

Dell EMC for VMware Horizon Cloud On-Premises Infrastructure – Reference Architecture

Integration of Horizon Cloud On-Premises Infrastructure with vSAN Ready Node and VxRail appliance clusters

Dell Cloud Client-Computing Engineering
September 2017

Revisions

Date	Description
August 2016	Initial release
September 2016	Updated Thin Client List
February 2017	Dell EMC VxRail, VSRN R630 support
April 2017	Section 6 Test results update
August 2017	Horizon Cloud On-Premises Update
September 2017	Thin Client Update

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
2 Solution Architecture Overview	7
2.1 Introduction.....	7
2.2 Physical Architecture Overview	7
2.3 Horizon Cloud On-Premises (HCOP).....	7
2.4 What is the Dell EMC VxRail Appliance?	9
2.4.1 What is included in Dell EMC VxRail 4.0?.....	10
2.5 What is a vSAN Ready Node (VSRN)?	12
2.6 Solution Layers.....	12
2.6.1 Networking.....	13
2.6.2 Compute and Management Server	14
2.6.3 Software Defined Storage (vSAN).....	15
3 Hardware Components	17
3.1 Network.....	17
3.1.1 Dell Networking S3048 (1Gb ToR Switch)	17
3.1.2 Dell Networking S4048 (10Gb ToR Switch)	18
3.2 Dell EMC VxRail VDI Optimized V Series Configurations.....	19
3.2.1 V470/ V470F-A3 Configuration.....	20
3.2.2 V470/ V470F-B5 Configuration.....	21
3.2.3 V470/ V470F-C7 Configuration	22
3.2.4 Dell EMC VxRail E Series Appliance (E460/ E460F)	23
3.3 vSAN Ready Nodes (VSRN)	23
3.3.1 vSRN R630.....	23
3.3.2 vSRN R730.....	27
3.3.3 vSRN R730XD.....	30
3.4 Dell Wyse Thin Clients	33
3.4.1 Wyse 3040 Thin Client (ThinOS, ThinLinux)	33
3.4.2 Wyse 5030 PCoIP Zero Client.....	33
3.4.3 Wyse 5040 AIO Thin Client with PCoIP	33

3.4.4	Wyse 5050 AIO PCoIP Zero Client	34
3.4.5	Wyse 7030 PCoIP Zero Client.....	34
3.4.6	Wyse 5060 Thin Client (ThinOS, ThinLinux, WES7P, WIE10)	34
3.4.7	Wyse 7020 Thin Client (Windows 10 IoT)	35
3.4.8	Wyse 7040 Thin Client with Windows Embedded Standard 7P.....	35
3.4.9	Latitude 3480 and 5280 Mobile Thin Clients (Win 10 IoT)	35
4	Software Components.....	37
4.1	VMware vSphere 6.x	37
4.2	VMware vSAN	38
4.3	HCOP Virtual Appliances	39
4.4	HCOP Control Plane	40
5	Solution Architecture for HCOP.....	41
5.1	Management Server Infrastructure.....	41
5.1.1	DNS	41
5.2	Storage Architecture Overview.....	42
5.2.1	vSAN Local Storage	42
5.3	Virtual Networking.....	42
5.3.1	HCOP VSRN Network Configuration	42
5.4	Scaling Guidance.....	45
5.5	Solution High Availability	46
5.5.1	vSAN HA/ FTT Configuration	46
5.5.2	vSphere HA	47
5.5.3	HCOP High Availability.....	47
6	Solution Performance and Testing.....	48
6.1	Summary	48
6.2	Test and performance analysis methodology.....	49
6.2.1	Testing Process.....	49
6.2.2	Resource monitoring	51
6.3	Test Configuration	52
6.3.1	Compute VM configurations	53
6.3.2	Platform Configurations	53
6.4	Test Results and Analysis	53
6.4.1	HCOP VSRN/ VxRail C7	55

Acknowledgements.....71
About the Authors72

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 Cloud On-Premises on VMware vSphere /VMware vSAN.

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 component selection guidance.

2 Solution Architecture Overview

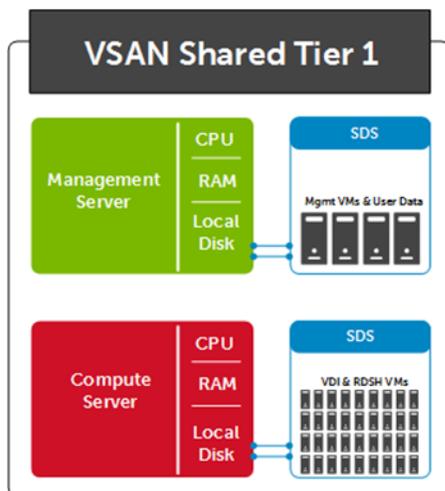
2.1 Introduction

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:

- Pooled Virtual Desktops (Non-persistent)

2.2 Physical Architecture Overview

The core vSAN Ready Node (vSRN) and VxRail 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 x 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.



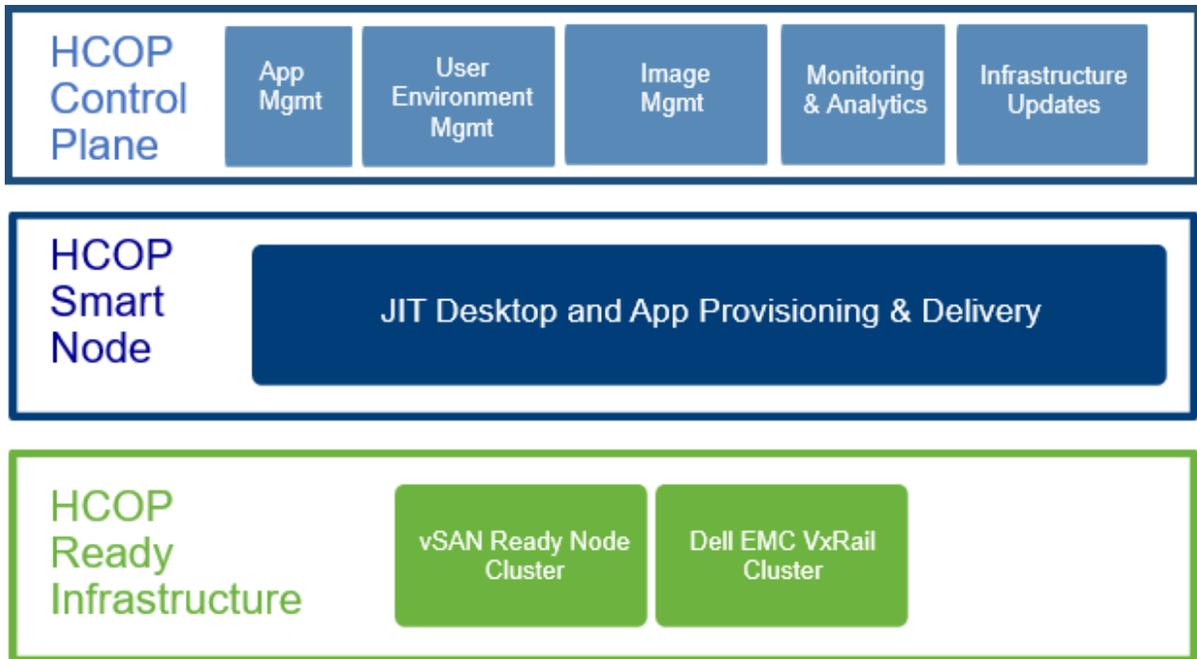
2.3 Horizon Cloud On-Premises (HCOP)

With a Horizon Cloud with On-Premises Infrastructure environment, end users can securely access their desktops and applications from any device. After you deploy the Horizon Cloud Node, you can use the Administration Console to configure desktop and application assignments.

HCOP consists of a cloud service, Horizon Cloud, which pairs with an on-premises component called the Horizon Cloud Node.

- **Horizon Cloud**– A control plane hosted in the cloud by VMware for management of virtual desktops, applications, and user profiles hosted on on-premises infrastructure.
- **Horizon Cloud Node** – Optimized hardware that is connected to the cloud control plane by way of integrated Horizon Cloud with On-Premises Infrastructure connector software and configured for the Horizon Cloud with On-Premises Infrastructure environment. vSAN Ready Nodes or Dell EMC VxRail appliances are types of optimized hardware you can configure as Horizon Cloud Smart Nodes for use with the Just-in-Time (JIT) Desktop and App Provisioning mechanisms.
- **Horizon Cloud Ready Infrastructure** – vSAN Ready & VxRail Appliance nodes that run on the customer’s site, connected to the Cloud Control Plane by way of integrated Horizon Air Hybrid-Mode Software. (see section 4.4)

The top layer below represents the Cloud Control Plane and the lower two layers show the Horizon Cloud Node, consisting of the Smart Node & vSAN Ready Node layers.



2.4 What is the Dell EMC VxRail Appliance?



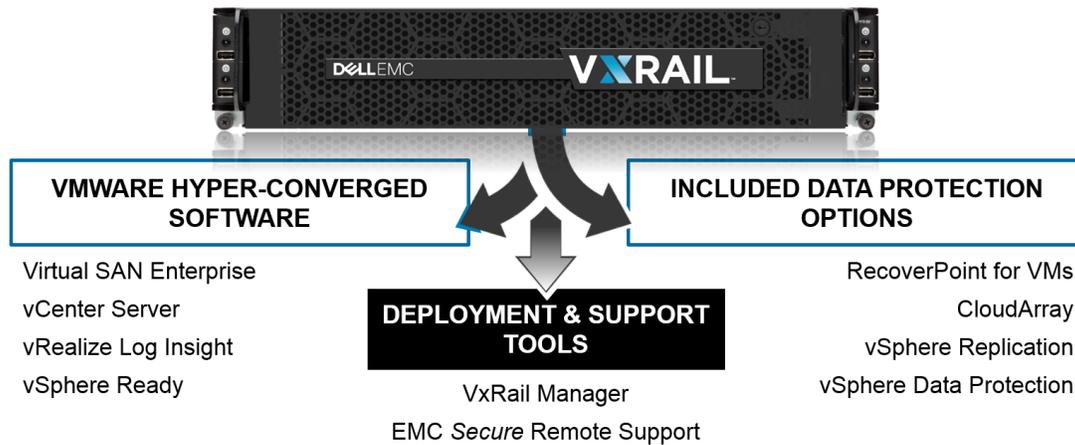
The Dell EMC VxRail appliances are very powerful Hyper Converged Infrastructure Appliances (HCIA) delivered in 1U/2U Rack Building Blocks. The appliances are built using VMware vSphere and vSAN technology with EMC software. VxRail allows the seamless addition of additional nodes to a cluster from the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

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

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

2.4.1 What is included in Dell EMC VxRail 4.0?

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



VxRail includes the following software from VMware and Dell EMC:

- VMware vSAN
- VMware vCenter
- VMware ESXi
- VMware vRealize Log Insight
- Dell EMC VxRail Manager

Optional software licenses included with VxRail:

- VMware vSAN
- VMware vCenter
- VMware vRealize Log Insight

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

Optional Software

VxRail also includes optional licenses for Dell EMC CloudArray and RecoverPoint which are not pre-installed but the customer is entitled to use of this software.

CloudArray

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

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

RecoverPoint

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

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

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

For more information on vSphere Data Protection visit [here](#)

2.5 What is a vSAN Ready Node (VSRN)?

A vSAN Ready Node 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 as any of the Dell configurations can be selected to meet your needs from the vSRN compatibility page. [Link](#)

vSRN also comes with pre-selected Horizon OEM VMware SKU's 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.

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

2.6 Solution Layers

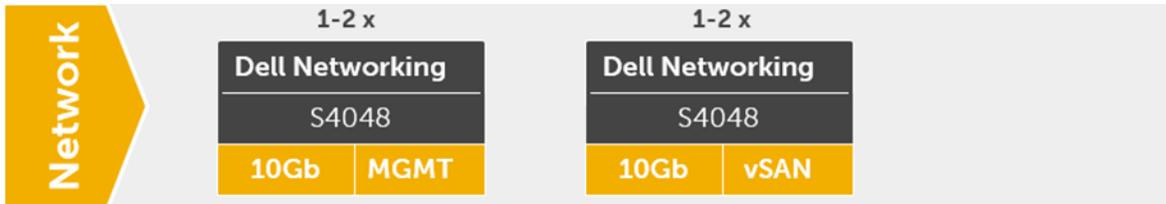
The vSAN Ready Node Solution leverages a core set of hardware and software components consisting of five primary layers:

- Networking Layer
- Compute & Management Server Layer
- Storage Layer (VMware vSAN)
- Thin Client Layer (please refer to section 3.3)

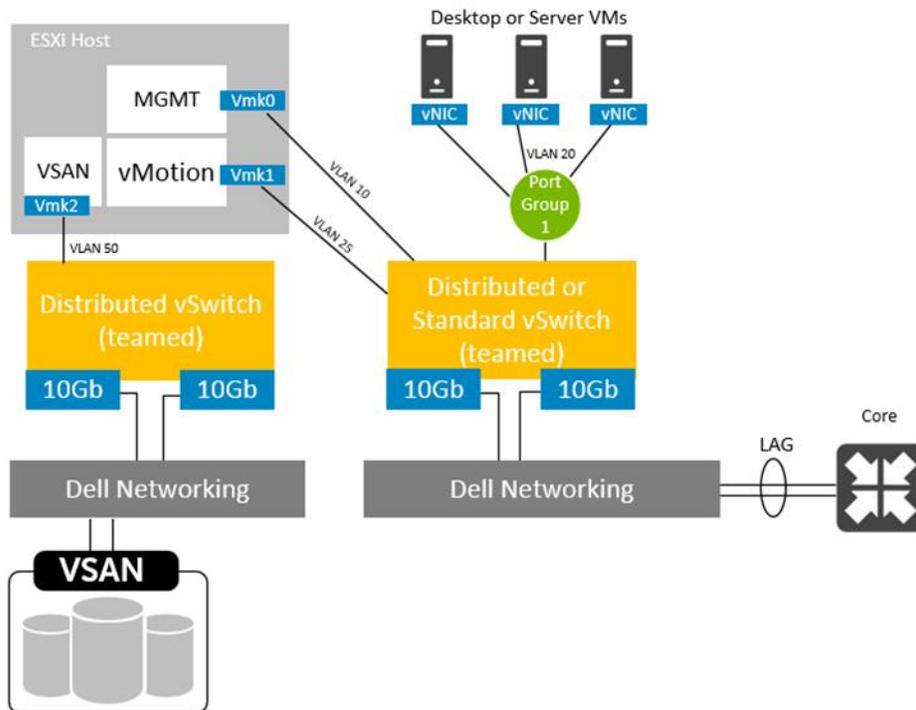
These components have been integrated and tested to provide the optimal balance of high performance and lowest cost per user. The vSAN Ready Node stack is designed to be cost effective allowing IT departments to implement high-performance fully virtualized desktop environments.

2.6.1 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.



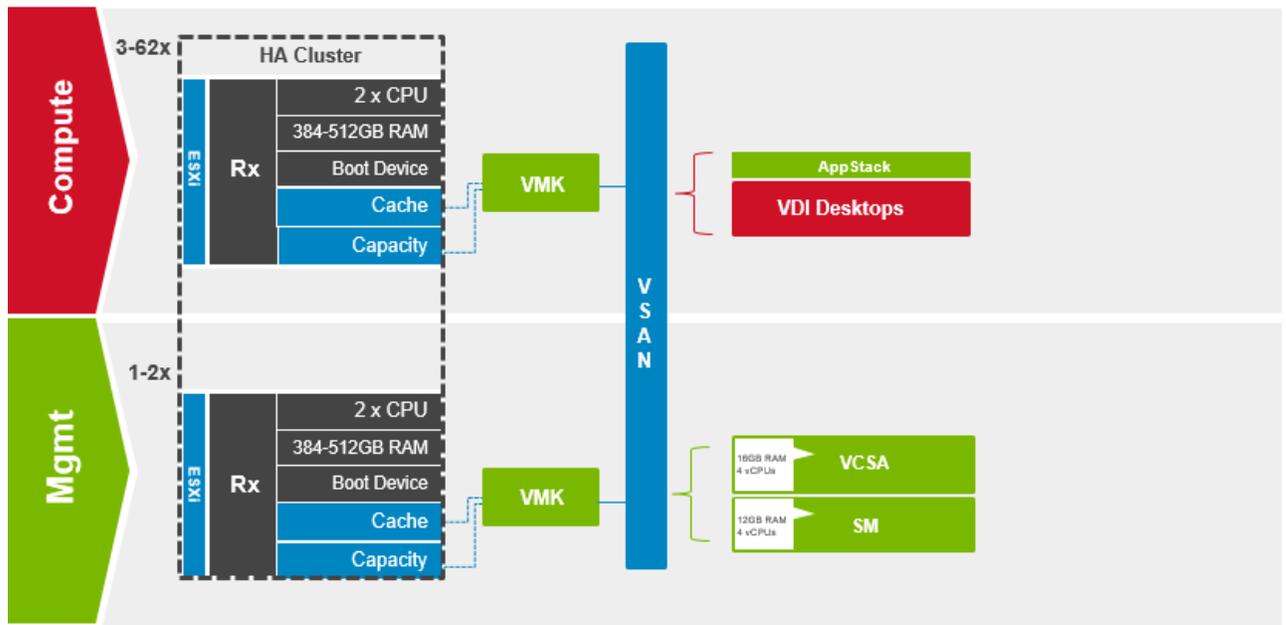
Additional switches are added and stacked as required to provide High Availability for the Network layer. 1Gb connectivity for iDRAC/ remote management can make use of an existing 1Gb ToR switch or adding a Dell Networking 1Gb 48-port switch for this function.



