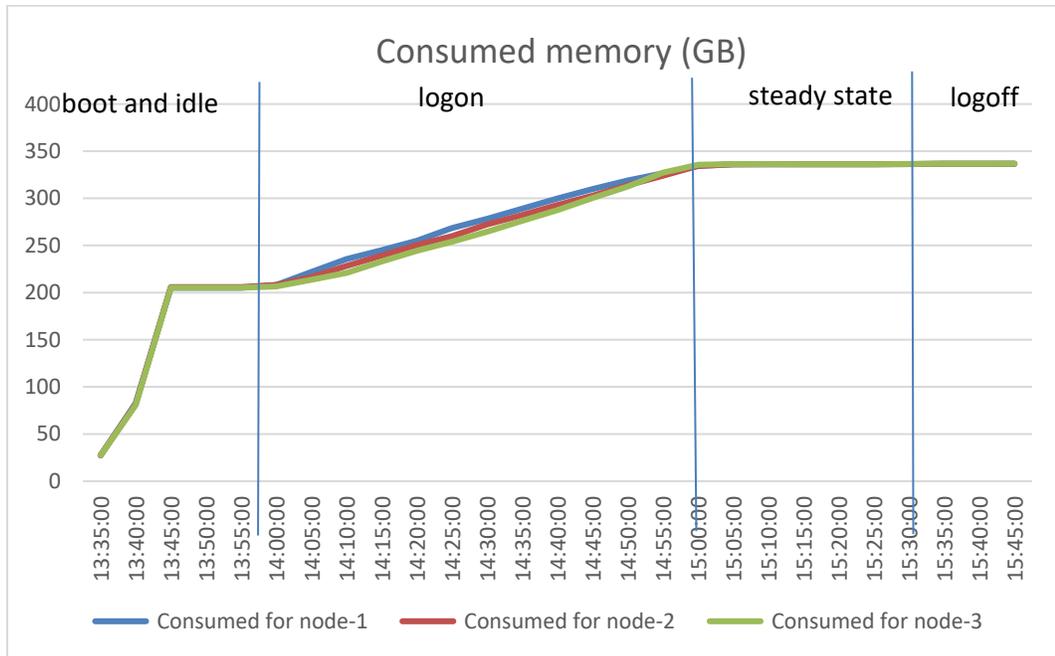
#### 6.3.3.1 Task Worker, 450 users, ESXi 6.5, Horizon 7 Linked clones

450 VM of Task worker workload configuration as described in section 2.3) deployed across the 3 FC430 compute nodes in the FX2 chassis using Horizon linked clone method. Each node hosts 150 VM. When all the VM powered on before test commence, the CPU usage was approximately 10%

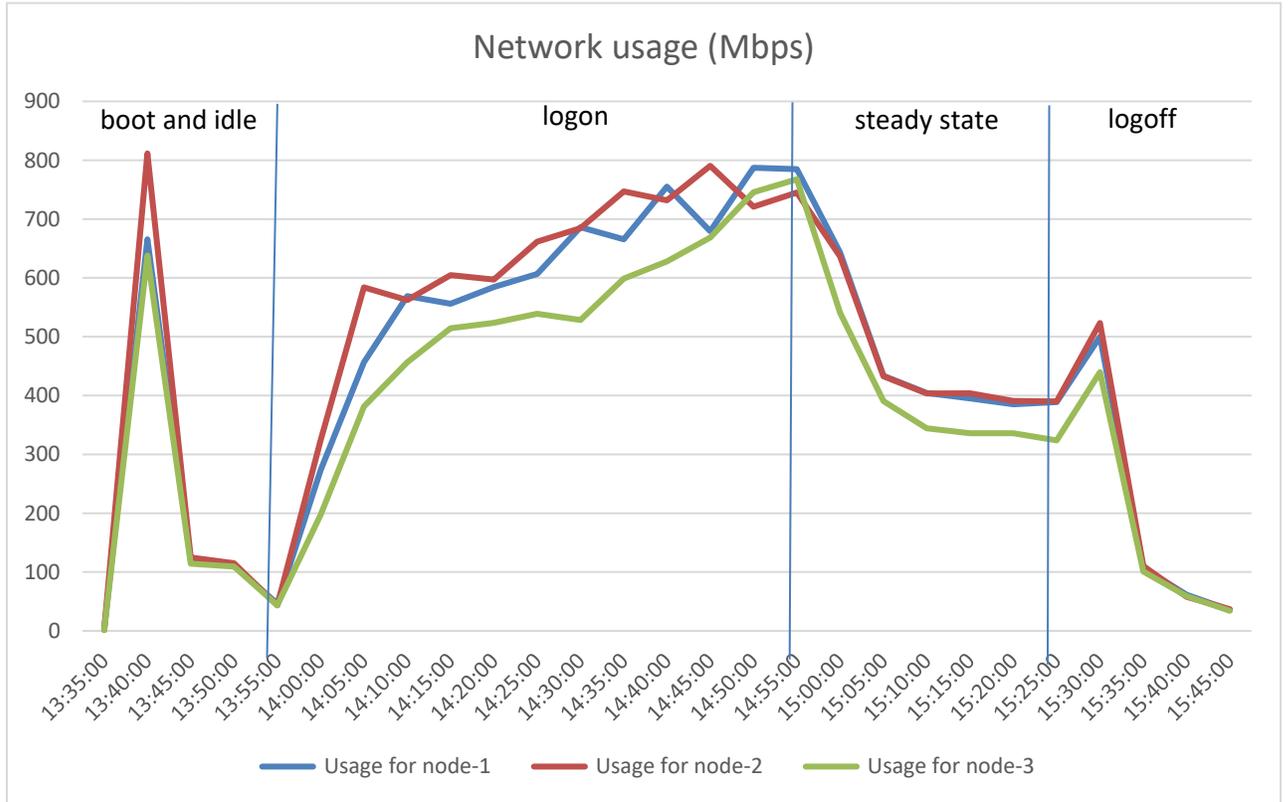
Below graph shows the CPU usage during the test session. The max CPU usage stays around 75% during steady state when all users logged on and running LoginVSI work load activities.



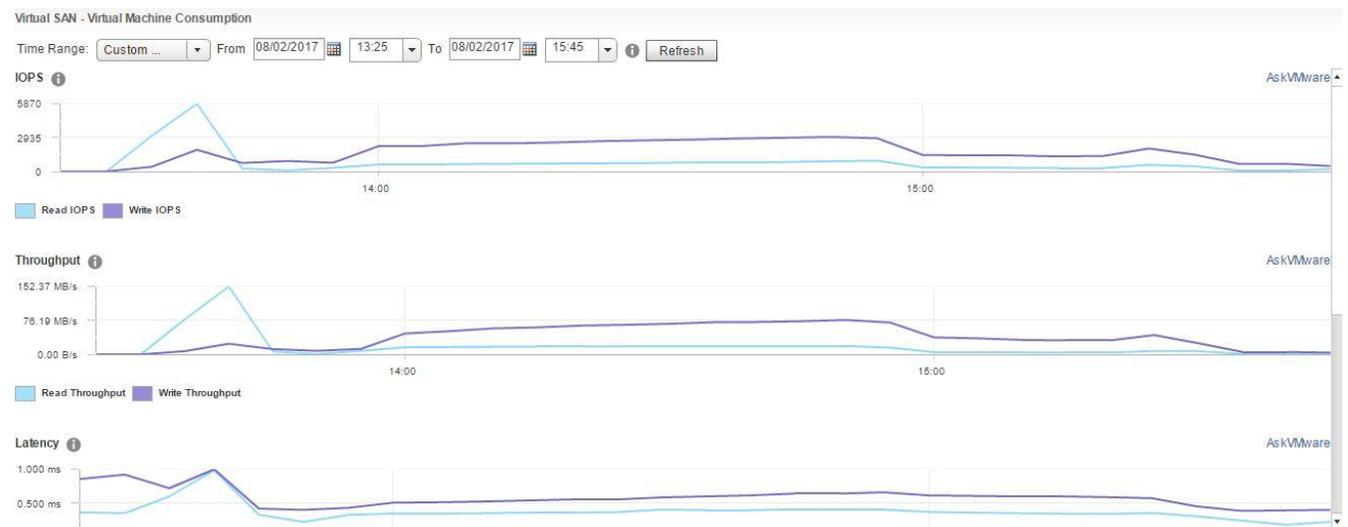
Below two graphs show the memory usage during the task worker workload test. Consumed memory maximized at around 335GB and active memory is around 130GB during steady state. There was no memory ballooning or swapping any time during the tests. Indicates that the 512GB physical on each of the node has enough capacity to support the 150 task worker users.



