Wyse Datacenter for VMware Horizon Reference Architecture

A Reference Architecture for the design, configuration and implementation of a VMware Horizon environment.

Dell Engineering
September 2017

A Dell Reference Architecture

Revisions

Date	Description
May 2014	Initial release (v.6.5)
July 2015	Updated density numbers for ESXi 6.0 and added PowerEdge C4130 (v.6.7)
April 2016	Document overhaul, new Broadwell MLK, networking, servers. (v.7.0)
August 2016	Updated list of Thin Clients, added NSX and M60 GPU (v.7.1)
September 2016	Update list of Thin Clients.
January 2017	Replaced Compellent and EqualLogic recommendations with Dell EMC storage
April 2017	Section 7 Test results update
September 2017	Horizon Apps section added and Thin Clients section updated

Contents

1 Introduction

1.1 Purpose

This document addresses the architecture design, configuration and implementation considerations for the key components of the architecture required to deliver virtual desktops via VMware Horizon on VMware vSphere 6.

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 that 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.

1.3 What's New

- Introduce support for Intel Broadwell processors (E5-2600v4)
- Introduce support for new Dell Networking options
- Introduce all-flash for Local Tier 1 rack solution
- Replaced Compellent and EqualLogic storage recommendations with Dell EMC storage

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)
- Full Clone Virtual Desktops (Persistent)
- Shared Sessions (RDSH)

Additionally, our solution includes options for users who require:

- Graphics Acceleration
- Unified Communications

2.2 Physical Architecture Overview

The core Dell Wyse Datacenter architecture consists of two models: Local Tier1 and Shared Tier1. “Tier 1” in the Dell Wyse Datacenter context defines from which disk source the VDI sessions execute. Local Tier1 includes rack servers or blades with SSDs while Shared Tier 1 can include rack or blade servers due to the usage of shared Tier 1 storage. Tier 2 storage is present in both solution architectures and, while having a reduced performance requirement, is utilized for user data and Management VM execution. Management VM execution occurs using Tier 2 storage for all solution models. Dell Wyse Datacenter is a 100% virtualized solution architecture.

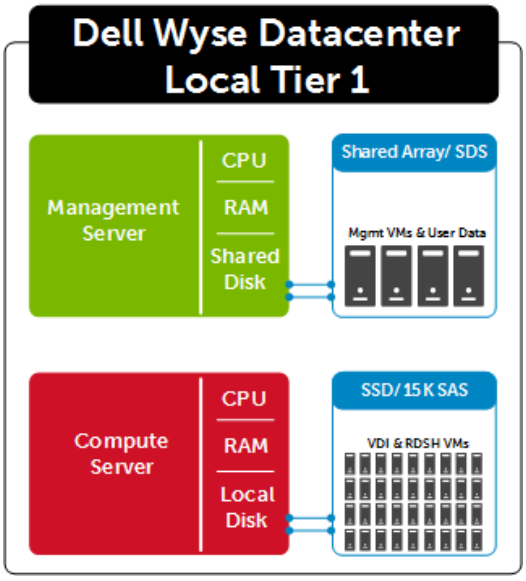

In the Shared Tier 1 solution model, an additional high-performance shared storage array is added to handle the execution of the VDI sessions. All compute and management layer hosts in this model are diskless.

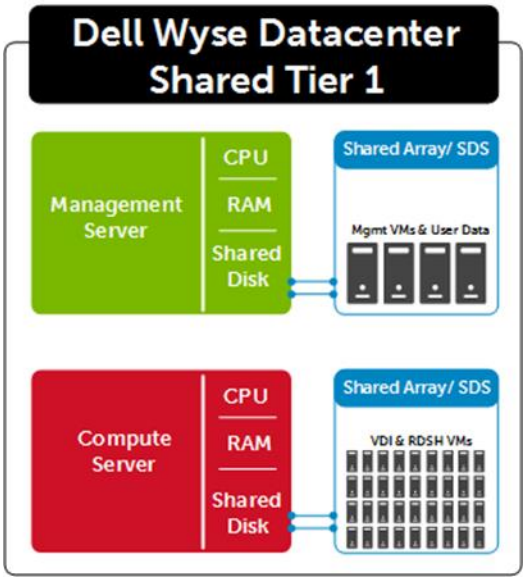

2.3 Solution Layers

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

- Networking Layer
- Compute Server Layer

- Management Server Layer
- Storage Layer
- Thin Client Layer (please refer to section 3.6)