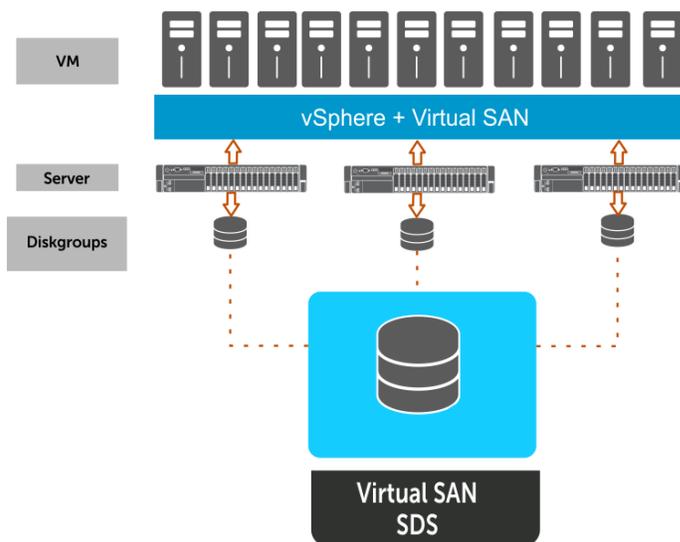
2.6.2 Compute and Management Server

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

The HCOP Cluster consist of Nodes which essentially perform management and compute functions. When deploying a pool with HCOP the VMs are spread evenly across all available Nodes in the Cluster. The maximum amount of VMs supported in this release is 1000 so will be unable scale to the 64 Nodes per Cluster. The minimum amount of Nodes recommended for HCOP is 4. The SmartNode(SM) appliance on the management node contains the management services which communicate to the HCOP control plane.



2.6.3 Software Defined Storage (vSAN)



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

From a hardware perspective, at least three ESXi hosts (four recommended) are required to form the vSAN cluster. For a hybrid configuration, 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 vSAN cluster. 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 SSDs.

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

Capacity devices have two roles in VMware vSAN. They make up the capacity of the VMware vSAN datastore as well as making up components for a stripe width. SAS and NL-SAS drives are supported.

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

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

3 Hardware Components

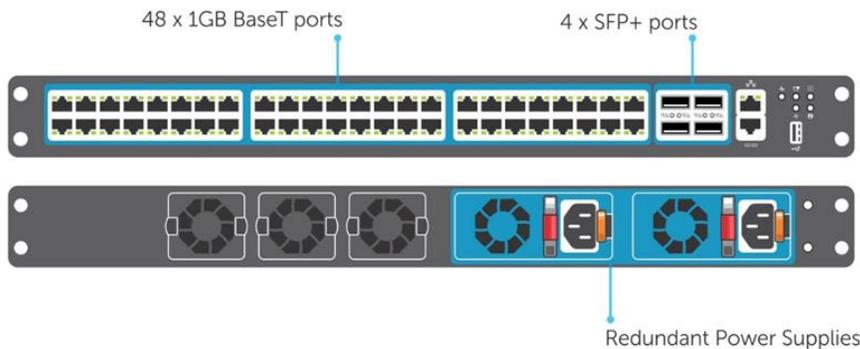
3.1 Network

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.

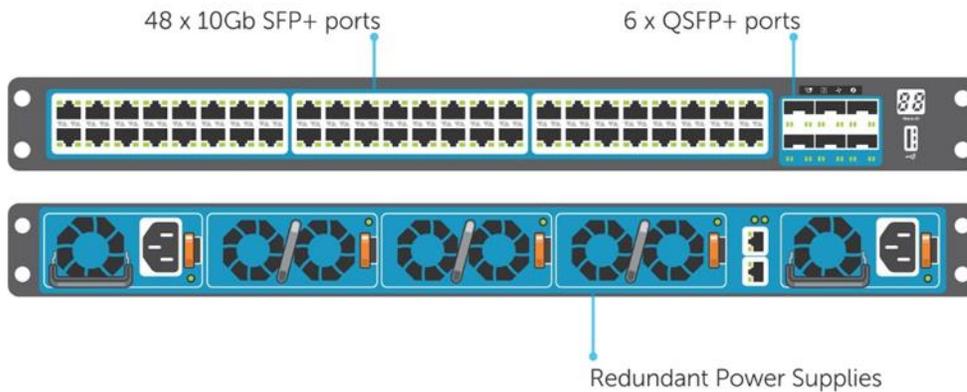
Model	Features	Options	Uses
Dell Networking S3048-ON	48 x 1000BaseT	Redundant hot-swap PSUs & fans	1Gb connectivity (iDRAC)
	4 x 10Gb SFP+	VRF-lite, Routed VLT, VLT Proxy Gateway	
	Non-blocking, line-rate performance	User port stacking (up to 6 switches)	
	260Gbps full-duplex bandwidth	Open Networking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



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



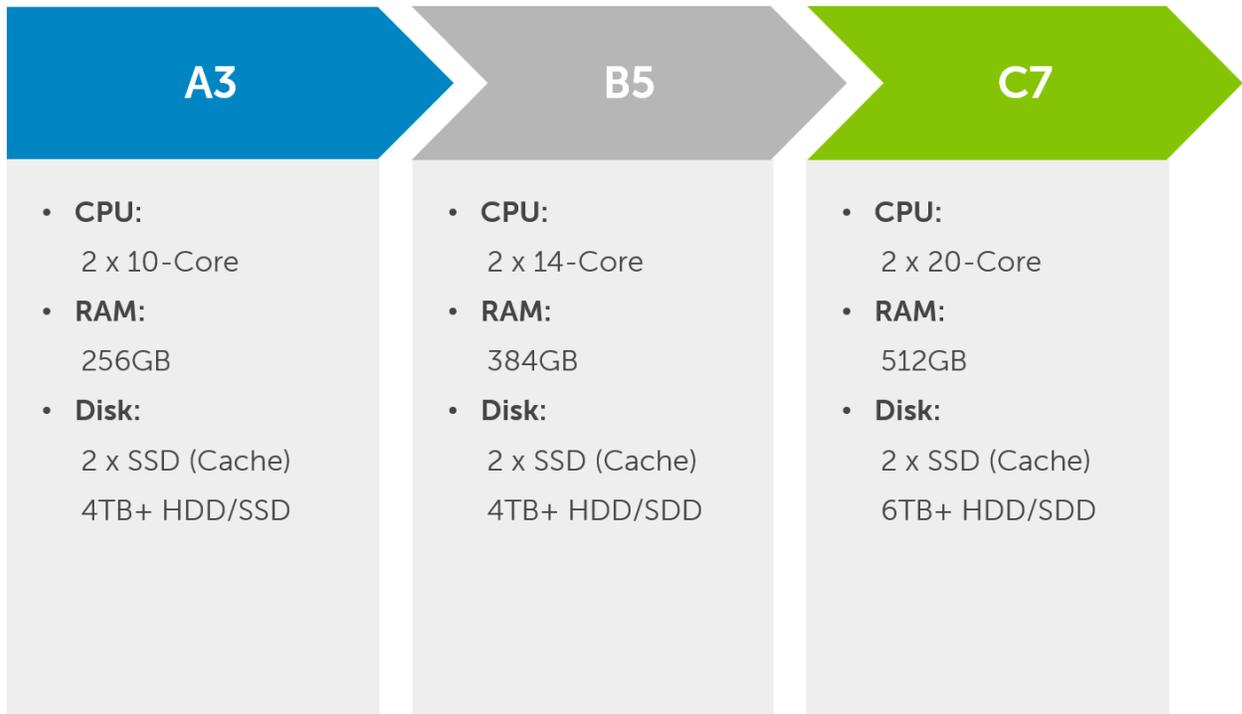
3.2 Dell EMC VxRail VDI Optimized V Series Configurations

The VDI-optimized 2U/ 1Node appliance with GPUs is optimized for graphics-intensive desktop deployments. The V Series nodes can be configured with or without GPUs as detailed in the A3, B5 & C7 optimized configurations below.

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

- **A5** configuration is perfect for small scale, POC or low density cost-conscious environments.
- **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.

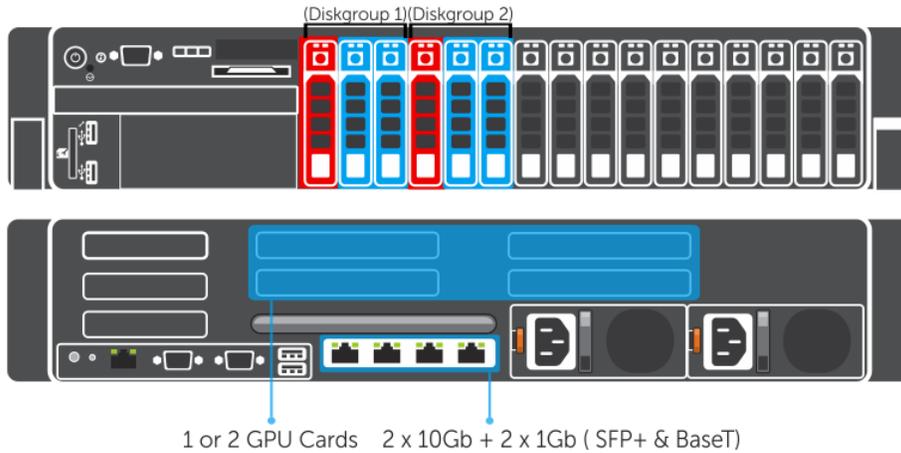
Note: GPU cards are not yet supported for Horizon Cloud On-Prem.



3.2.1 V470/ V470F-A3 Configuration

The V470/ V470F-A3 configuration is equipped with 256GB of Memory and 2 x E5-2640v4 CPUs with the option of 2 x NVIDIA M60 GPU cards. The drive configuration consists of two disk groups, 1 cache disk and 2 capacity disks in each. The cache disks are populated in slots 0 & 4 as shown below.

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

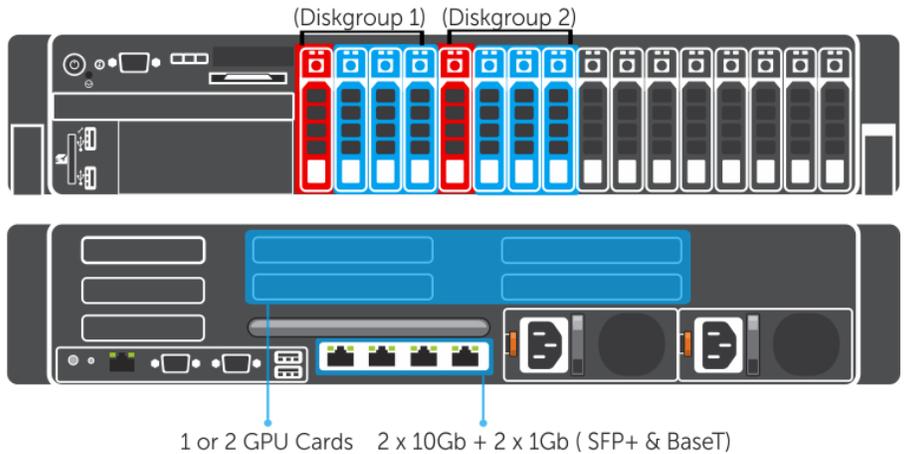


V470	A3
CPU	2 x E5-2640v4 (10C, 2.4GHz)
Memory	16 x 16GB 2400MT/s RDIMMs Effective speed: 2400MT/s @ 256GB
GPU Boards	2 x NVIDIA GRID M60 vGPU Profile: M60-xQ
Storage Ctrls	PERC H330 – no RAID
Storage	64GB SATADOM (ESXi Boot) 16GB SD Module (PowerTools) 2 x 400GB SSD 2.5" (Cache) 4 x 1.2TB HDD 2.5" (Capacity-Hybrid)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

3.2.2 V470/ V470F-B5 Configuration

The V470/ V470F-B5 configuration is equipped with 384GB of Memory and 2 x E5-2660v4 CPUs with the option of 2 x NVidia M60 GPU cards. The drive configuration consist of two disk groups, 1 cache disk and 3 capacity disks in each. The cache disks are populated in slots 0 & 4.

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

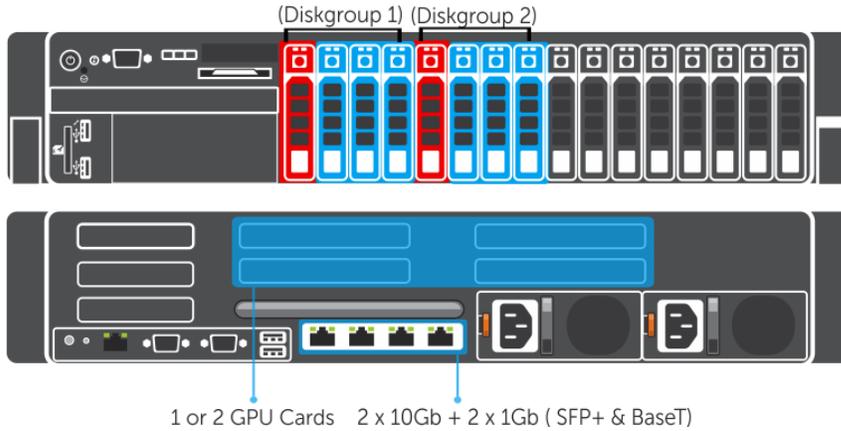


V470	B5
CPU	2 x E5-2660v4 (14C, 2.0GHz)
Memory	24 x 16GB 2400MT/s RDIMMs Effective speed: 2133MT/s @ 384GB
GPU Boards	2 x NVIDIA GRID M60 vGPU Profile: M60-xQ
Storage Ctrlrs	PERC H330 – no RAID
Storage	64GB SATADOM (ESXi Boot) 16GB SD Module (PowerTools) 2 x 400GB SSD 2.5" (Cache) 4 x 1.2TB HDD 2.5"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

3.2.3 V470/ V470F-C7 Configuration

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

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

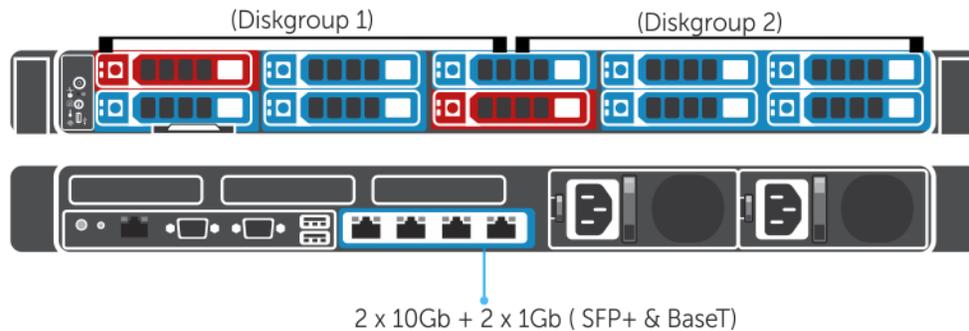


V470	C7
CPU	2 x E5-2698v4 (20C, 2.2GHz)
Memory	16 x 32GB 2400MT/s RDIMMs Effective speed: 2400MT/s @ 512GB
GPU Boards	2 x NVIDIA GRID M60 vGPU Profile: M60-xQ
Storage Ctrls	PERC H330 – no RAID
Storage	64GB SATADOM (ESXi Boot) 16GB SD Module (PowerTools) 2 x 800GB SSD 2.5" (Cache) 6 x 1.2TB HDD 2.5" (Capacity-Hybrid) Or 4 x 1.92TB SSD 2.5" (Capacity-AF)
Network	Intel X520 DP 10Gb SFP+ & i350 DP 1Gb NDC
iDRAC	iDRAC8 Ent w/ vFlash, 8GB SD
Power	2 x 1100W PSUs

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

The E Series is the entry level platform, offered with single or dual socket processors in a 1U configuration. These configurations are geared toward basic workloads, remote or branch office scenarios. The minimum amount of memory needed for a one CPU configuration is 64GB and the maximum supported is 768GB. The minimum for a two socket CPU configuration is 128GB and the maximum supported is 1536GB. The minimum drive configuration is 1 x cache disk and 1 x capacity disk in a single disk group configuration. The maximum for this configuration is 2 x cache disks and 8 x capacity disks in a two disk group configuration. Slot 0 and Slot 5 are used for Cache disks.