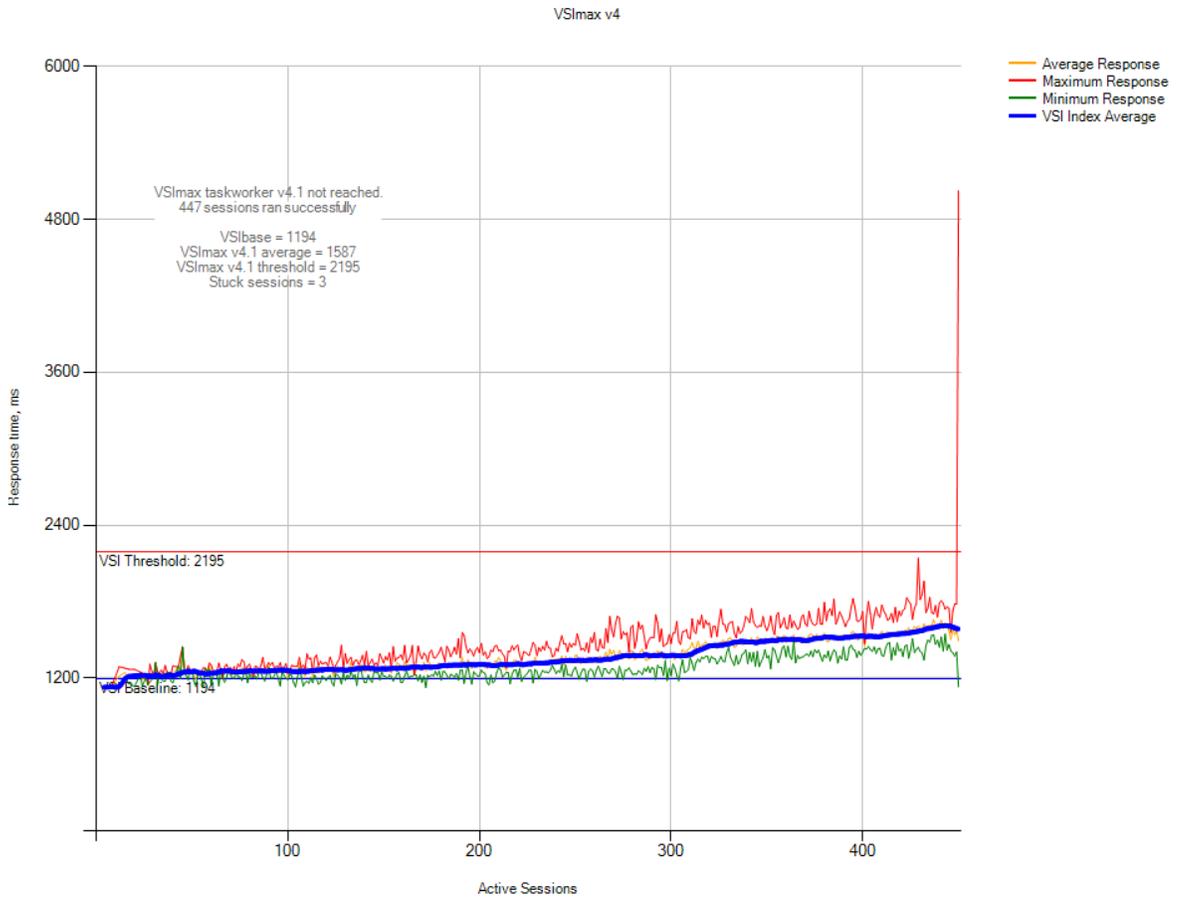
Below graph shows the network usage during 450 task worker workload test, the network usage peaked at around 800 Mbps, this is over two 10Gbps NIC on each node, so it is not near any saturation point.



Below screenshot shows the VSAN data store IOPS, throughput and latency during the test. The latency is under 1ms throughout the test.



The Login VSI Max user experience score shown below for this test was not reached indicating there was little deterioration of user experience during testing and manually interacting with the test sessions backed this up, mouse and window responses were fast and video play back was of good quality.



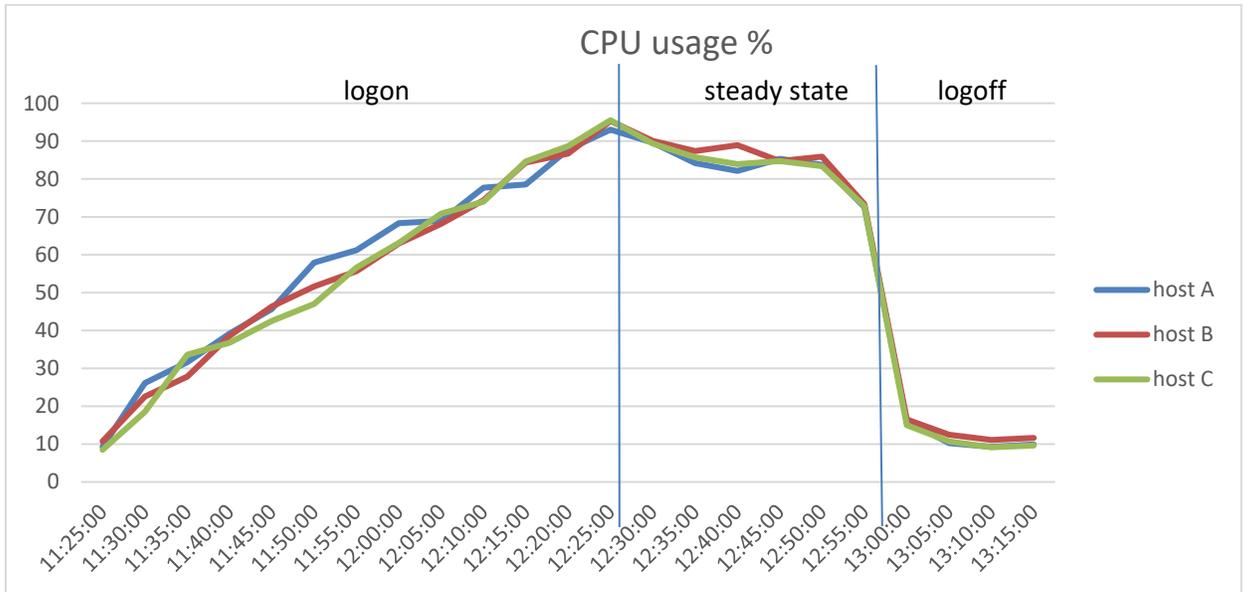
**Note:**

There were three sessions that failed to complete the test cycle. Upon investigation, it was concluded that the failure was not related to host performance reasons. Since it was only 3 out of 450 sessions, it was below the 2% acceptable results threshold. The LoginVSI test “VSI max not reached” result is considered valid.