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

2.3.1 Networking

Only a single high performance Dell Networking 48-port switch is required to get started in the network layer for a combined pilot/POC configuration. For all other configurations, you can start with a single Dell Networking 48-port switch for 10 GB LAN traffic along with a single Brocade fibre channel switch for SAN connectivity. Additional switches are added and stacked as required to provide High Availability for the Network layer.

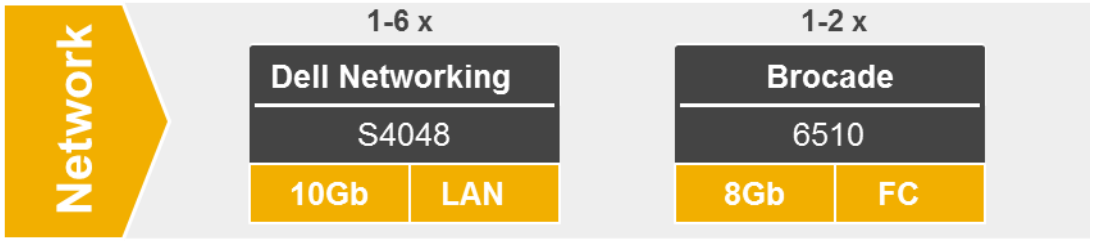

2.3.2 Compute

The compute layer consists of the server resources responsible for hosting the Horizon user sessions, hosted via the VMware vSphere hypervisor, local or shared tier 1 solution models (local Tier 1, all-flash, pictured below).

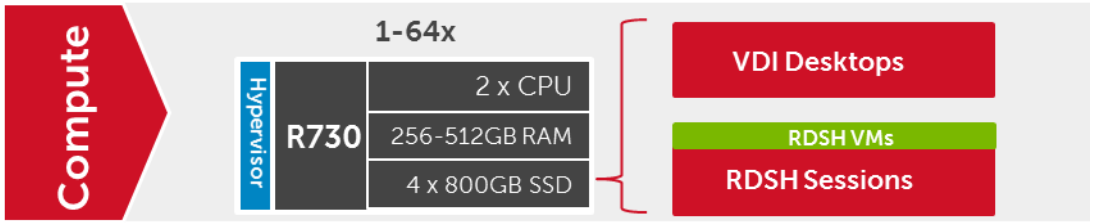

2.3.3 Management

VDI management components are dedicated to their own layer so that they do not negatively affect the user sessions running in the compute layer. This physical separation of resources provides clean, linear, and predictable scaling without the need to reconfigure or move resources within the solution as you grow. The management layer will host the entire server VMs necessary to support the VDI infrastructure.

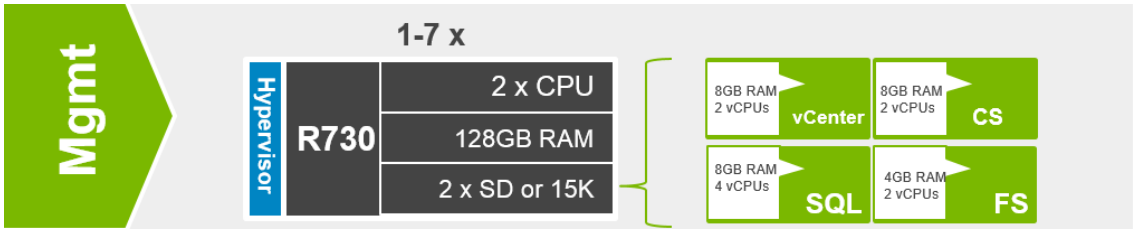

2.3.4 Storage

The storage layer consists of options provided by Dell EMC XtremIO arrays for combined shared T1, T2, and file storage (optionally for up to 500 users only) or XtremIO arrays for capacity and scaling in discrete shared T1 configurations. Dell EMC Unity arrays are used for discrete Tier 2 management VM storage and user data file storage. The typical configurations of XtremIO X-Brick arrays for shared T1 and Unity arrays for T2 are depicted below.