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



3.3 vSAN Ready Nodes (VSRN)

vSAN Ready Nodes are preconfigured and optimized server nodes with certified components designed to remove the guess work from vSAN deployments.

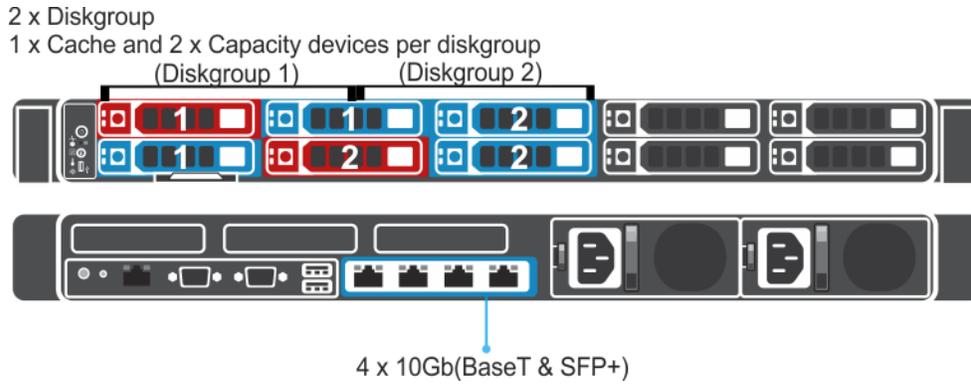
Note: All-Flash configurations have not yet been validated with deduplication and compression enabled.

3.3.1 vSRN R630

The Dell R630 is a 1U platform with a broad range of configuration options. Each optimized 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).

3.3.1.1 vSRN R630-A3 Configuration

There are two diskgroups in this configuration, as depicted below, consisting of 1 x Cache SSD and 2 x Capacity HDD/SSD disks per diskgroup. There are two ESXi boot options, a single 64GB SATADOM or a 16GB mirrored pair of SD modules. Each platform can be outfitted with SFP+ or BaseT NICs.



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

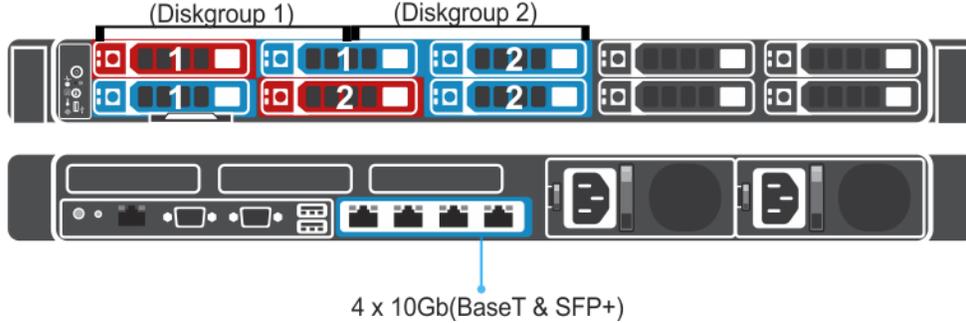
R630	A3
CPU	2 x E5-2630v4 (10C, 2.2GHz)
Memory	16 x 16GB 2400MT/s RDIMMs Effective speed: 2400MT/s @ 256GB
Storage Ctrlrs	PERC H730/H330 – no RAID
Boot Device	64GB Satadom or 2 X 16GB SD Module
Storage Hybrid	2 x 400GB SSD (Cache) 4 x 1.2TB HDD (Capacity)
Storage All-Flash	2 x 400GB SSD (Cache) 2 x 1.92TB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

3.3.1.2 vSRN R630-B5 Configuration

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

2 x Diskgroup

1 x Cache and 2 x Capacity devices per diskgroup



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

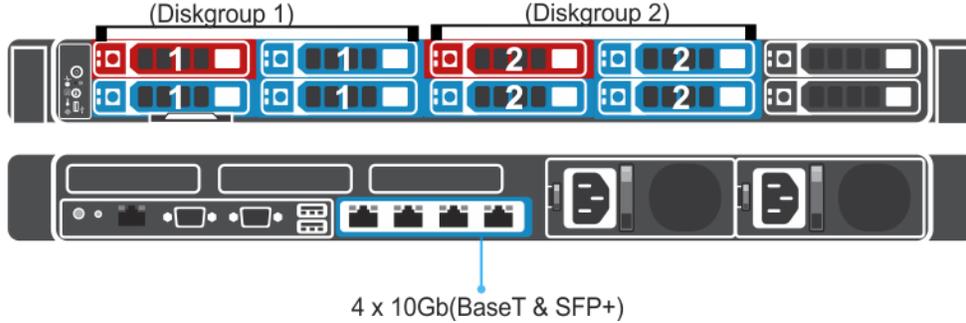
R630	B5
CPU	2 x E5-2660v4 (14C, 2.0GHz)
Memory	24 x 16GB 2400MT/s RDIMMs Effective speed: 2133MT/s @ 384GB
Storage Ctrlrs	PERC H730/H330 – no RAID
Boot Device	64GB Satadom or 2 X 16GB SD Module
Storage Hybrid	2 x 400GB SSD (Cache) 4 x 1.2TB HDD (Capacity)
Storage All-Flash	2 x 400GB SSD (Cache) 4 x 1.92TB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

3.3.1.3 vSRN R630-C7 Configuration

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

2 x Diskgroup

1 x Cache and 3 x Capacity devices per diskgroup



This configuration 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 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 750W PSUs

3.3.2 vSRN R730

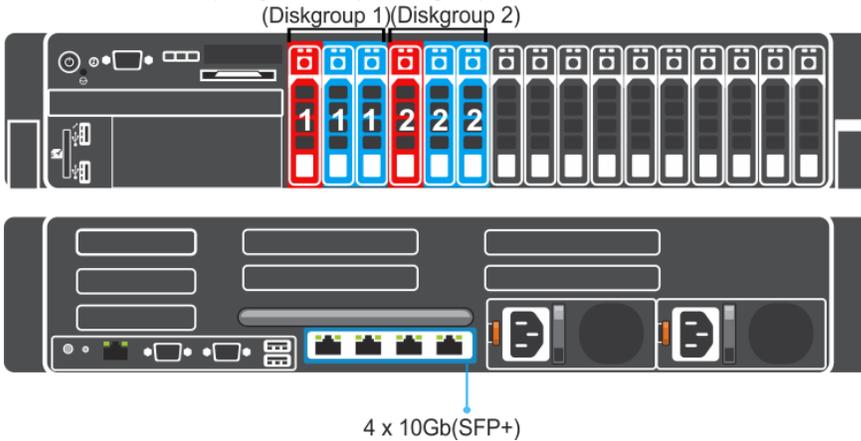
This dual socket CPU platform runs the fastest Intel Xeon E5-2600 v4 family of processors, can host up to 24 DIMMs of DDR4 RAM, supports up to 16 x 2.5" SAS disks and can be outfitted with 2 double-wide NVIDIA GPU accelerators.

3.3.2.1 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, as depicted below, consisting of 1 x Cache SSD and 2 x Capacity HDD/SSD disks per diskgroup.

2 x Diskgroups

1 x Cache and 2 x Capacity devices per diskgroup



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

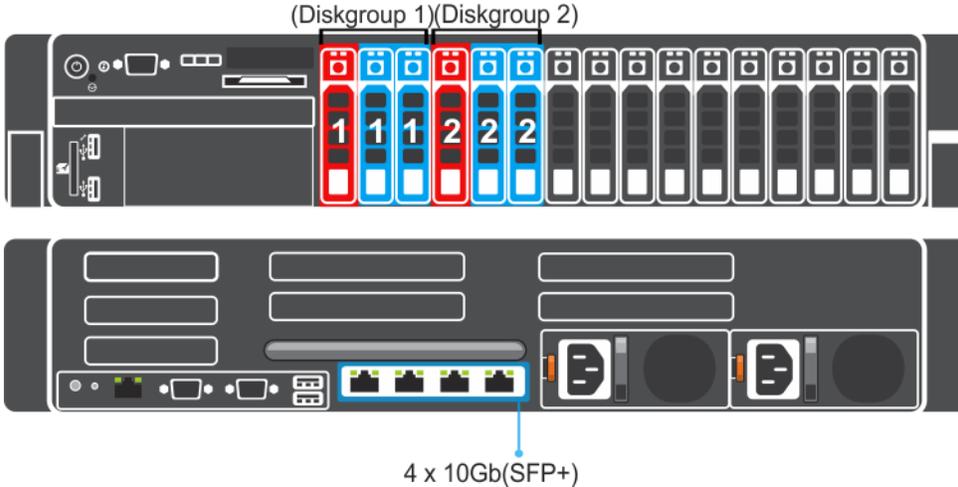
R730	A3
CPU	2 x E5-2630v4 (10C, 2.2GHz)
Memory	16 x 16GB 2400MT/s RDIMMs Effective speed: 2400MT/s @ 256GB
Storage Ctrlrs	PERC H730/H330 – no RAID
Boot Device	64GB Satadom or 2 X 16GB SD Module
Storage Hybrid	2 x 400GB SSD (Cache) 4 x 1.2TB HDD (Capacity)
Storage All-Flash	2 x 400GB SSD (Cache) 2 x 1.92TB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

3.3.2.2 vSRN R730-B5 Configuration

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

2 x Diskgroups

1 x Cache and 2 x Capacity devices per diskgroup



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

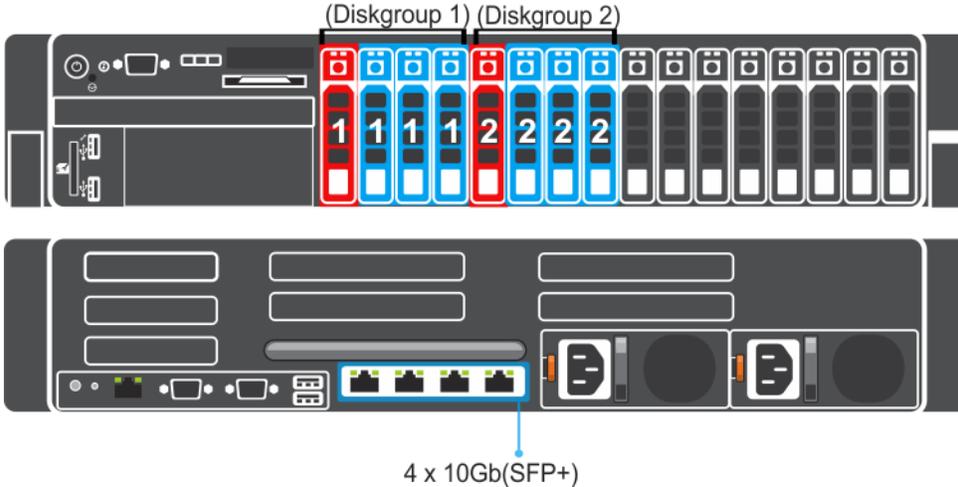
R730	B5
CPU	2 x E5-2660v4 (14C, 2.0GHz)
Memory	24 x 16GB 2400MT/s RDIMMs Effective speed: 2133MT/s @ 384GB
Storage Ctrlrs	PERC H730/H330 – no RAID
Boot Device	64GB Satadom or 2 X 16GB SD Module
Storage Hybrid	2 x 400GB SSD (Cache) 4 x 1.2TB HDD (Capacity)
Storage All-Flash	2 x 400GB SSD (Cache) 4 x 1.92TB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

3.3.2.3 vSRN R730-C7 Configuration

There are two diskgroups in this configuration, as depicted below, consisting of 1 x Cache SSD and 3 x Capacity HDD/SSD disks per diskgroup.

2 x Diskgroups

1 x Cache and 3 x Capacity devices per diskgroup



This configuration 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 Ctrlrs	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

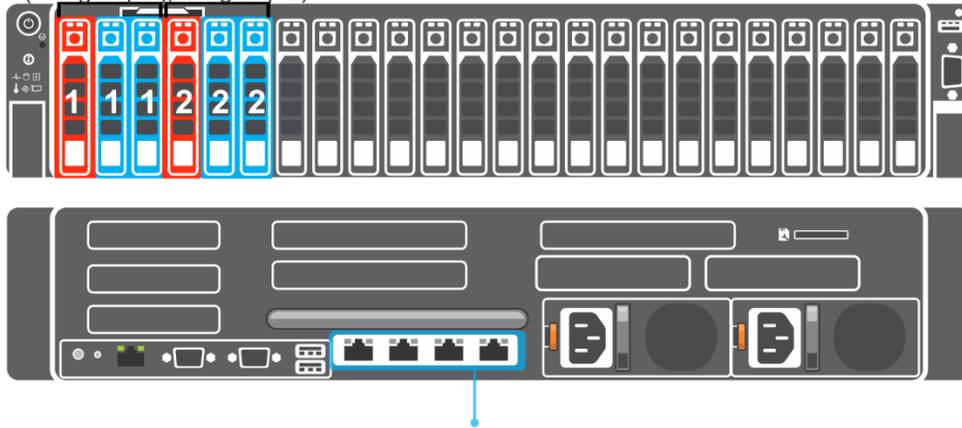
3.3.3 vSRN R730XD

The R730XD can be configured as hybrid or all-flash. For all-flash the cache tier uses write intensive SSD and mixed use SSDs for the capacity tier. There are two boot options, a single 64GB SATADOM or a 16GB mirrored pair of SD modules.

3.3.3.1 vSRN R730XD-A3 Configuration

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

2 x Diskgroups
 1 x Cache SSD and 2 x Capacity devices per diskgroup
 (Diskgroup 1 | Diskgroup 2)



4 x 10Gb (BaseT or SFP+)

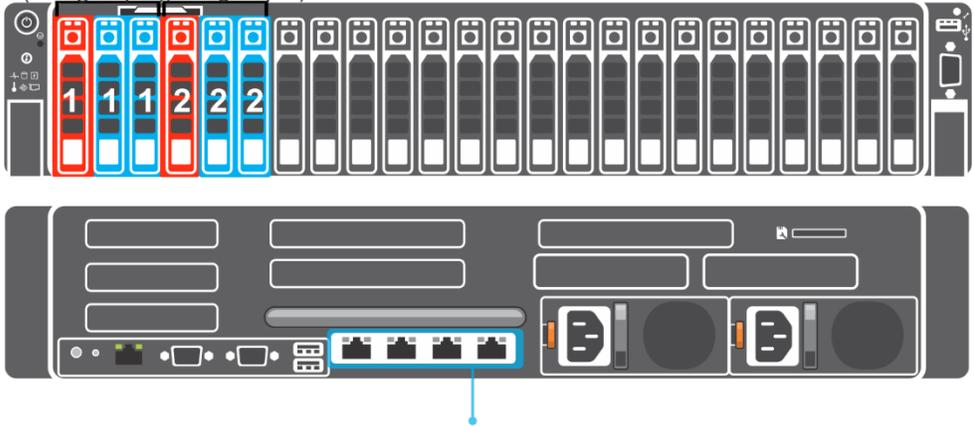
This would be classed as a vSRN HY-6 (Hybrid) or AF-4 (All-Flash) configuration on the vSRN HCL.