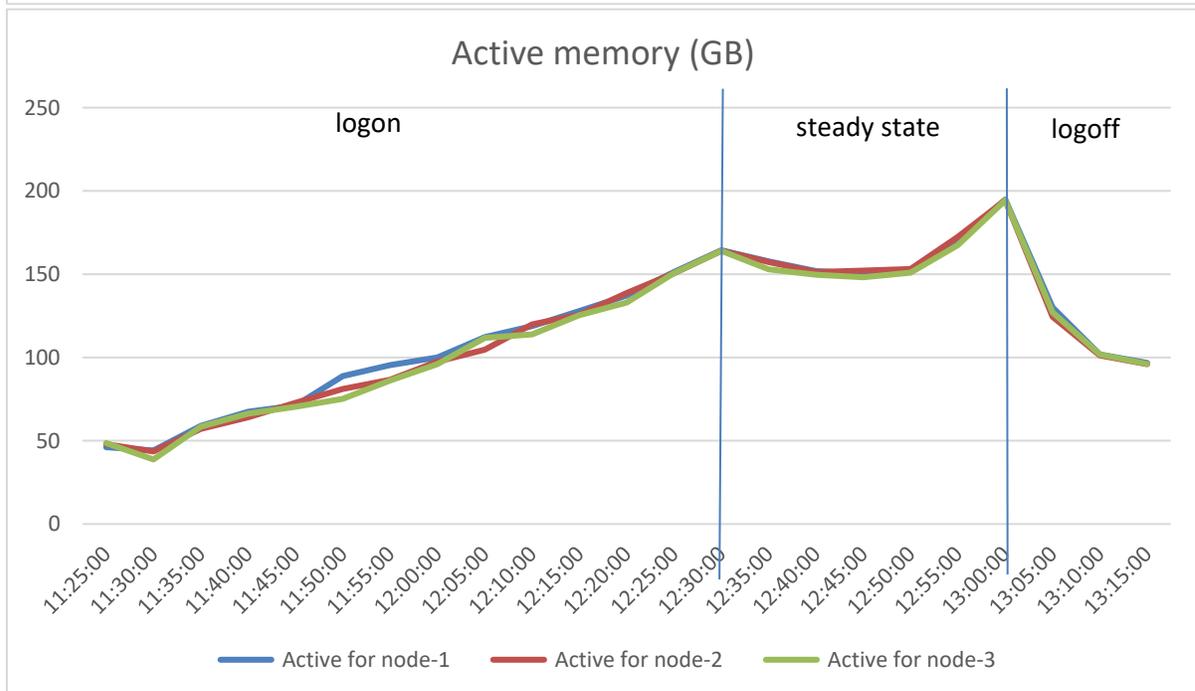
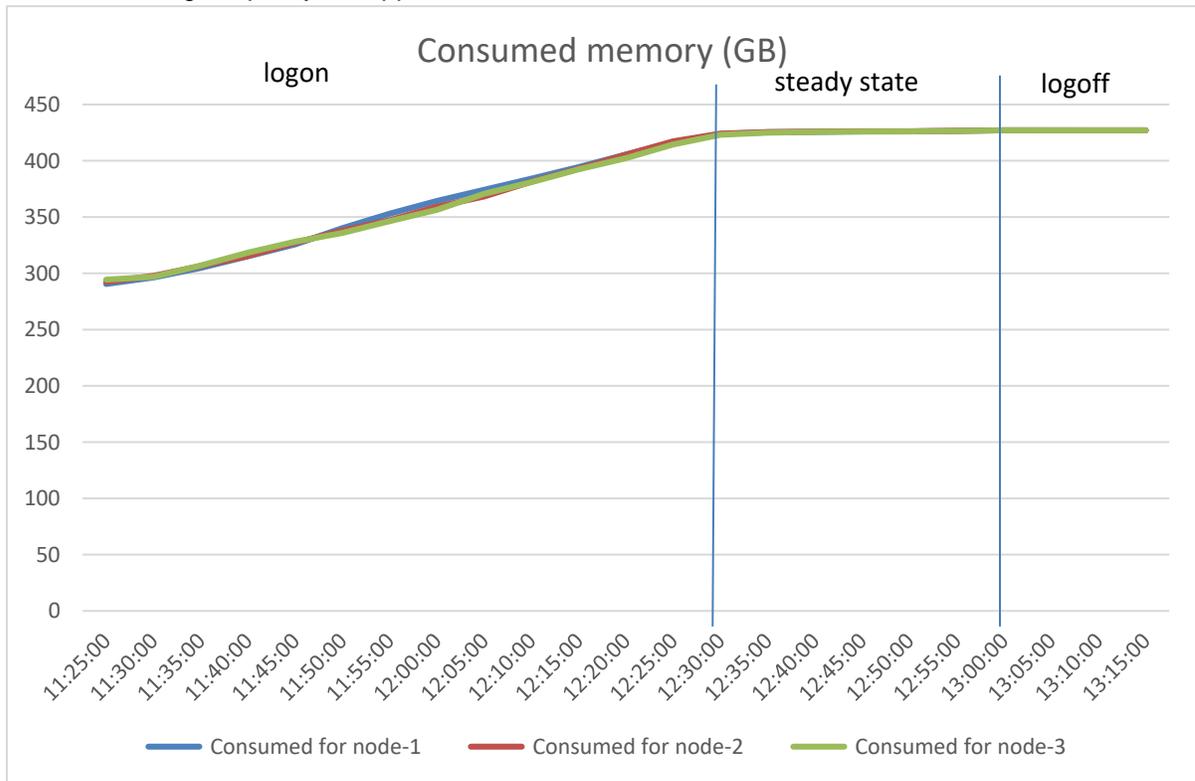
### 6.3.3.2 Knowledge Worker, 390 users, ESXi 6.5, Horizon 7 linked clones

390 VM (of knowledge worker workload configuration as described in section 2.3) deployed across the 3 FC430 compute nodes in the FX2 chassis using Horizon linked clone method. Each node hosts 130 VM. When all the VM powered on in the idle state before test commence, the CPU usage was approximately 9%

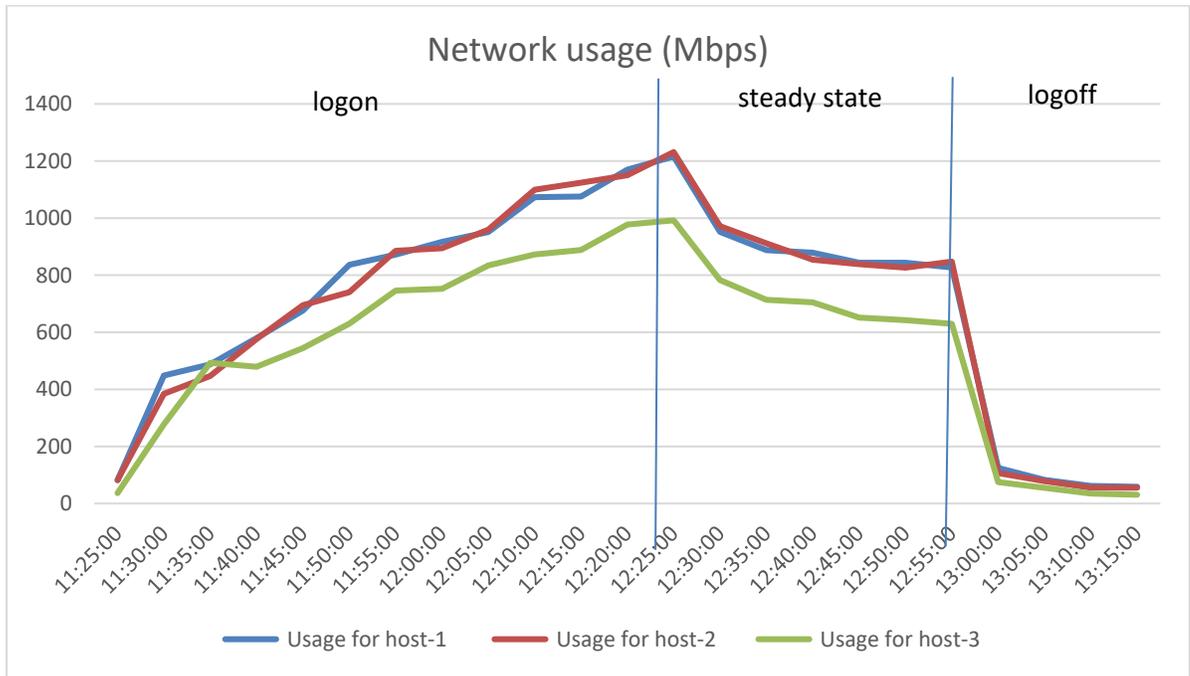
Below graph shows the CPU usage during the test. The max CPU usage was just over 90% when all users logged on, and drop to around 85% during steady state after all users logged on and running LoginVSI knowledge work load activities.