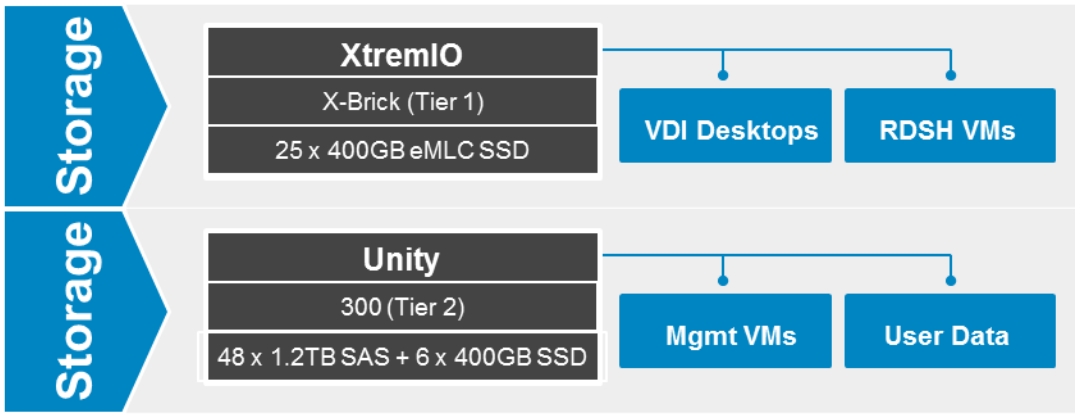

2.4 Local Tier 1

2.4.1 Base Offerings

2.4.1.1 Local Tier 1 – Combined Pilot/POC (Up to 150 Users)

For pilot or small deployments, a single server can be used. This architecture is non-distributed with all VDI, Management, and storage functions on a single host. If additional scaling is desired, you can grow into a larger distributed architecture seamlessly.