R730XD	A3
CPU	2 x E5-2630v4 (10C, 2.2GHz)
Memory	16 x 16GB 2400MT/s RDIMMs Effective speed: 2400MT/s @ 256GB
Storage 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

3.3.3.2 vSRN R730XD-B5 Configuration

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

2 x Diskgroups
 1 x Cache SSD and 2 x Capacity devices per diskgroup
 (Diskgroup 1 | Diskgroup 2)



4 x 10Gb (BaseT or SFP+)

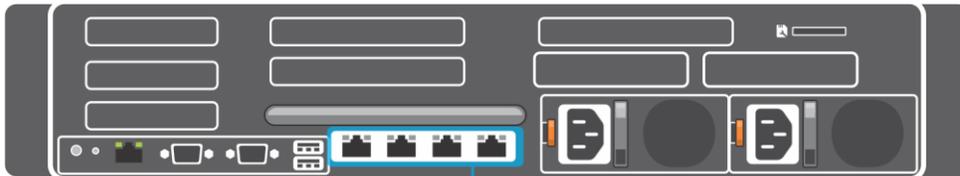
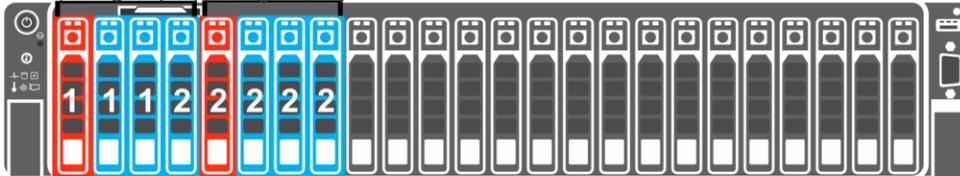
This would be classed as a vSRN HY-6 (Hybrid) or AF-4 (All-Flash) configuration on the vSRN HCL.

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 (Cache) 4 x 1.2TB HDD (Capacity)
Storage All-Flash	2 x 400GB SSD (Cache) 4 x 1.92TB SSD (Capacity)
Network	4 x 10Gb SFP+
iDRAC	iDRAC8 Ent w/ vFlash, 8GB SD
Power	2 x 750/1100W PSUs

3.3.3.3 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.

2 x Diskgroups
 1 x Cache SSD and 3 x Capacity devices per diskgroup
 (Diskgroup 1) (Diskgroup 2)



4 x 10Gb (BaseT or SFP+)

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 Ctrlrs	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

3.4 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 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 all Dell Wyse Datacenter solutions, 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 VM or newer 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](#).

4.2 VMware vSAN

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

Deduplication and compression: vSAN 6.2 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 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 6.2 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.

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

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

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.

For more information on what’s new in vSAN [Link](#).

4.3 HCOP Virtual Appliances

The first configuration step for the HCOP Appliances is to deploy the Horizon Air Link (HAL) OVA to the vSRN/VxRail Cluster. This appliance provides the initial connection the HAL control plane (see section 4.3). The HAL once deployed and configured is then followed by the onboarding process. This process requires the configuration details of the vSRN Cluster/SmartNode to be entered and once complete will then deploy the HCOP appliance from the Cloud to the on premise vSRN/VxRail Cluster.

The management tier of a Horizon Cloud with On-Premises Infrastructure deployment has been consolidated into a single virtual appliance. The Horizon Cloud Node appliance manages all of the critical functions including App Volumes, Instant Clone creation, and communication with the Horizon Cloud control plane. This change was made to keep the Horizon Cloud Node footprint small and efficient.

The appliance that reside on the vSRN Cluster consist of the SmartNode Manager, and with the 1.3 release the Tenant Appliance and Resource Manager services are now running on the one smartnode appliance. These VMs connect to the cloud service and this intercommunication between the on-premise vSRN and the Cloud Management service manage the HAH infrastructure.

The SmartNode on-prem appliance consists of:

Tenant Service: This component brokers the end user connection to the VDI pool. The user connects via the Horizon client (same as Horizon View client).

SmartNode Service: This is the component that connects the SmartNode to the cloud service and also hosts the App Volumes manager service.

Resource Manager Service: This component has the Instant Clone Manager and Hypervisor manager services.

The administrator does not have direct access to these appliances and all management is done from the Cloud Plane, which includes deploying Desktop Pools, Appstacks and managing the assignments of which groups/users have access to these.

The user connects to the desktop pool via a floating IP address (to the broker) which acts as the broker and this is located on the tenant appliance. The floating IP will failover to the standby appliance in the case of a failure on the controlling appliance.

4.4 HCOP Control Plane

Horizon Cloud Services with Horizon Cloud Manager is a control plane hosted by VMware in the cloud; it enables the central orchestration and management of virtual desktops, applications, and user profiles on on-premises vSRN Cluster. Horizon Cloud Services is a set of micro services which enables end-to-end management workflows such as inventory management and App Volumes.

VMware is responsible for hosting the service and providing feature updates and enhancements for a true software-as-a-service experience. These updates are deployed to the on-premise vSRN Cluster by a specified or suitable time for the customer. Only the application metadata is stored in the Cloud everything else is stored on-premise.

The Cloud Control Plane also hosts a common management UI called Horizon Cloud Manager. The Horizon Cloud Manager is accessible by way of all major browsers and provides IT administrators a single location for managing desktop images, applications, user data, profiles, and assignments. The Horizon Cloud Manager is accessible from anywhere at any time, providing maximum flexibility.

The Cloud Control Plane was built to be a highly scalable, multi-tenant management service, leveraging the architecture of the Horizon Cloud On-Premises Platform.

5 Solution Architecture for HCOP

5.1 Management Server Infrastructure

The Management role requirements for the base solution are summarized below. Use data disks for role-specific application files and data, logs, IIS web files, etc. in the Management directory located on the vSAN datastore.

Role	vCPU	RAM (GB)	NIC	OS + Data vDisk (GB)	Tier 2 Volume (GB)
VMware vCenter Server Appliance	2	8	1	290	-
SmartNode Appliance	2	12	2	30	-
Total	4	32	3	310	

5.1.1 DNS

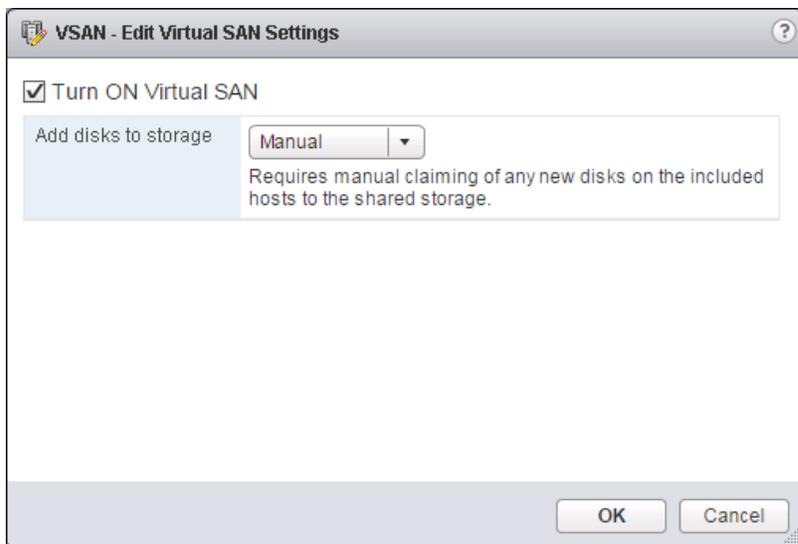
DNS plays a crucial role in the environment not only as the basis for Active Directory but will be used to control access to the various VMware software components. All hosts, VMs and consumable software components need to have a presence in DNS, preferably via a dynamic and AD-integrated namespace. 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.

5.2 Storage Architecture Overview

5.2.1 vSAN Local Storage

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 vSAN is setup in automatic mode, then all empty local disks will be used in the creation of the shared data store. If configured in manual mode, then disks can be manually selected during the creation of the data store.



5.3 Virtual Networking

5.3.1 HCOP VSRN Network Configuration

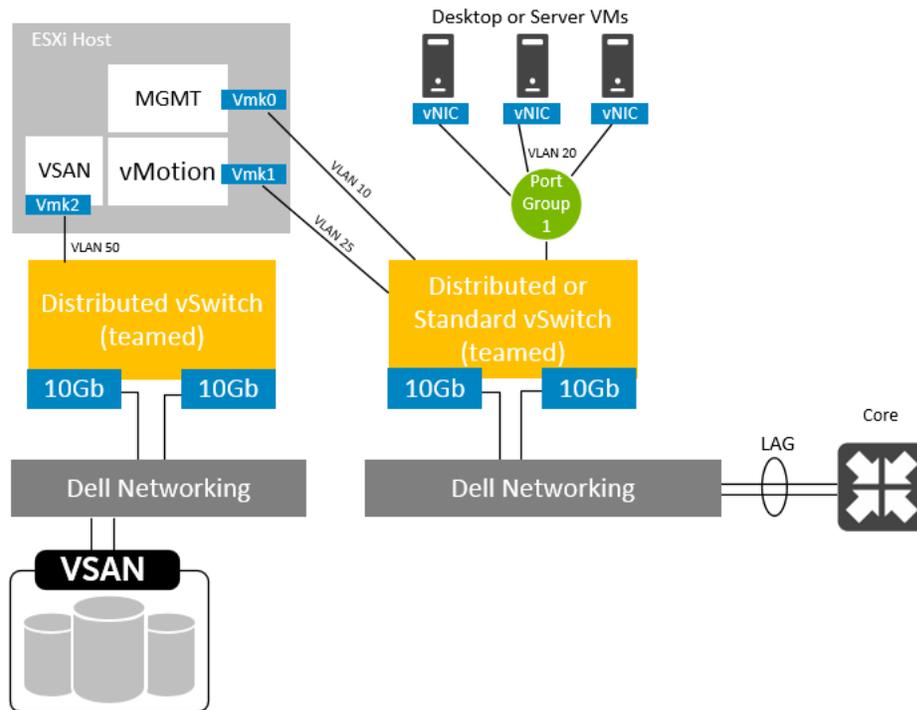
The network configuration for the Dell VSRN appliances utilizes a 10Gb converged infrastructure model. The management, vMotion and VDI traffic are configured across 2 x 10Gb NICs configured in an active/active team ensuring that the vSAN VLAN is configured across 2 physical 10Gb interfaces.

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 Migration. The following outlines the VLAN requirements for the Compute and Management hosts in this solution model:

- Management and Compute Layer Networks
 - Management VLAN: Configured for hypervisor infrastructure traffic – L3 routed via core switch
 - VDI VLAN: Configured for VDI session traffic and tenant appliances – L3 routed via core switch
 - Backbone/Appliance VLAN: Private VLAN-L2 switched only (IPv4 & IPv6 enabled)
 - vSAN VLAN: Configured for vSAN traffic – L2 switched only via 10 Gb 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

The following image shows the VMkernel adapters for the management network (vmk0), vMotion network (vmk1) on a standard or distributed vSwitch and vSAN Network (vmk2) on a distributed vSwitch.



5.3.1.1 vSphere Distributed Switches

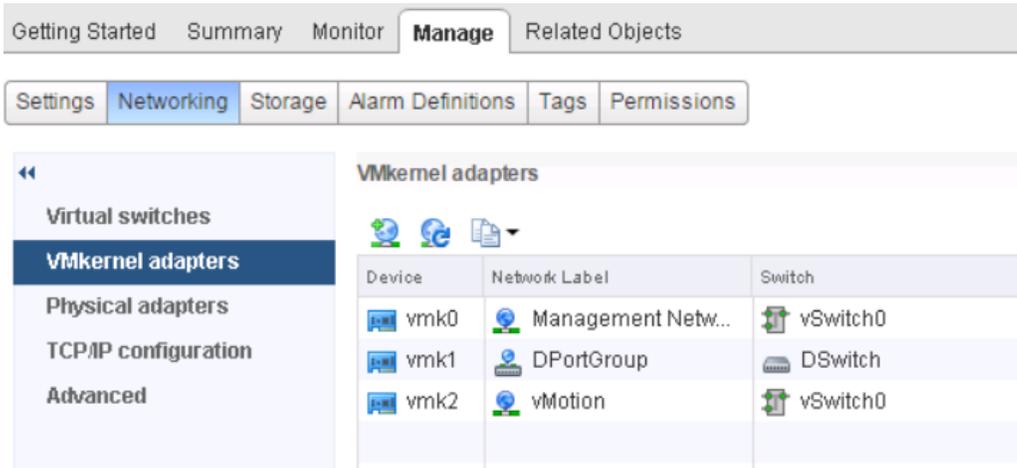
The VMware vSphere Distributed Switch (vDS) brings a consistent configuration across all hosts. The vDS is configured at the vCenter level and provides central management and monitoring to all hosts configured on the vDS.

vDS 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 VMkernel ports. Network share values should be configured equally among the VMkernel port groups that share a physical set of network adapters.

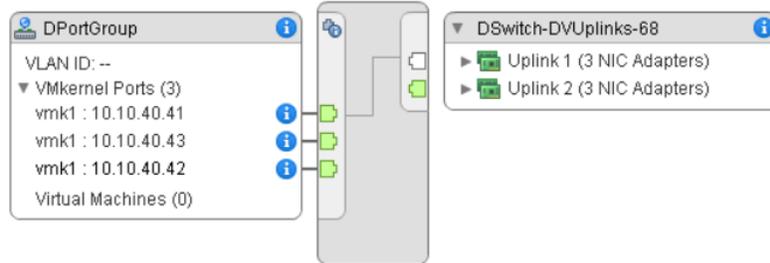
vSAN cluster networking includes at least two VMkernel ports, one for management traffic and one for vSAN traffic. If vMotion, Storage vMotion or High Availability functionality is required in addition, a third VMkernel port should be configured for this.

vSAN traffic can be used on 1Gb networks as well as 10Gb networks for Hybrid configuration but 10Gb is recommended and is required for all-flash configurations. Standard switch configuration can be used for Proof of Concept, while VMware distributed virtual switch configuration is highly recommended for production.

Network VMkernel adapter configuration for the host management traffic uses a 10Gb network with standard vSwitch. It is recommended that the network configuration for the vSAN storage is a 10Gb network with distributed vSwitch configuration.

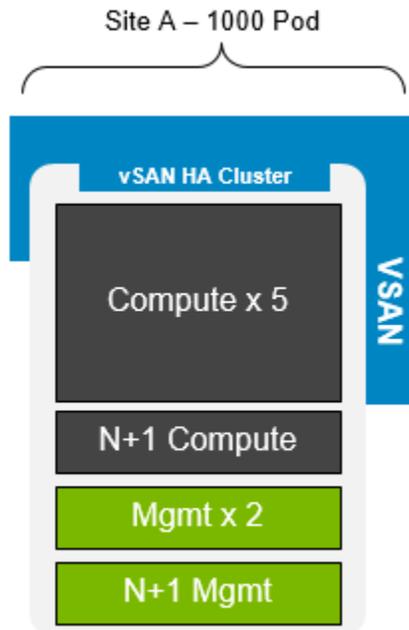


The distributed switch configuration is the same on all vSAN storage hosts. It is recommended to have at least two uplinks for each host to provide load balancing and fail back redundancy.



5.4 Scaling Guidance

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 in this case 64 nodes in total per cluster. The supported amount of VMs per SmartNode at the time of release is 1000 so taking this into consideration we need 5 compute nodes per block for the C7 Task Worker user @200 per Host. The image below shows a 1000 user HCOP vSAN block.



- 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.

Currently, only one SmartNode (on premise cluster) is supported, so the HCOP appliances are not scalable outside of this. The below table shows the scalability options for each component.

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

5.5 Solution High Availability

High availability (HA) for vSAN is offered to protect each layers of the solution architecture, individually if desired. Following the N+1 model, additional ToR switches for LAN or vSAN, are added to the Network layer and stacked to provide redundancy as required, additional hosts are added to their respective layers. All compute and management nodes should be part of the same vSphere HA cluster.