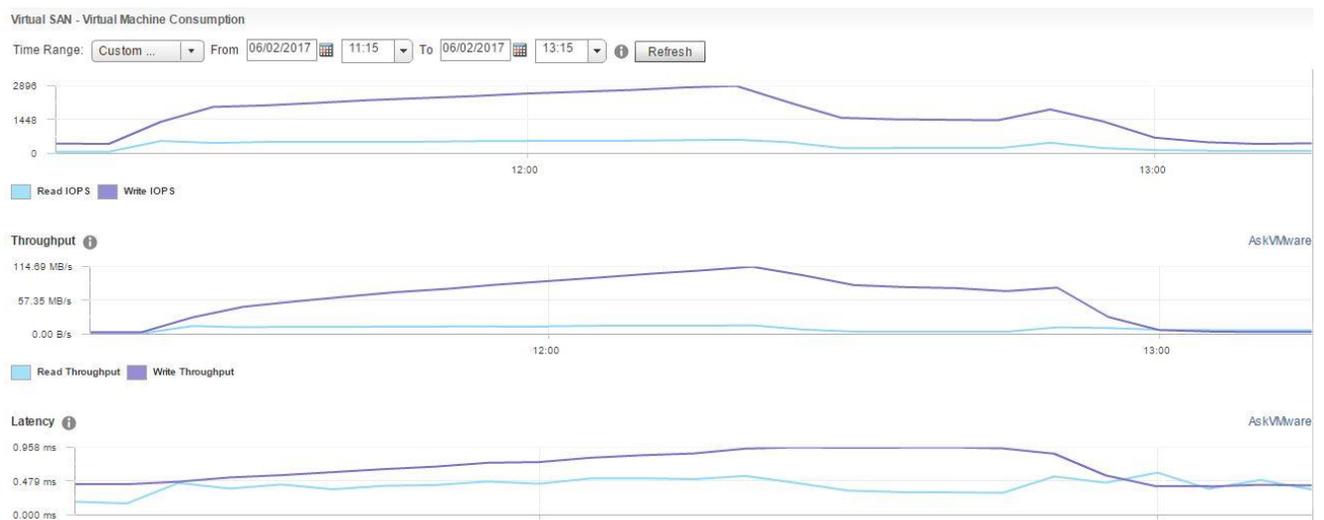
Below two graphs show the memory usage during the task worker workload test. Consumed memory maximized at around 425GB and active memory is just over 150GB during steady state. There was no memory ballooning or swapping any time during the tests. Indicates that the 512GB physical on each of the node has enough capacity to support the 130 task worker users.



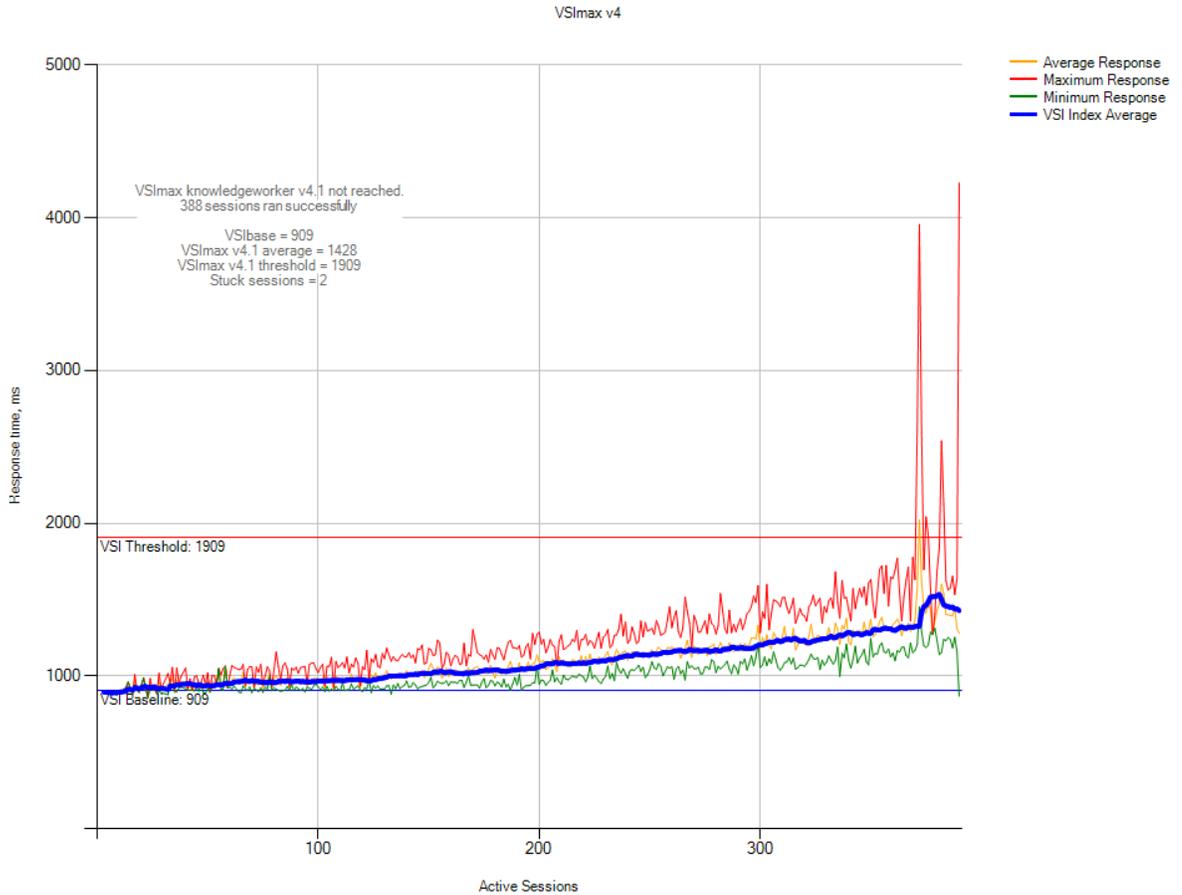
Below graph shows the network usage during 390 knowledge worker workload test, the network usage peaked at around 1,200 Mbps on each node, this is over two 10Gbps NIC on each node, so it is not anyway near the saturation point.



Below screenshot shows the VSAN data store IOPS, throughput and latency during the test. The latency is under 1ms throughout the test.



The Login VSI Max user experience score shown below for this test was not reached indicating there was little deterioration of user experience during testing and manually interacting with the test sessions backed this up, mouse and window responses were fast and video play back was of good quality.