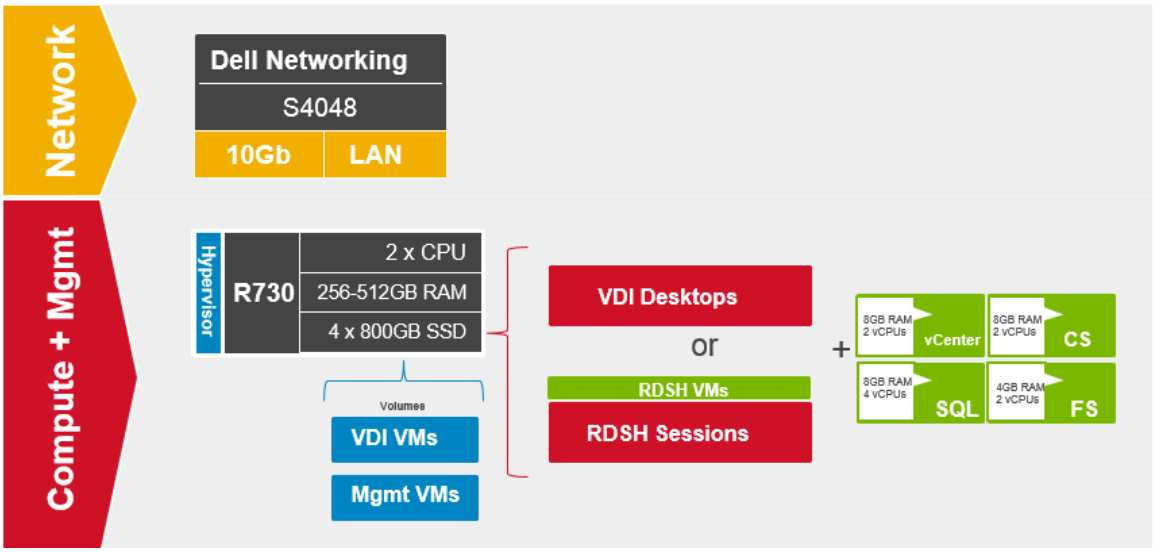

Note: 150 user density is based on the Task Worker workload

2.4.2 Local Tier 1 for Rack Servers

The Local Tier 1 solution model provides a scalable rack-based configuration that hosts user VDI sessions on local SSD or spinning disk in the compute layer, All-flash pictured below, if spinning disk is desired substitute the SSDs with 10-12 x 600GB 15K SAS HDDs.

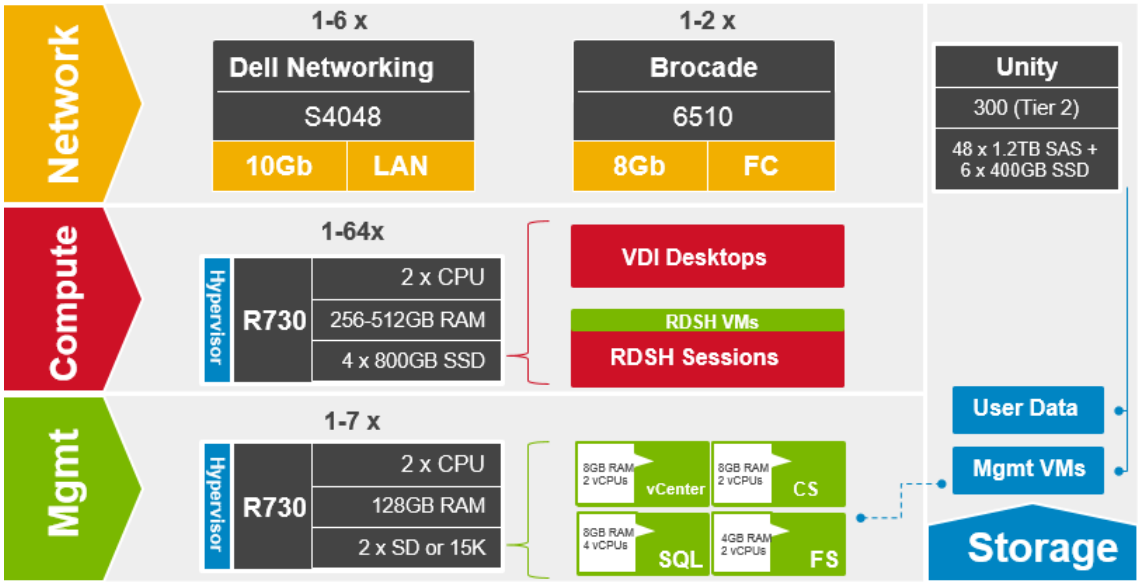

2.4.2.1 Local Tier 1 – Network Architecture

In the local Tier 1 architecture, the network fabrics are separated to isolate LAN and FC storage traffic. Dell Engineering recommends making each switch stack redundant. Only the management servers connect to FC

storage in this model. All Top of Rack (ToR) traffic is layer 2 (L2)/ switched locally, with all layer 3 (L3)/ routable VLANs trunked from a core or distribution switch. The following diagrams illustrate the logical data flow in relation to the core switch.