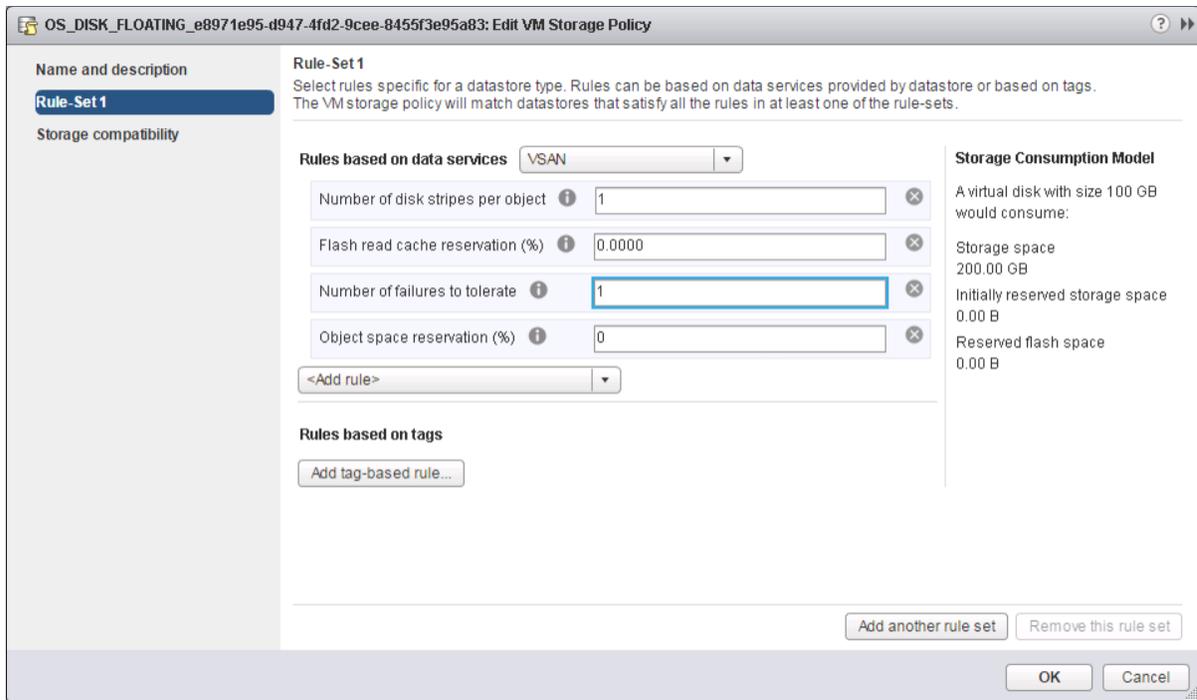
The HA options provide redundancy for all critical components in the stack while improving the performance and efficiency of the solution as a whole.

- Additional switches added to the existing thereby equally spreading each host's network connections across multiple switches.
- Additional ESXi hosts added in the VSRN Cluster to provide N+1 protection.

5.5.1 vSAN HA/ FTT Configuration

The minimum configuration required for vSAN is 3 ESXi hosts. A consideration of the 3-Node cluster model is if one node fails there is nowhere to rebuild the failed components, so 3 node clusters should be used only for POC or non-production.

The virtual machines that are deployed via HCOP are policy driven and one of these policy settings is Number of failures to tolerate (FTT). The default value is FTT=1 which will create a mirrored copy of the Virtual Machine's objects. For example, if the VM is 40Gb in size then 80Gb of storage space is required.



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 and leverage shared vSAN storage so can make full use of vSphere HA and the vSRN Compute hosts can be configured in an HA cluster following the boundaries of vCenter with respect to limits imposed by VMware. This will result in multiple HA clusters managed by multiple vCenter servers. HCOP supports 1000 VMs per SmartNode.

5.5.3 HCOP High Availability

The HCOP appliances are deployed in pairs so for High Availability it is recommended to keep each appliance pair on separate hosts within the vSRN Cluster. This can be configured via a VM Affinity Rule, for more information on how to configure Virtual Machine Affinity Rules [Link](#)

6 Solution Performance and Testing

6.1 Summary

At the time of publication, here are the available density recommendations as tested. The below user densities were achieved by following the VMware/HCOP best practices. The testing was completed on the VSRN C7 configuration but the user densities are the same for VxRail V470/V470F-C7 since this configuration has the same memory and CPU configuration. Please refer to [Section 3](#)

Table 1 User density summary

Hypervisor	Provisioning	Profile	Template OS	Config	User Density
6.0 Update 2	Instant Clones	Task	Windows 10	VSRN/VxRail -C7	200
6.0 Update 2	Instant Clones	Power	Windows 10	VSRN/VxRail -C7	109
6.0 Update 2	Instant Clones	Task	Windows 10	VSRN/VxRail -C7	200
6.0 Update 2	Instant Clones	Power	Windows 10	VSRN/VxRail -C7	109

For detailed validation results and analysis of these reference designs and more see below.

6.2 Test and performance analysis methodology

6.2.1 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.

6.2.1.1 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.

6.2.1.2 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.

6.2.2 Resource monitoring

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

6.2.2.1 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.

6.2.2.2 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.

6.2.2.3 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. The threshold details are in the bellow table

Resource utilization thresholds

Parameter	Pass/Fail Threshold
Physical Host CPU Utilization (AHV & ESXi hypervisors)*	100%
Physical Host CPU Utilization (Hyper-V)	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.

6.3 Test Configuration

The following components were used to complete the validation testing for the solution: VxRail currently only supports ESXi 6.0 U2.

Table 2 Hardware and software test components

Component	Description/Version
Hardware platform(s)	VSRN R730 AF/HY & VxRail V470/V470F
Hypervisor(s)	ESXi 6.0 Update 2/ ESXi 6.5
Broker technology	HCOP
Broker database	HCOP
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

6.3.1 Compute VM configurations

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

Table 3 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

Table 4

6.3.2 Platform Configurations

The hardware configuration details are listed in [Section 3](#)

6.4 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.

Table 5 Test result summary

Platform Config	Hypervisor	Broker & Provisioning	Client OS	Login VSI Workload	Density Per Host	Avg CPU	Avg Mem Consumed	Avg Mem Active	Avg IOPS / User
C7	ESXi 6.0 U2	HCOP	Win 7	Task Worker	200	67%	370GB	223GB	12.4
C7	ESXi 6.0 U2	HCOP	Win 7	Power Worker	109	68%	370GB	249GB	19
C7	ESXi 6.0 U2	HCOP	Win 10	Task Worker	200	80%	369GB	256GB	15
C7	ESXi 6.0 U2	HCOP	Win 10	Power Worker	109	78%	369GB	261GB	18.1
C7*	ESXi 6.0 U2	HCOP	Win 10	Task Worker	200	66%	370GB	210GB	14
C7*	ESXi 6.0 U2	HCOP	Win 10	Power Worker	109	68%	370GB	276GB	15.2

*This test was performed with Office 2016 embedded in the Client OS and not via an Appstack.

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 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: ESXi consumed memory is the amount of host physical memory consumed by a host.

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.

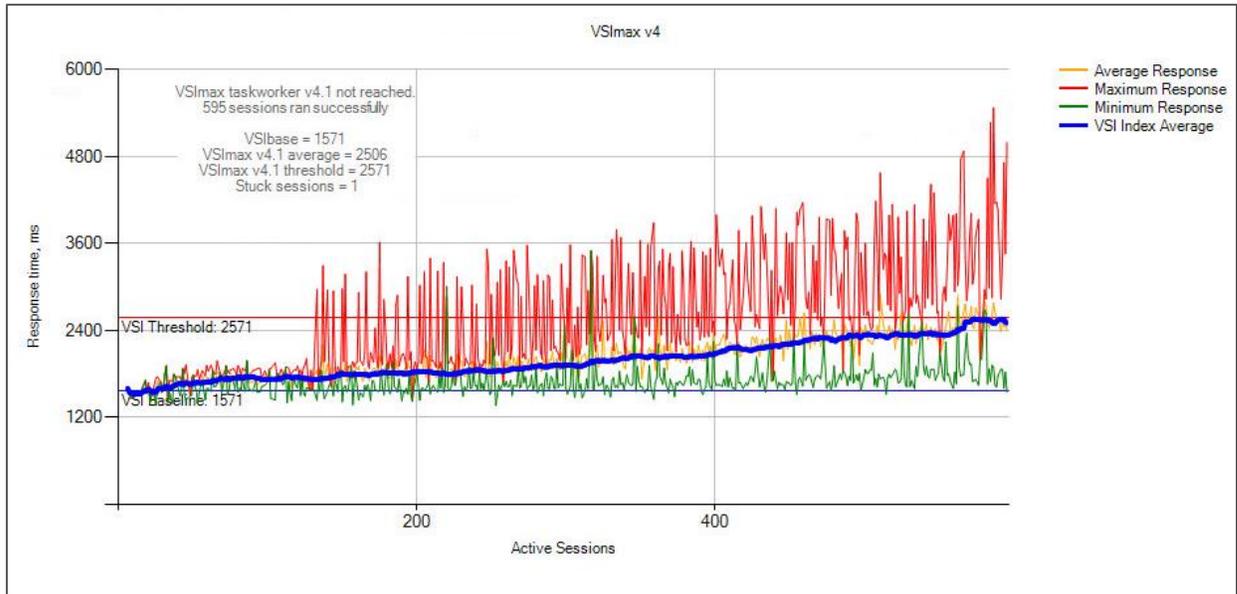
Avg Net Mbps/User: Amount of network usage over the steady state period divided by the number of users. For clusters, this is the combined average of all compute hosts over the steady state period divided by the number of users on a host.

6.4.1 HCOP VSRN/ VxRail C7

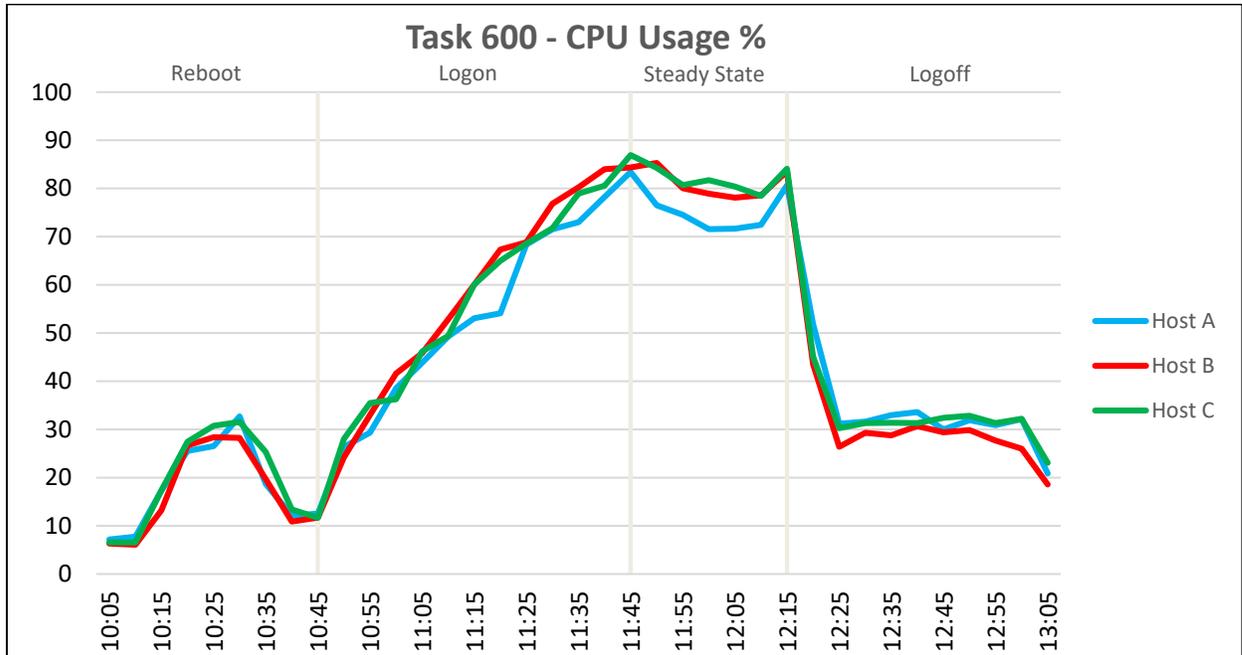
The hardware configuration details are listed in [Section 3](#)

6.4.1.1 Task Worker, 600 user, HCOP instant clones, Windows 10, Office 2016

VSI Max was not reached on this test run indicating no deterioration in user experience.



The CPU usage on this test run reached a steady state average of 69%. This is a bit short of the 85% capacity but with the memory usage close to maximum on the config the number of users is appropriate. The CPU would be capable of handling extra users with an increase in the amount of physical memory.

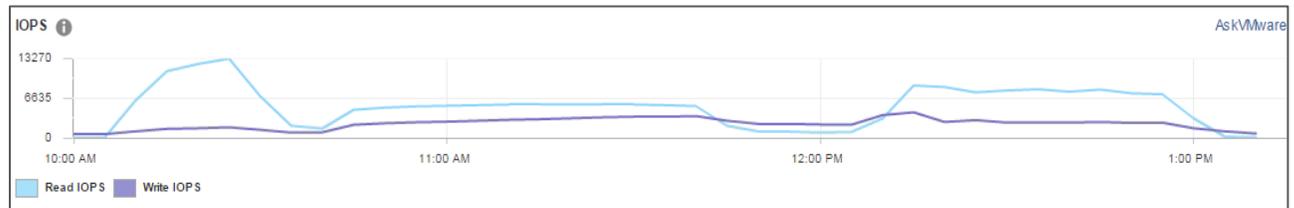


Latency on the datastore showed no significant spikes on this test run. It reached approximately 6ms read latency at the beginning of steady state, well below the 20ms threshold. Write latency was not an issue.

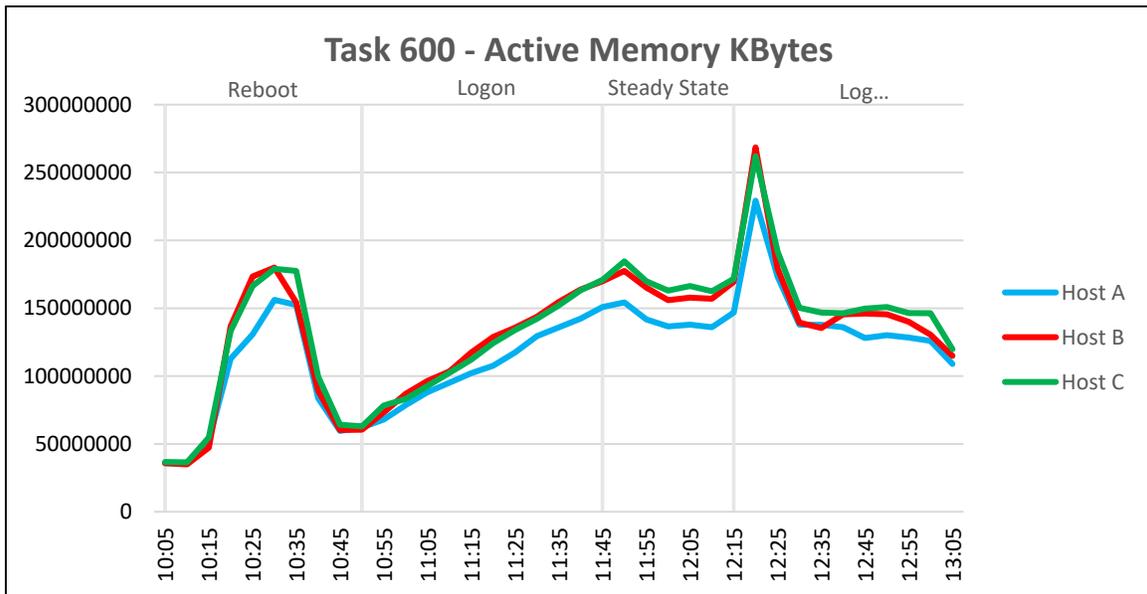
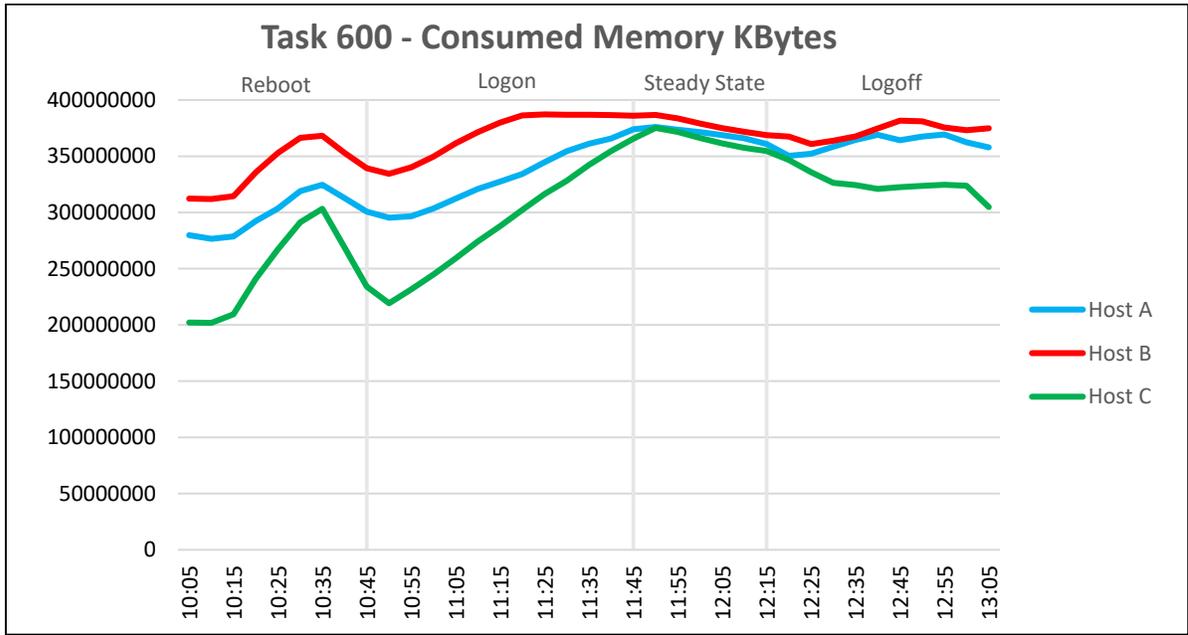
The IOPS peaked during the reboot phase before starting test and during the rebalancing of the pool after the users had logged off the desktops but was steady during the login and steady state phases and showed no spikes at any point.