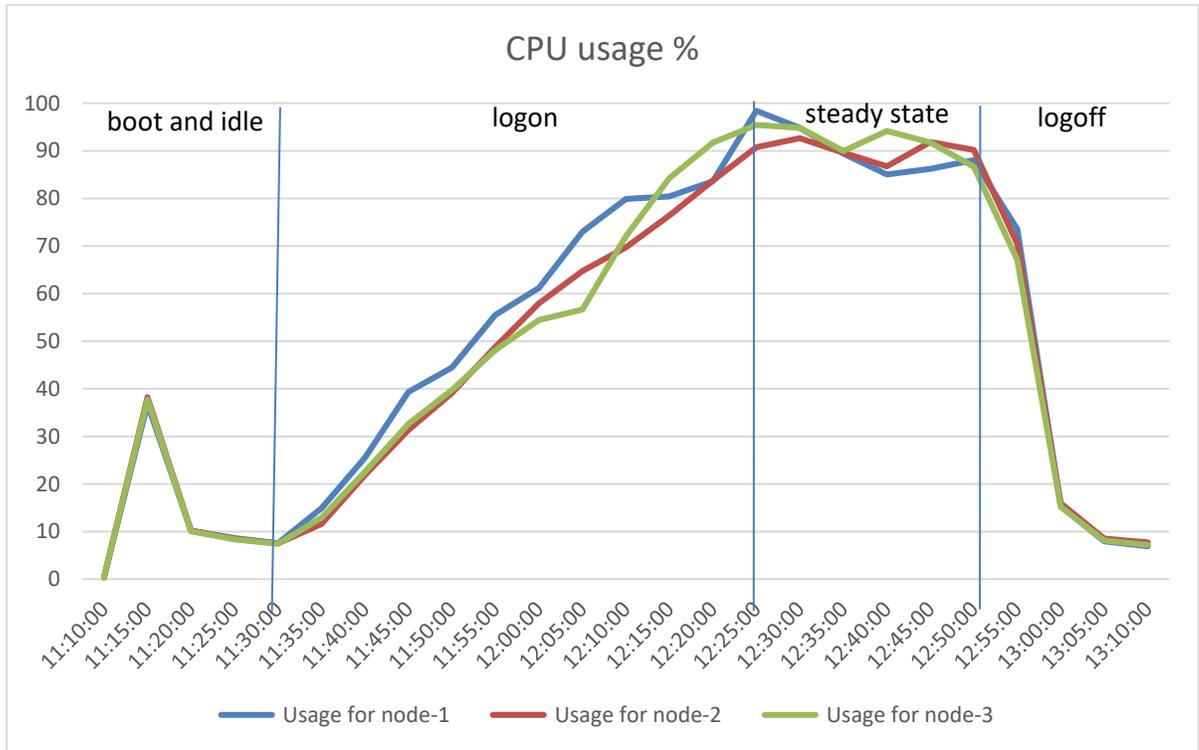
**Note:**

There were two sessions that failed to complete the test cycle. Upon investigation, it was concluded that the failure was not related to host performance reasons. Since there were only 2 out of 390 sessions, it was below the 2% acceptable results threshold. The LoginVSI test "VSImax not reached" result is considered valid.

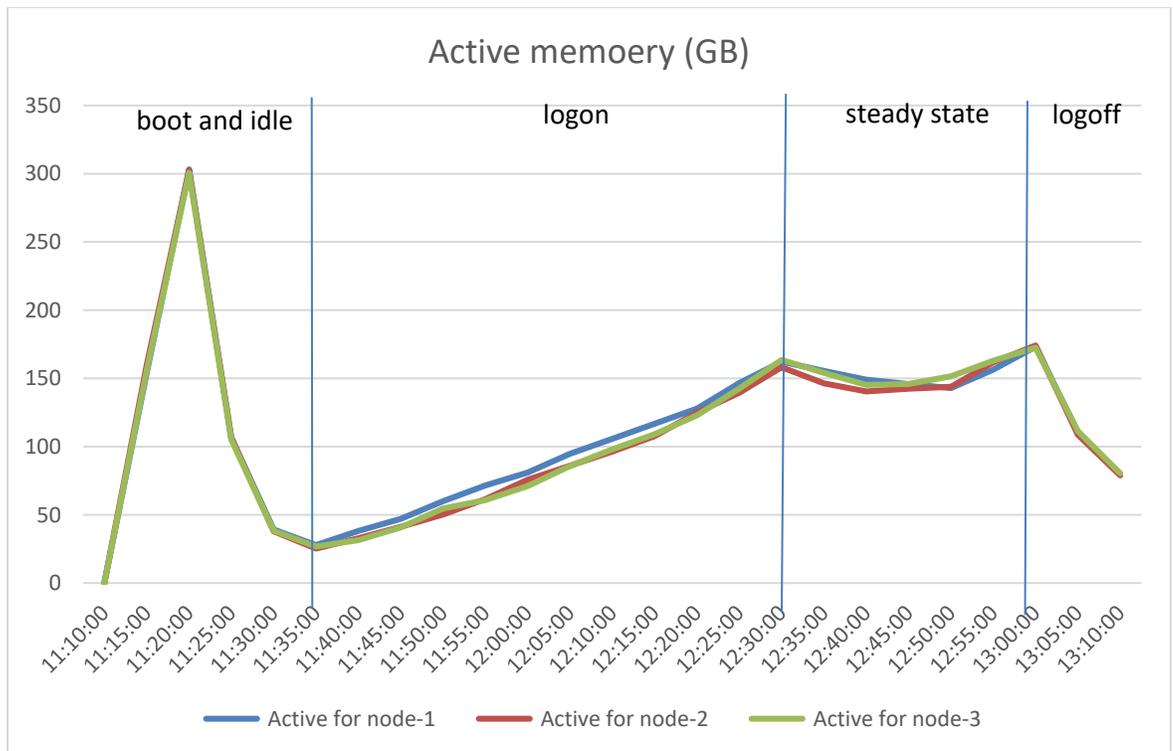
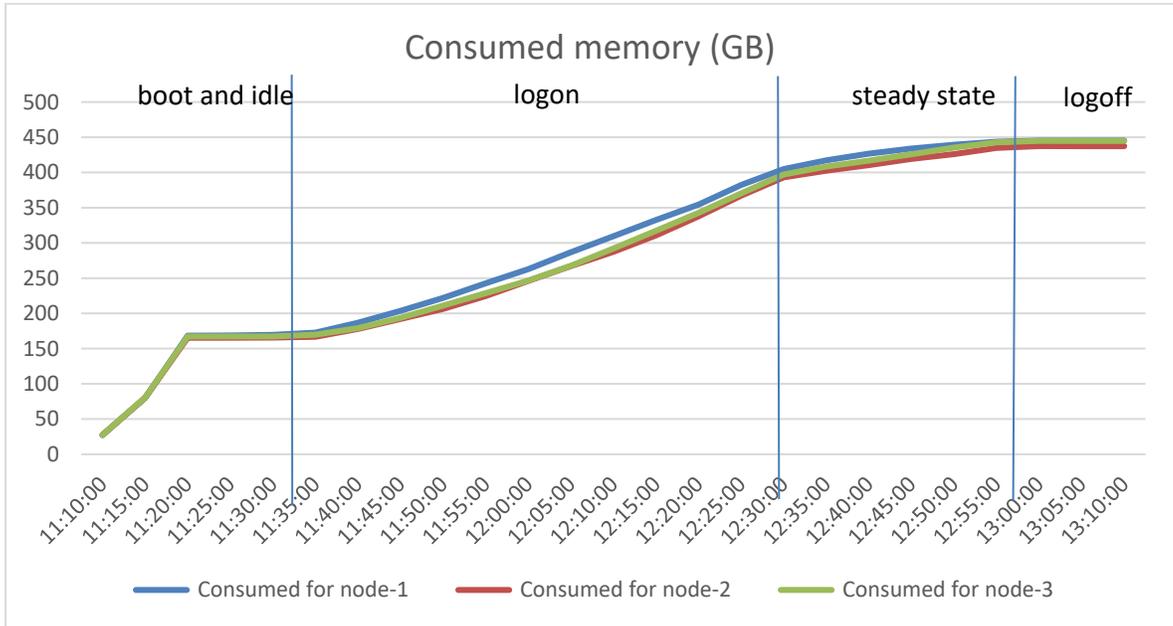
### 6.3.3.3 Power Worker, 315 users ESXi 6.5, Horizon 7 linked clones

315 Power worker workload configuration VMs deployed across the 3 FC430 compute nodes in the FX2 chassis using Horizon linked clone method. Each node hosts 105 VM. When all the VM powered on before test commence, the CPU usage was approximately 8%.