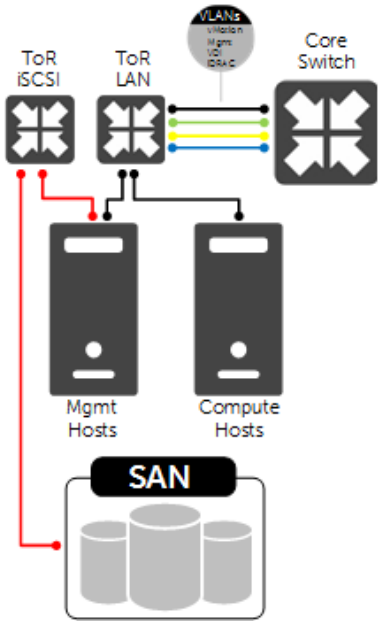

2.4.2.2 Local Tier 1 – Cabling (HA)

The following diagram depicts the LT1 rack solution including optional components for HA:

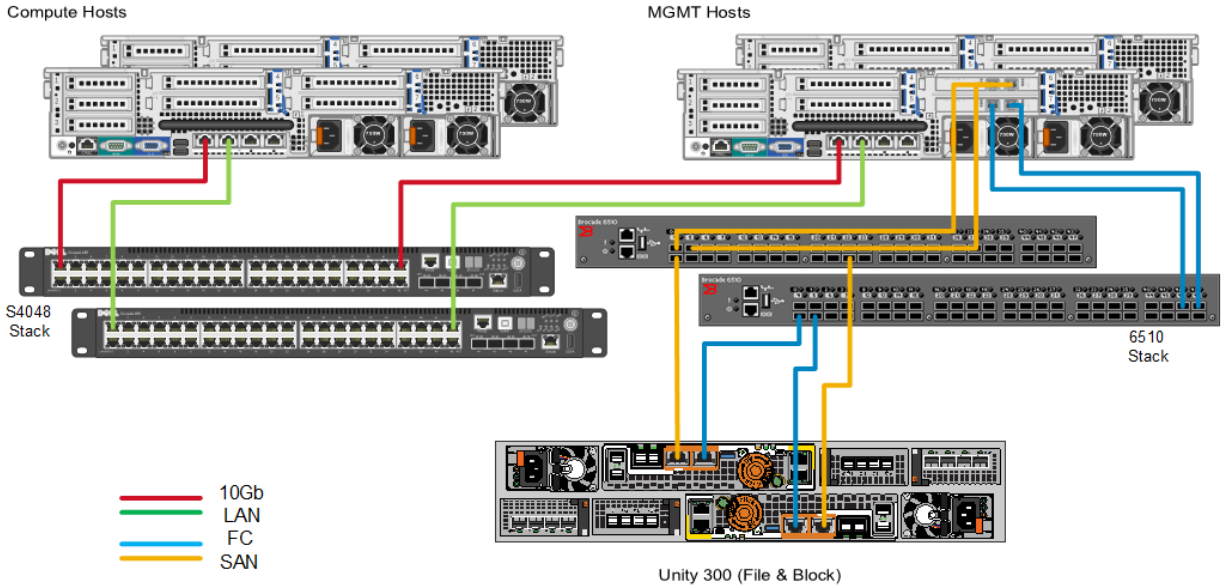

Refer to the Tier 2 – Unity 300 section for a diagram of disk enclosure connections via mini-SAS HD

2.4.2.3 Local Tier 1 – Rack Scaling Guidance

Local Tier 1 HW Scaling – Rack Servers (FC)			
User Scale	NAS & T2	ToR LAN	ToR FC
0-3,000	1 x Unity 300	S4048	6510
3,001-6,000	2 x Unity 300		
6,001-9,000	3 x Unity 300		
9,001-10,000	4 x Unity 300		

Note: For deployments over 10,000 users, create additional pods using sizing guidance contained herein.

2.4.3 Local Tier 1 for Blade Servers

The Local Tier 1 solution model for blade servers provides a high-performance 800GB SSD configuration that does not require shared storage but Tier 2 is added to provide HA to the management layer infrastructure. User VDI sessions are hosted locally on SSDs in each blade server using Horizon Connection Server (HCS) for desktop delivery. A pair of PowerEdge M I/O Aggregator switches is required in the A Fabric. To support the B Fabric as shown, blade chassis interconnects must be added along with FC HBAs in the servers. ToR FC switching is optional if a suitable FC infrastructure is already in place.