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.

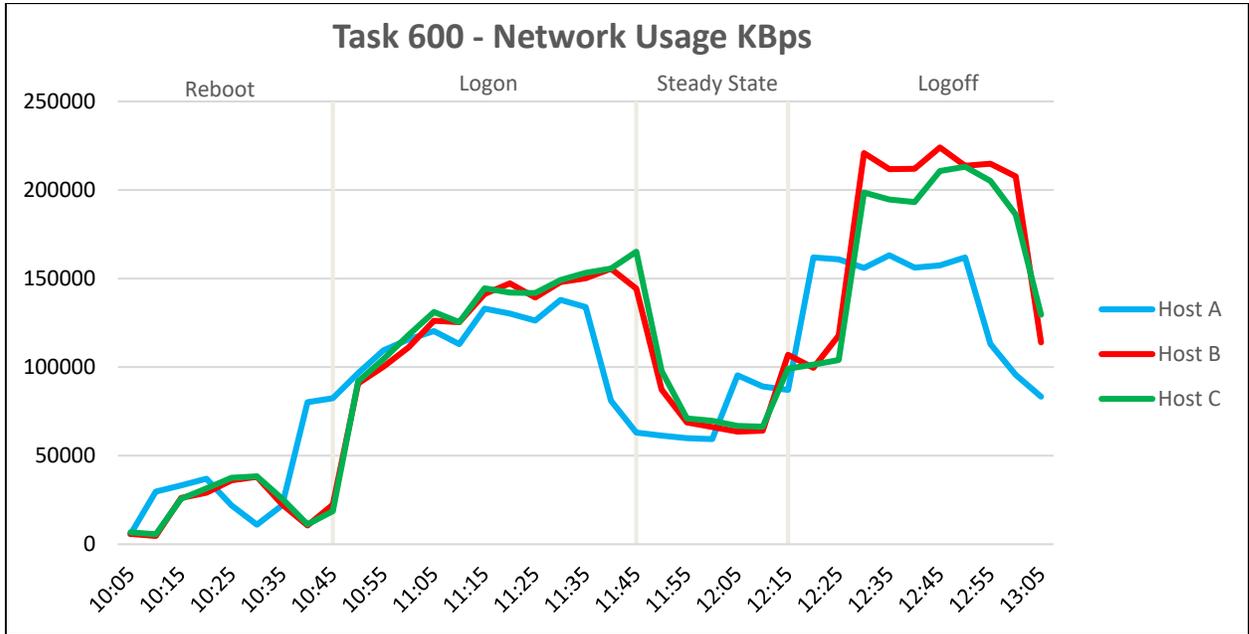
The statistics shown are for three hosts.



With 384 GB of memory installed in the host servers, memory usage came close to topping out, with 369 GB of memory consumed. A small amount of memory ballooning took place on Host B at the start of steady state for a few minutes but there was no other ballooning or swapping on the other hosts. This suggests the number of users tested on this configuration is close to the maximum number possible.

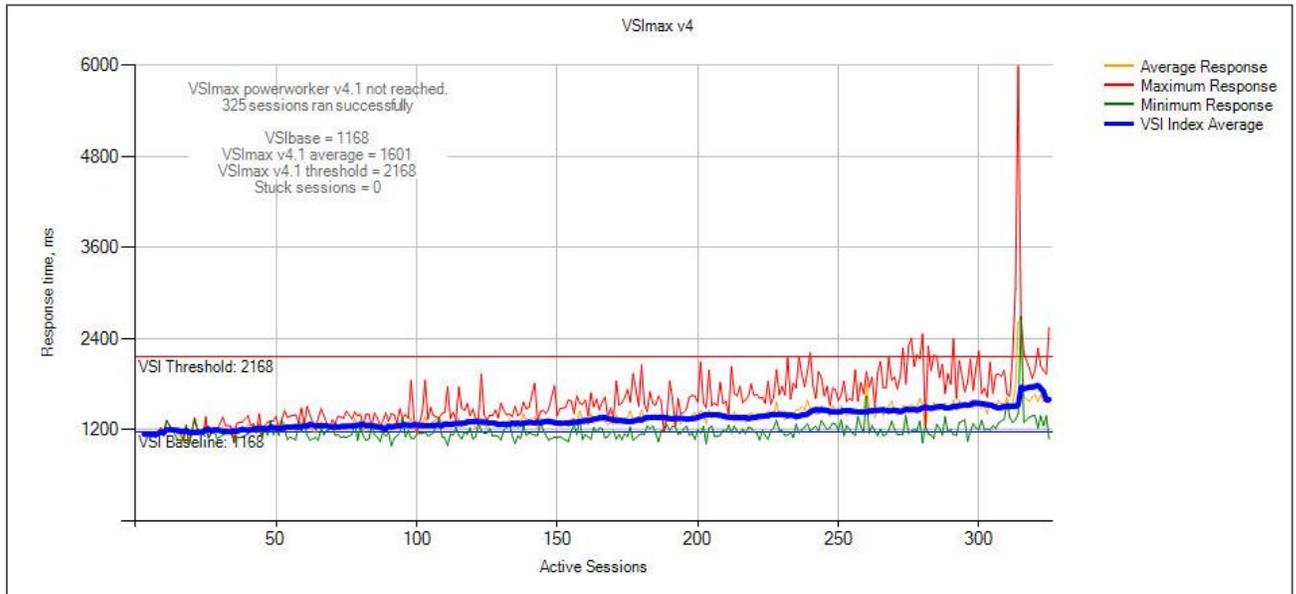


Network utilization was not an issue in this test run. Host usage reaching a maximum of approximately 224,000 KBps and this was during the logoff phase.

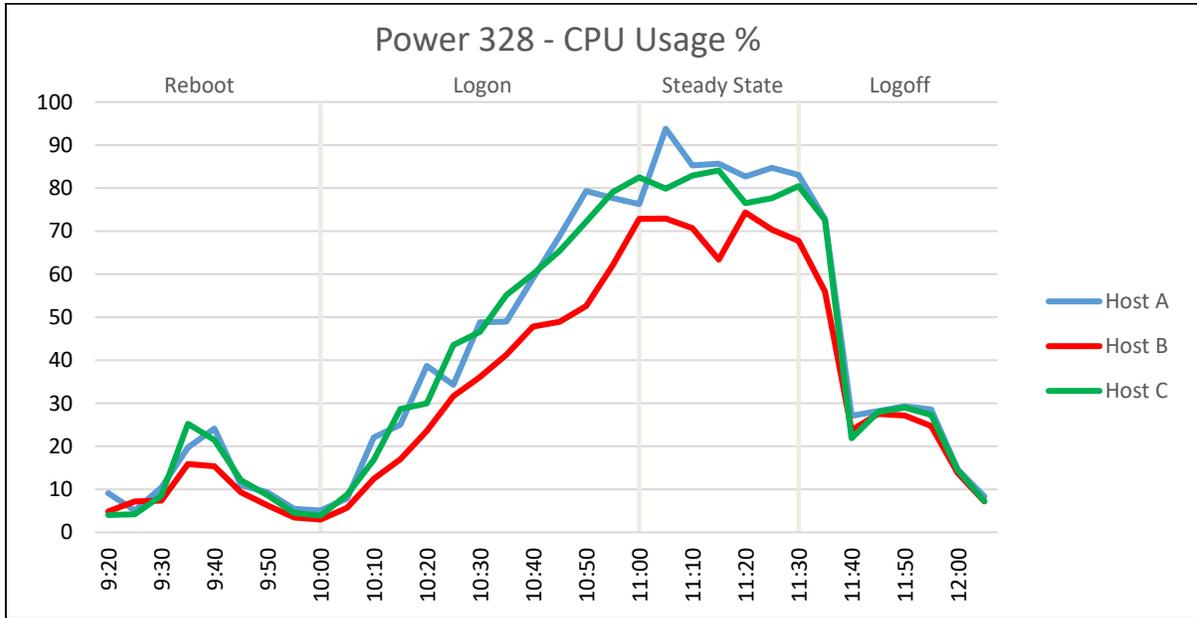


6.4.1.2 Power Worker, 328 user, HCOP instant clones, Windows 10, Office 2016

VSI Max was not reached on this test run indicating no deterioration in user experience.



The CPU usage on this test run reached a steady state average of 84%. The CPU would be capable of handling extra users with an increase in the amount of physical memory. Host B only reached a steady state average of 70% but Host B did not host any HAH management VM's while both other servers hosted management VM's.

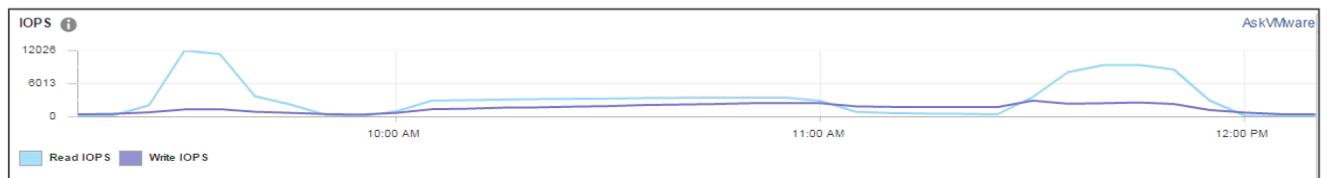


Latency on the datastore showed no significant spikes on this test run. It reached approximately 5ms read latency during the reboot of the virtual machine pool before starting testing, well below the 20ms threshold. Write latency was not an issue.

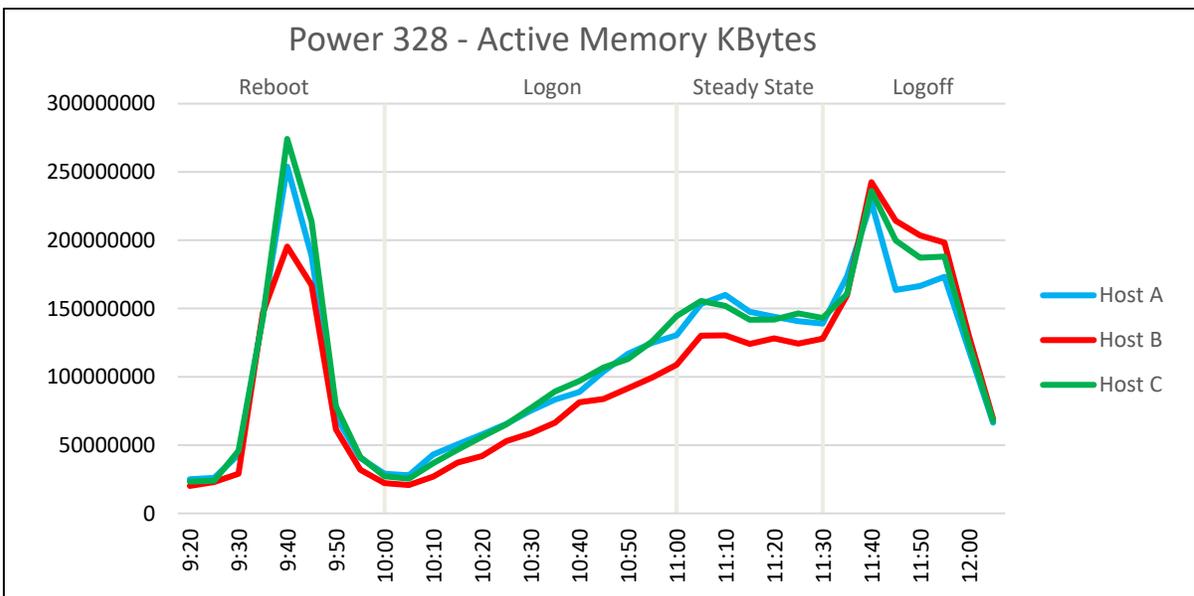
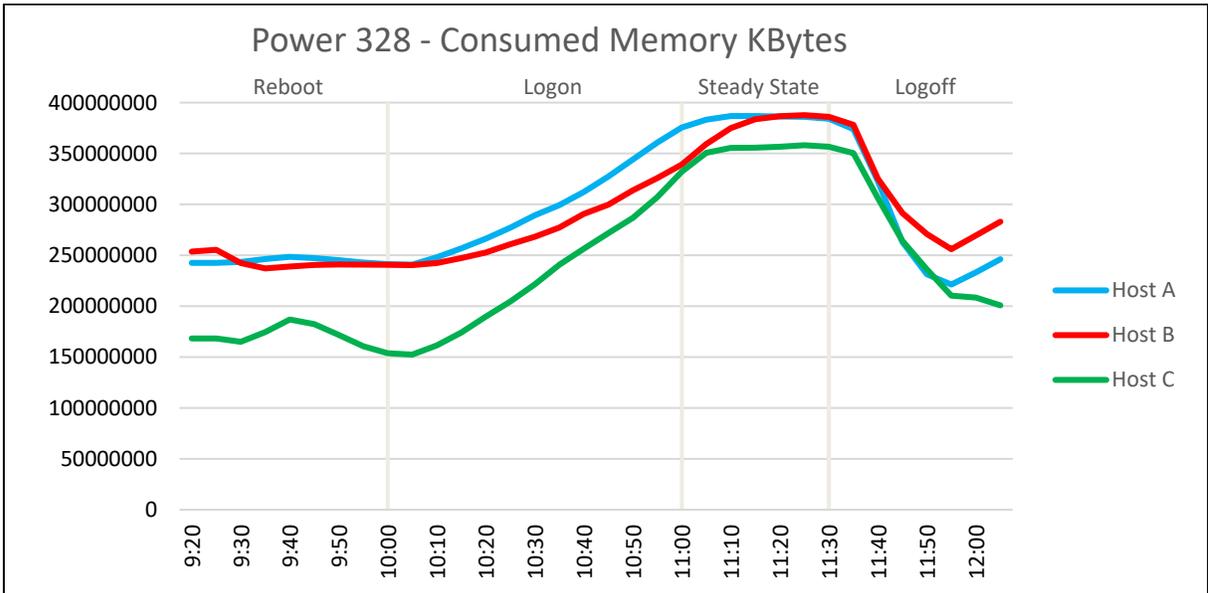
The IOPS peaked during the reboot phase before starting test and during the rebalancing of the pool after the users had logged off the desktops but was steady during the login and steady state phases and showed no spikes at any point.

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.