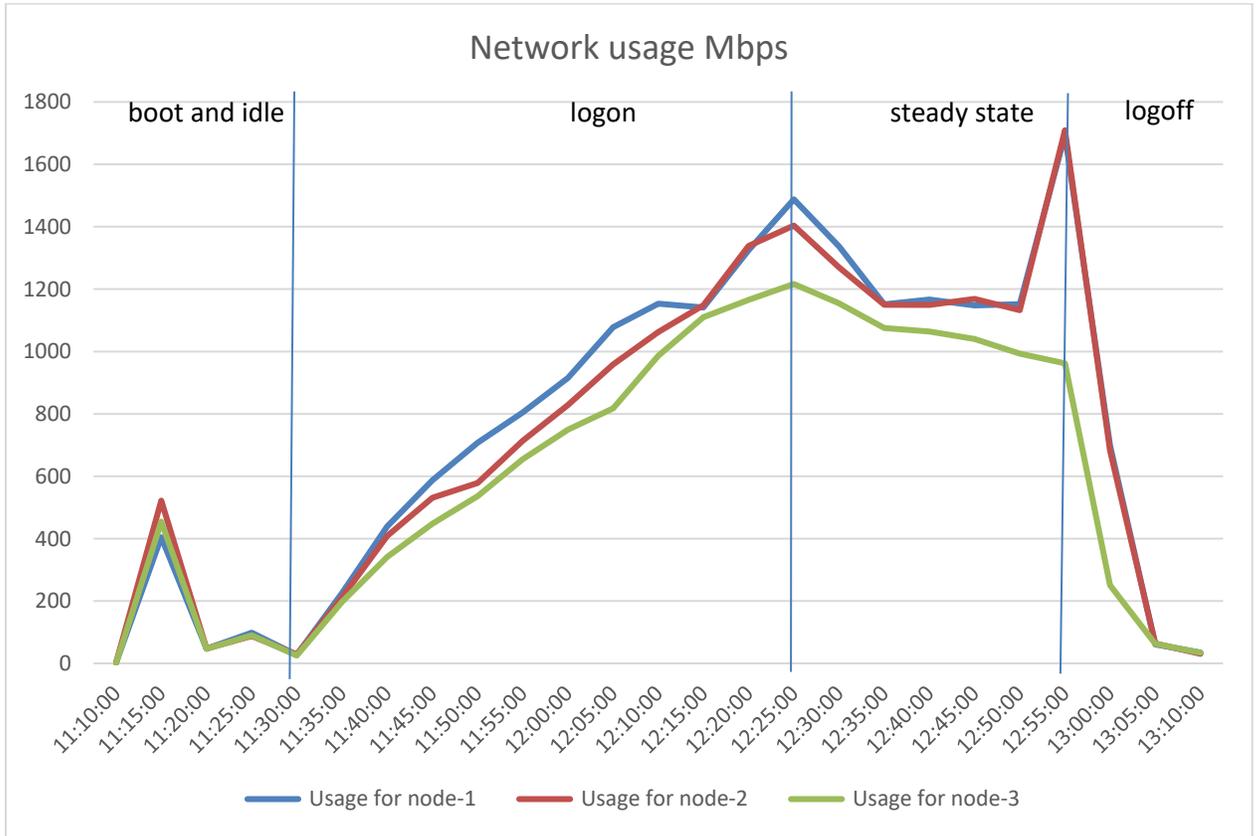
Below graph shows the CPU usage during the test. The max CPU usage was around 95% when all users logged on, and drop to around 90% during steady state after all users logged on and running LoginVSI Power worker workload activities.



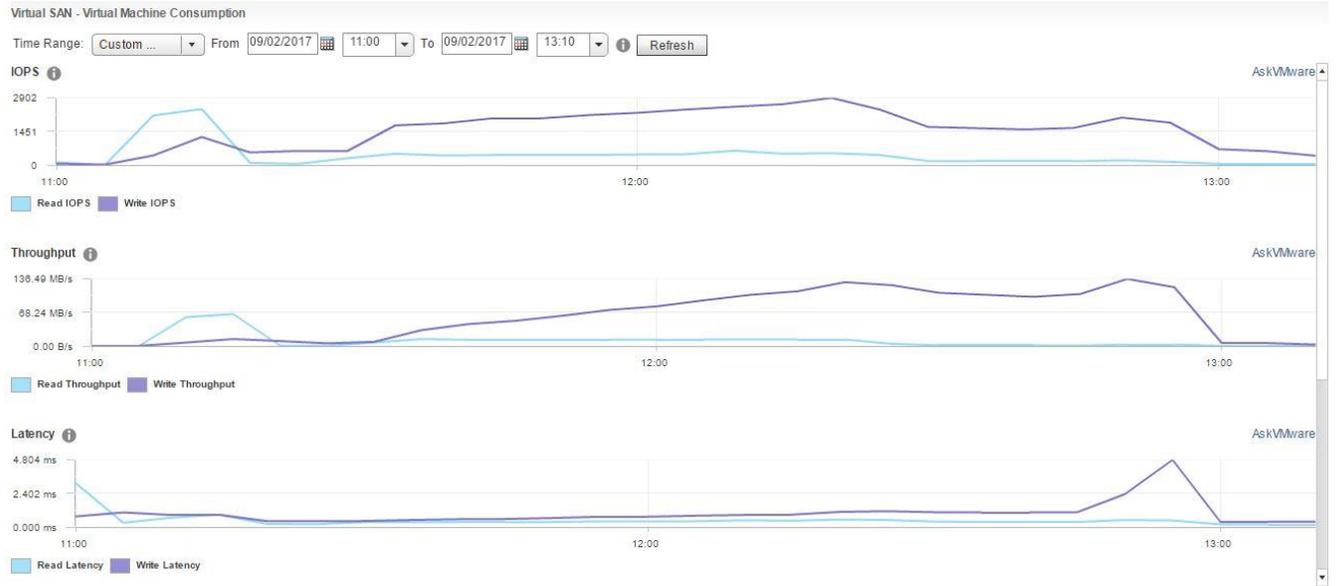
Below two graphs show the memory usage during the task worker workload test. Consumed memory maximized at nearly 450GB and active memory is just over 150GB during steady state. There was no memory ballooning or swapping any time during the tests. Indicates that the 512GB physical on each of the node has enough capacity to support the 130 task worker users.