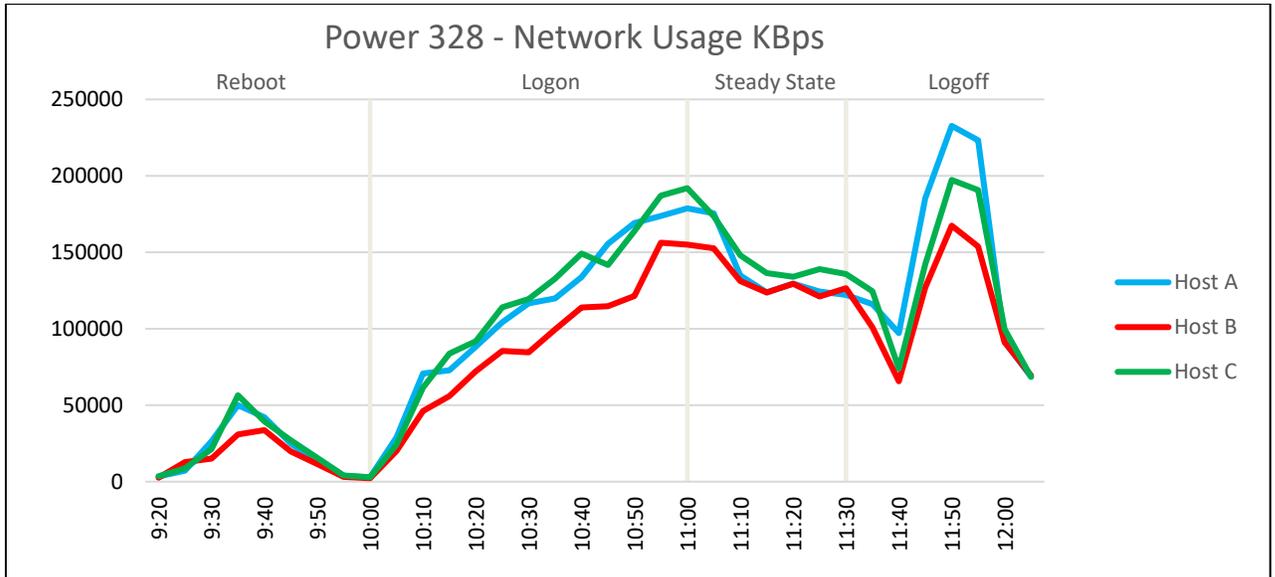
The statistics shown are for three hosts.



With 384 GB of memory installed in the host servers, memory usage came close to topping out, with 369 GB of memory consumed. There was a small amount of memory ballooning at the beginning of steady state for a few minutes and small amount of swapping took place on Host C. This suggests the number of users tested on this configuration is close to the maximum number possible.



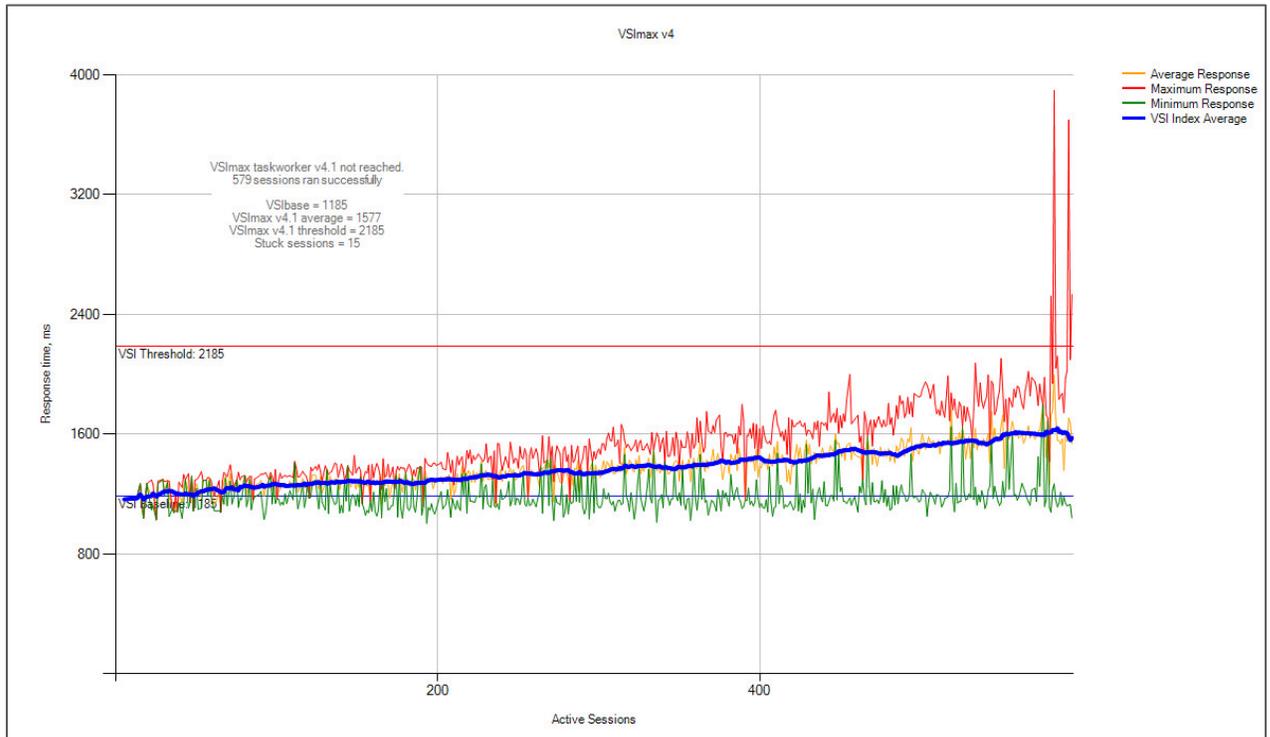
Network utilization was not an issue in this test run. Host usage reaching a maximum of approximately 232,000 KBps and this was during the logoff phase.



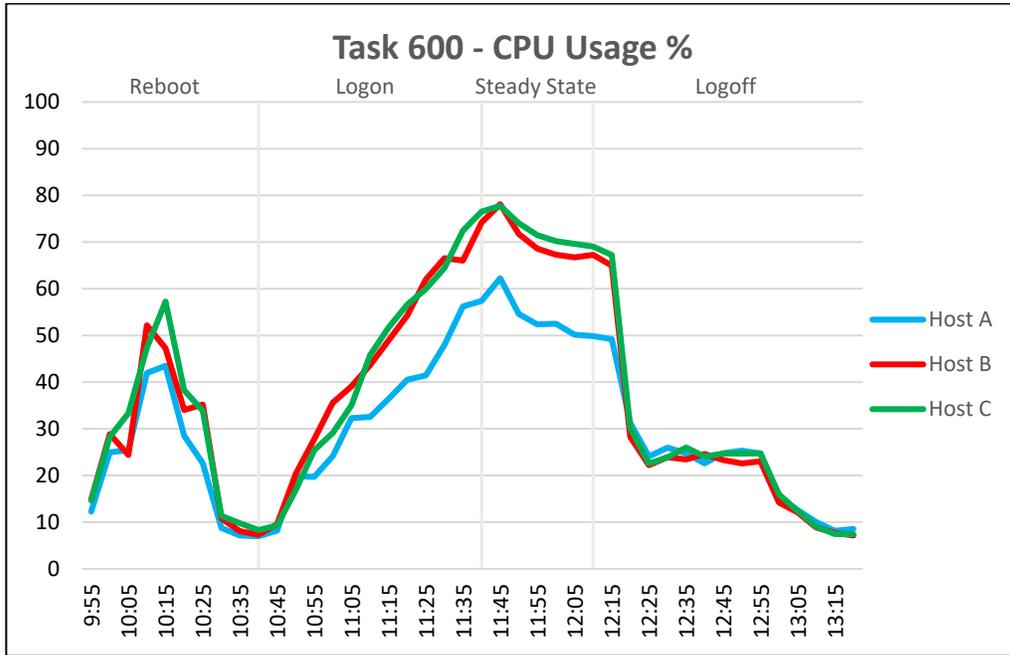
6.4.1.3 Task Worker, 600 user, HCOP instant clones, Windows 10, Office 2016*

*Note that this test run was performed with Office 2016 installed in the master virtual machine image rather than providing the office suite through AppVolumes as indicated in section 1.4.2

VSI Max was not reached on this test run indicating no deterioration in user experience.



The CPU usage on this test run reached a steady state average of 66%. The CPU would be capable of handling extra users with an increase in the amount of physical memory. Lower CPU usage is commonly seen on test runs where the test virtual machines have only a single vCPU.

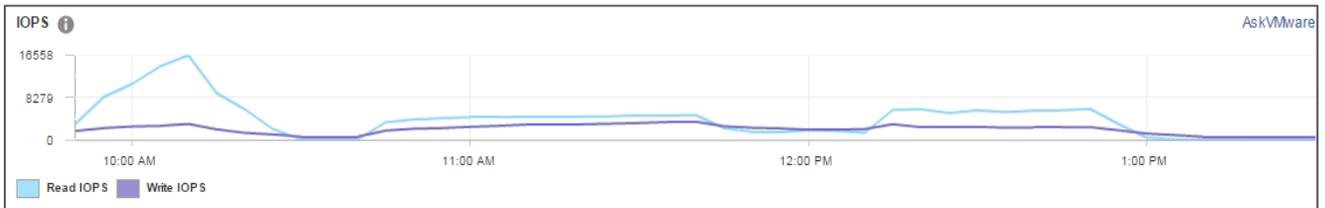


Latency on the datastore showed no significant spikes on this test run. It reached approximately 3.5ms read latency close to the end of the steady state period. Write latency was not an issue.

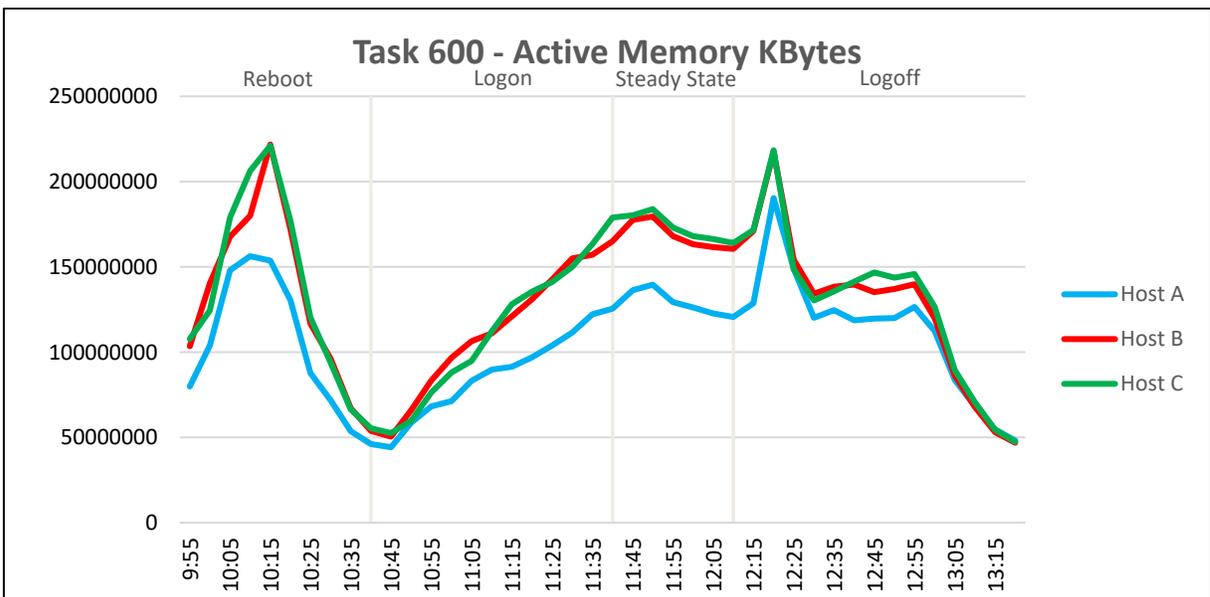
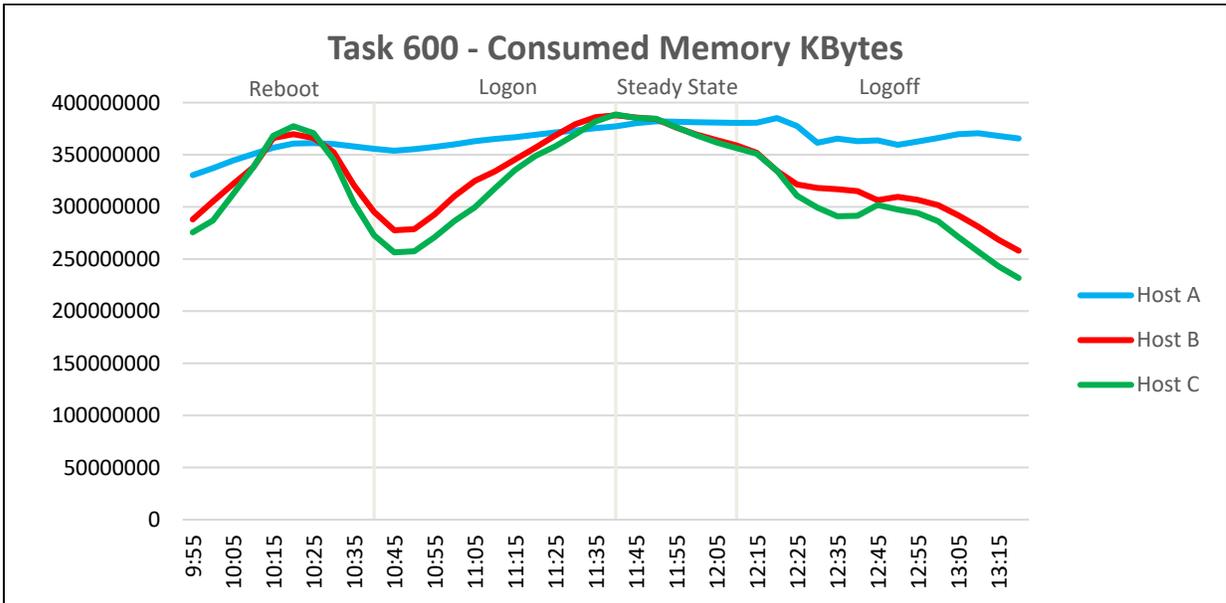
The IOPS peaked during the reboot phase before starting test and during the rebalancing of the pool after the users had logged off the desktops but was steady during the login and steady state phases and showed no spikes at any point.

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.

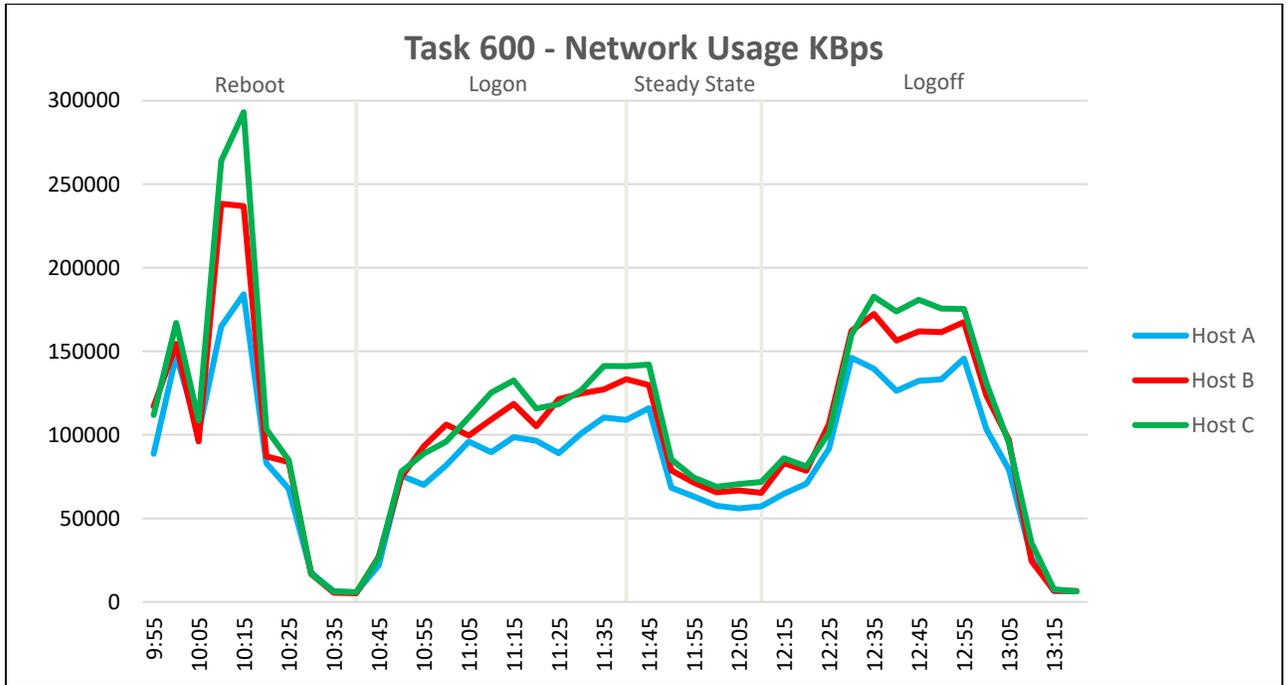
The statistics shown are for three hosts.