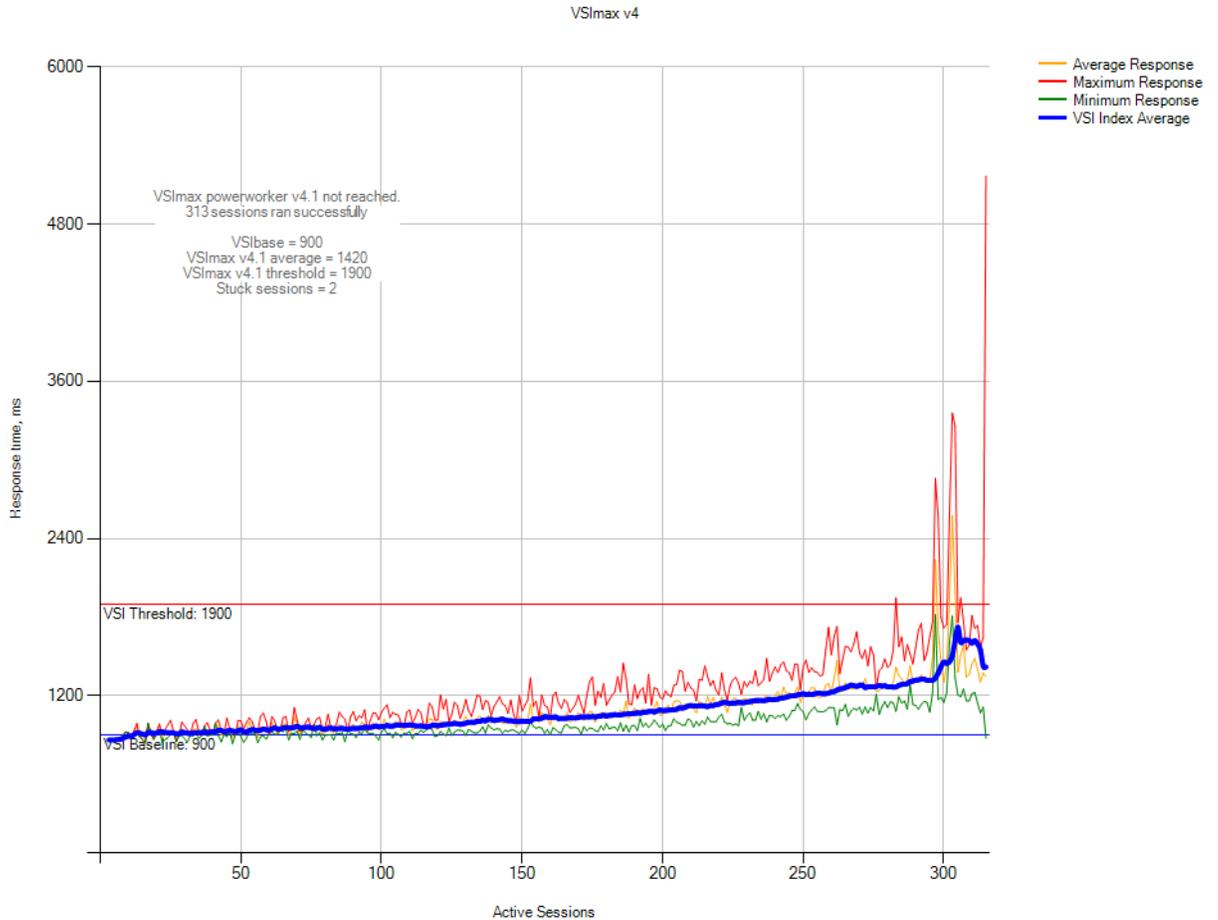
Below graph shows the network usage during 315 Power worker workload test, the network usage peaked at around 1,500 Mbps on each node, this is over two 10Gbps NIC on each node, so it is not anyway near the saturation point.



Below screenshot shows the VSAN data store IOPS, throughput and latency during the test. The latency is under 1ms throughout the test, there was one latency spike when user log off, possibly caused by user profile write back to file share, causing high wrote queue. But it is only 4ms, still well under 20ms latency threshold.



The Login VSI Max user experience score shown below for this test was not reached indicating there was little deterioration of user experience during testing and manually interacting with the test sessions backed this up, mouse and window responses were fast and video play back was of good quality.



**Note:**

There were two sessions that failed to complete the test cycle. Upon investigation, it was concluded that the failure was not related to host performance reasons. Since it was only 2 out of 390 sessions, it is below the 2% acceptable results threshold. The LoginVSI test “VSI max not reached” result is considered valid.

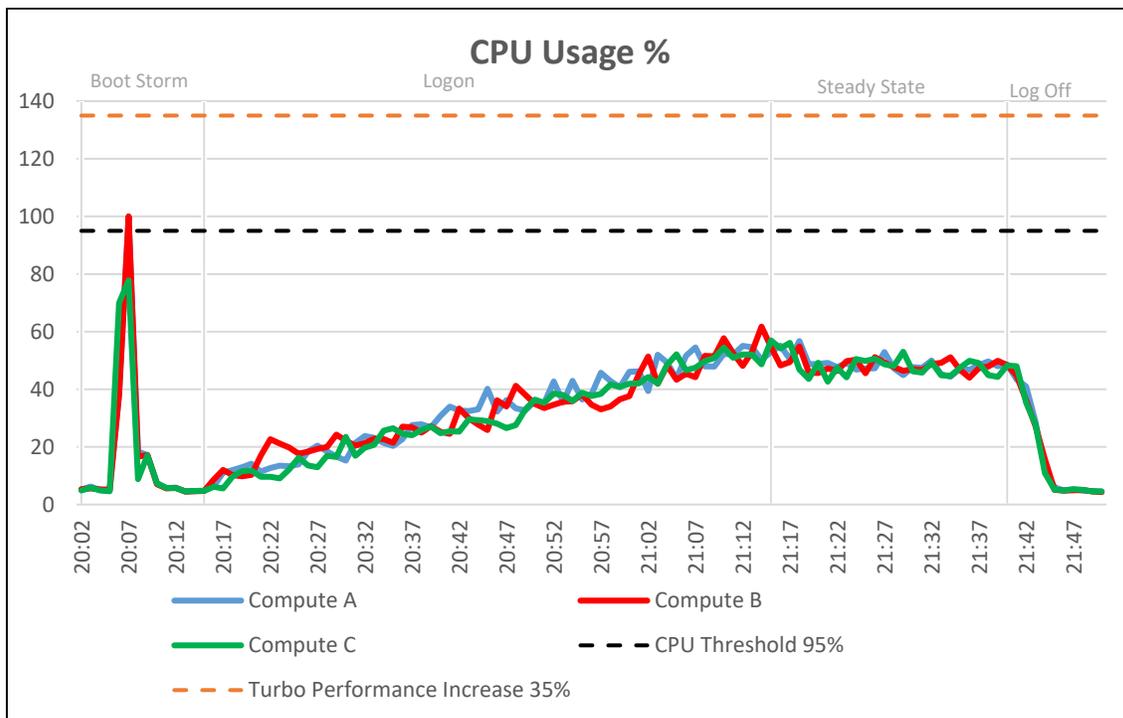
### 6.3.4 vSAN-FC630 B5

This would be classed as a vSRN AF-6 (All-Flash) configuration on the vSRN HCL

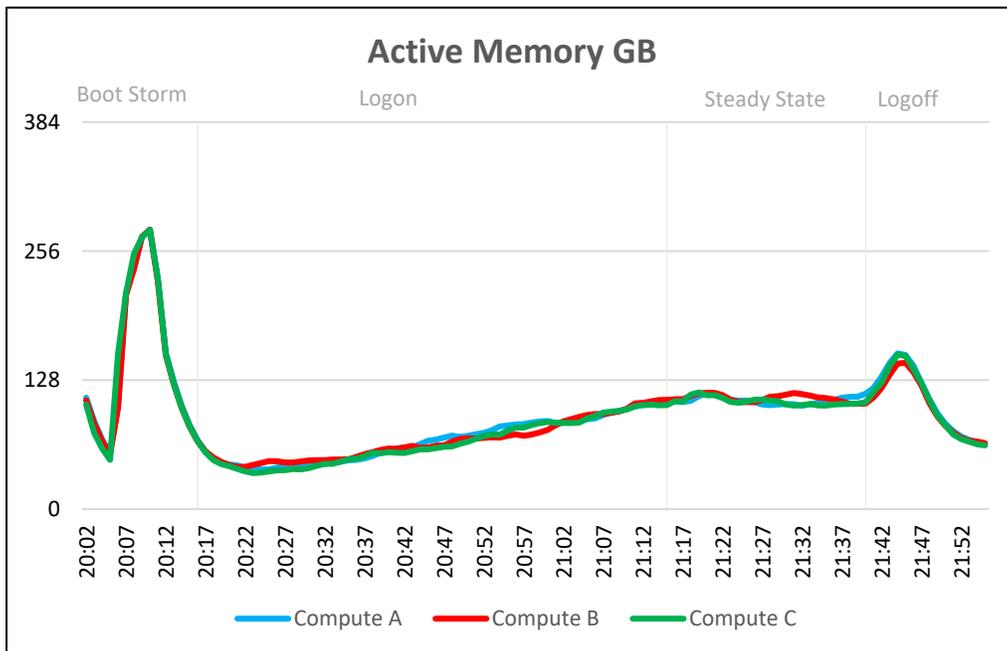
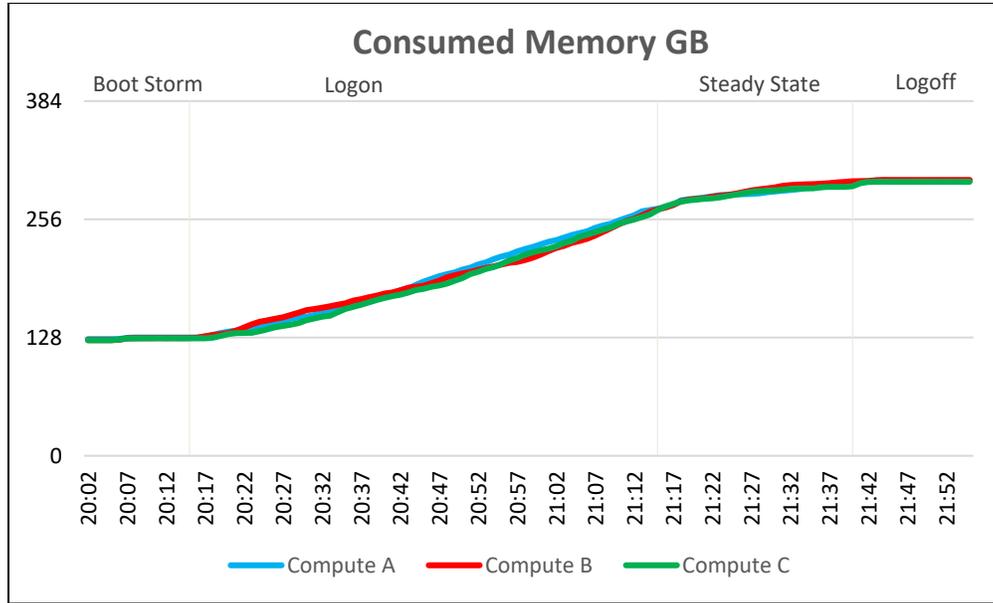
#### 6.3.4.1 Power Worker, 210 users ESXi 6.5, Horizon 7 linked clones

Each compute host was populated with 70 virtual machines per host. With all user virtual machines powered on and before starting test, the CPU usage was approximately 8%.

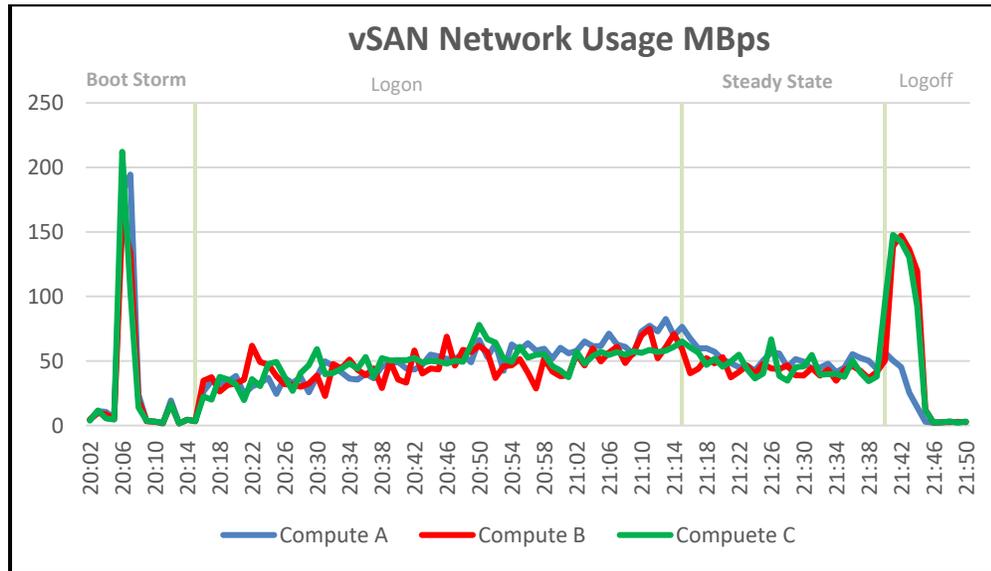
The below graph shows the performance data for 70 user sessions per host. The CPU reaches a steady state average of 48% across the three compute hosts during the test cycle when 70 users are logged on to each host.



In regards to memory consumption for the cluster, out of a total of 384 GB available memory per node, memory never reached maximum. The compute hosts reached a maximum memory consumption of 296 GB with active memory usage reaching a max of 277 GB. Average memory consumption of 285 GB and active memory usage 108 GB during steady state period. Zero ballooning occurred. There was no memory swapping on any of the hosts.

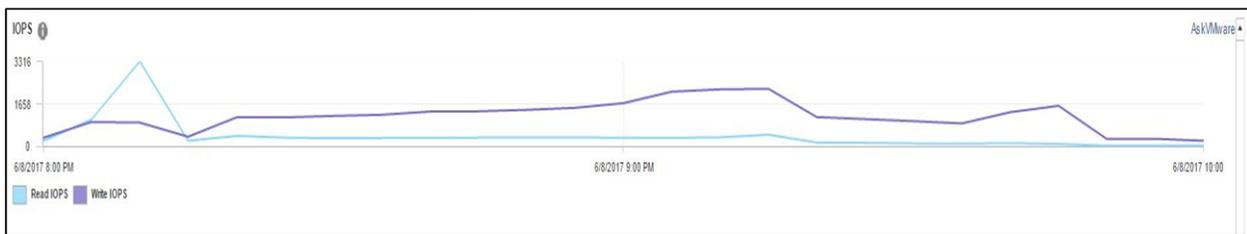


Network bandwidth is not an issue on this test run with a steady state average of approximately 48 MBps. One of the hosts reached a peak of 212 MBps during boot storm.

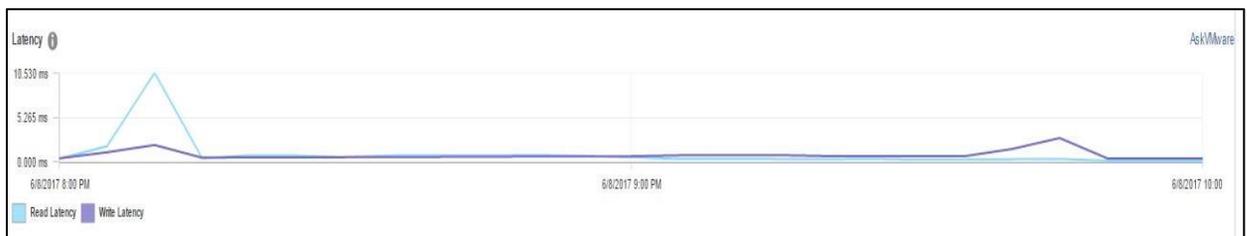


The IOPS graphs and IOPS numbers are taken from the vCenter web console and the graphs clearly show the initial boot storm during a reboot of all the desktops, followed by a settle period, logon, steady state and logoff phases. The graph displays the Disk IOPS figure for the VSAN cluster.

The cluster reached a maximum of 3,316 Disk IOPS during the boot storm before test start and 2,411 IOPS at the start of steady state.



Disk I/O Latency was not an issue. The maximum latency reached was approximately 10.5 ms during the boot storm and .8 ms at the beginning of steady state. This was well below the 20 ms threshold that is regarded as becoming potentially troublesome.





## 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 Andrew Breedy, Yinglong Jiang & Colin Byrne for their help in the validation of the solutions in this RA.

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

Colin Byrne, is a staff engineer in the Dell Wyse Datacenter engineering group. Colin has extensive experience in storage solutions with a primary focus on block storage technologies and server hardware.

Yinglong is a Senior Solution Engineer in Dell Wyse Datacenter engineering team, primarily responsible for testing the VDI solutions based on Dell solutions with a particular focus on software define storage and hyper-converged data center infrastructure.

Thanks to Mike Hayes from the Limerick CSC team for his help and support with the Graphics Functionality Testing that was completed on vSRN. 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](https://twitter.com/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 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](http://www.ExitTheFastLane.com).