With 384 GB of memory installed in the host servers, memory usage came close to topping out, with 370 GB of memory consumed and a small amount of swapping took place on Host C only. There was no memory ballooning on any host. This suggests the number of users tested on this configuration is close to the maximum number possible.



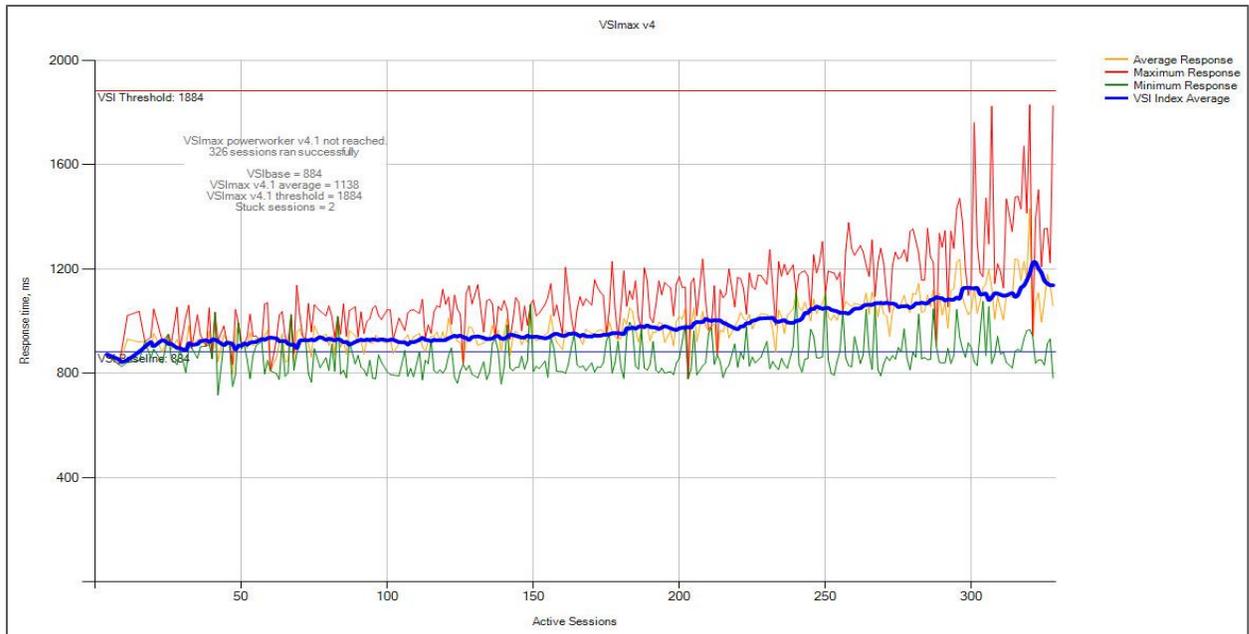
Network utilization was not an issue in this test run. Host usage reaching a maximum of approximately 293,000 KBps reboot of the desktop pool.



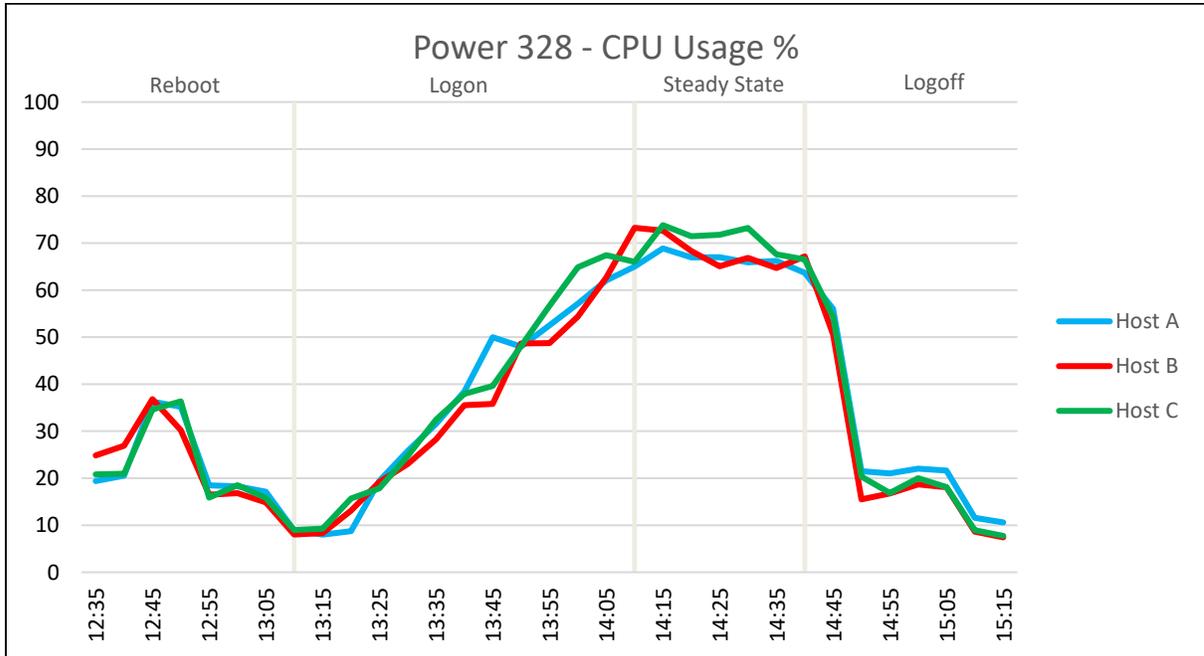
6.4.1.4 Power Worker, 328 user, HCOP instant clones, Windows 10, Office 2016*

*Note that this test run was performed with Office 2016 installed in the master virtual machine image rather than providing the office suite through AppVolumes as indicated in section 1.4.2

VSI Max was not reached on this test run indicating no deterioration in user experience.



The CPU usage on this test run reached a steady state average of 68%. The CPU would be capable of handling extra users with an increase in the amount of physical memory. Host B did not host any HAH management VM's while both other servers hosted management VM's but this did not have an impact on performance on this case.

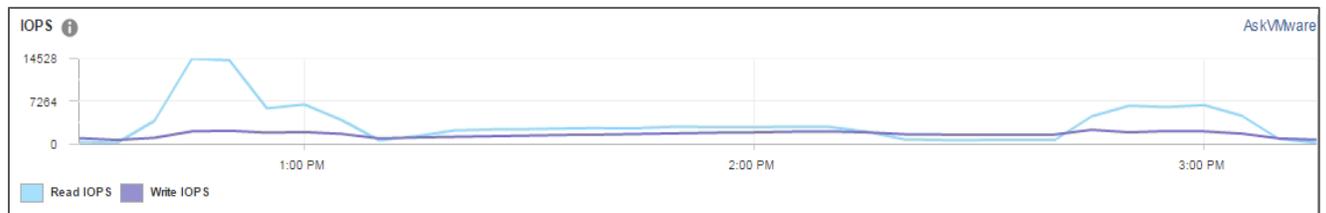


Latency on the datastore showed no significant spikes on this test run. It reached approximately 1.9 ms read latency close to the end of the steady state period. Write latency was not an issue.

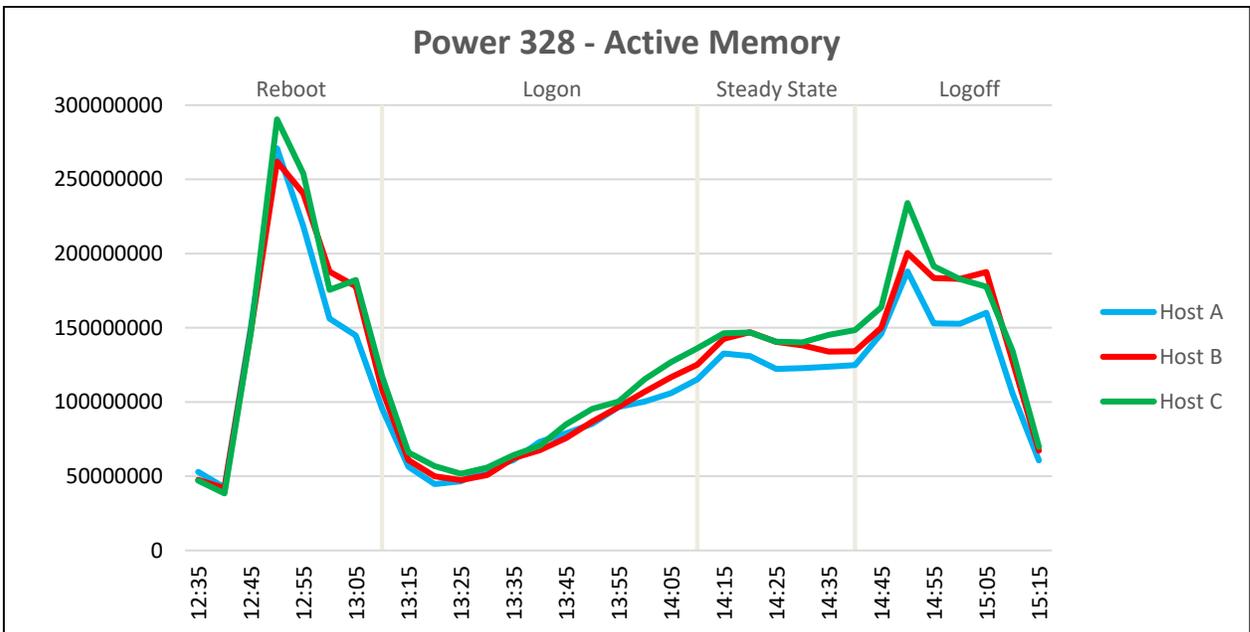
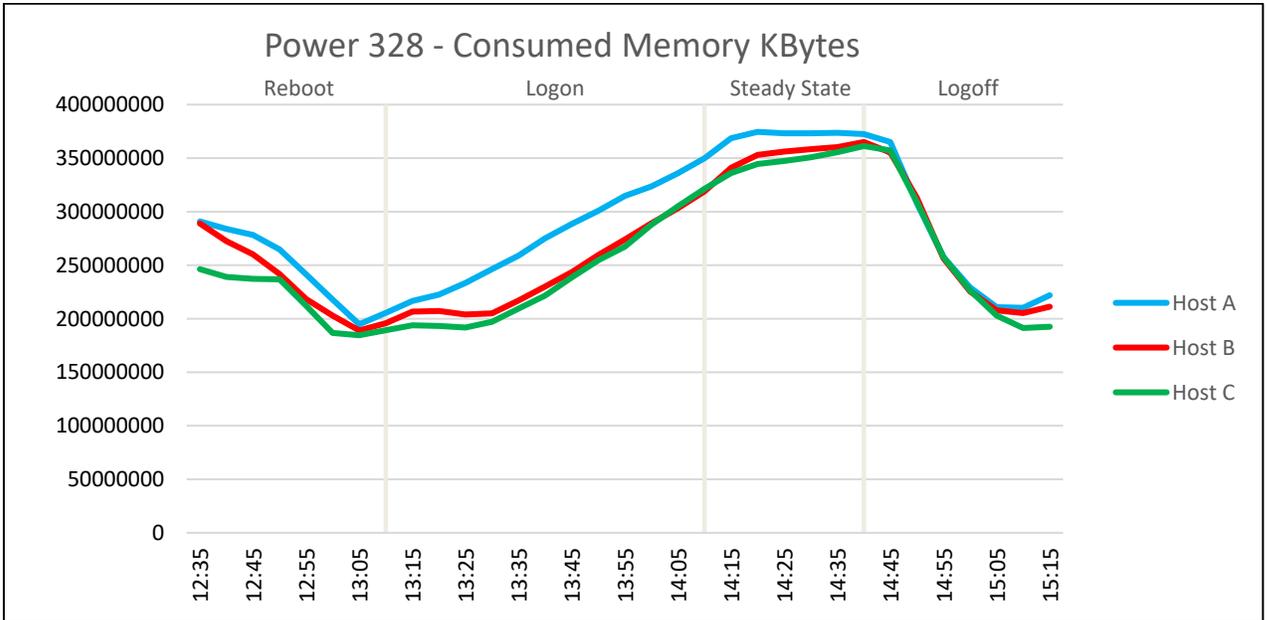
The IOPS peaked during the reboot phase before starting test and during the rebalancing of the pool after the users had logged off the desktops but was steady during the login and steady state phases and showed no spikes at any point.

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.

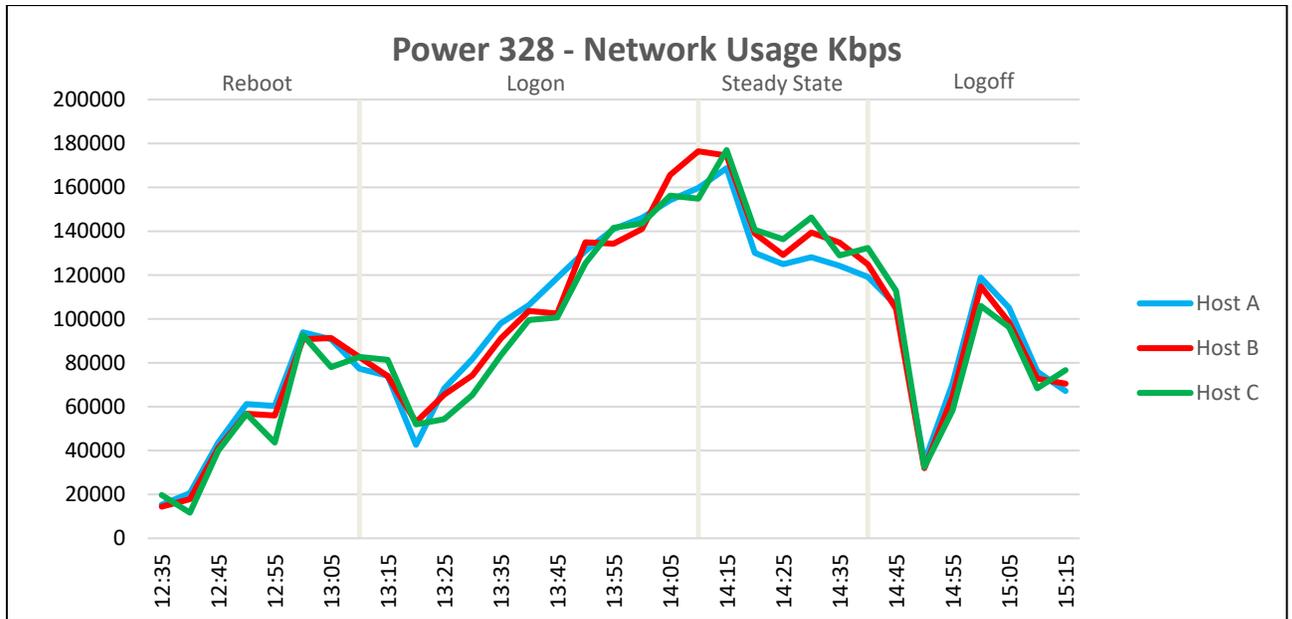
The statistics shown are for three hosts.



With 384 GB of memory installed in the host servers, memory usage came close to topping out, with 357 GB of memory consumed and a small amount of swapping took place on Host C only. There was no memory ballooning on any host. This suggests the number of users tested on this configuration is close to the maximum number possible.



Network utilization was not an issue in this test run. Host usage reaching a maximum of approximately 177,000 KBps at the beginning of the steady state period.



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 and Yinglong Jiang for their help in the validation of the solutions in this RA.

Andrew 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.

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 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 a Lead Architect 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.