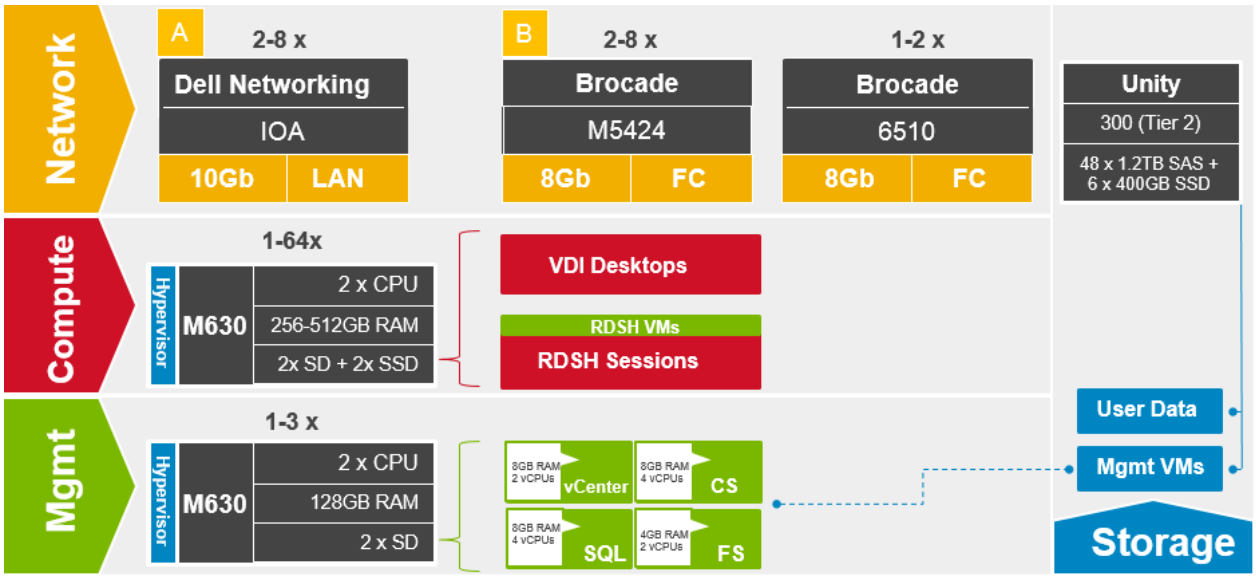

2.4.3.1 Local Tier 1 – Network Architecture

In the Local Tier 1 architecture for blades, a ToR switch is not required; however, a separate switching infrastructure is required for FC. The A Fabric IOA switches can connect directly to the core or distribution

network layer. Management servers connect to shared storage using FC switched via chassis interconnects. Both Management and Compute servers connect to all VLANs in this model. All ToR traffic has designed to be layer 2/ switched locally, with all layer 3/ routable VLANs routed through a core or distribution switch. The following diagrams illustrate the server NIC to ToR switch connections, vSwitch assignments, as well as logical VLAN flow in relation to the core switch.

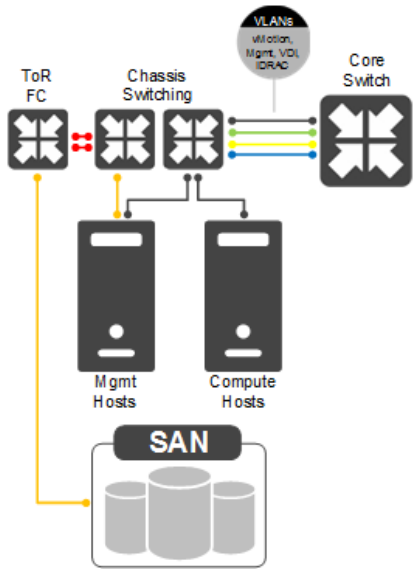

2.4.3.2 Local Tier 1 – Cabling (HA)

The following diagram depicts the Local Tier 1 blade solution including optional components for HA. The A Fabric, B Fabric and ToR switches are stacked, respectively.

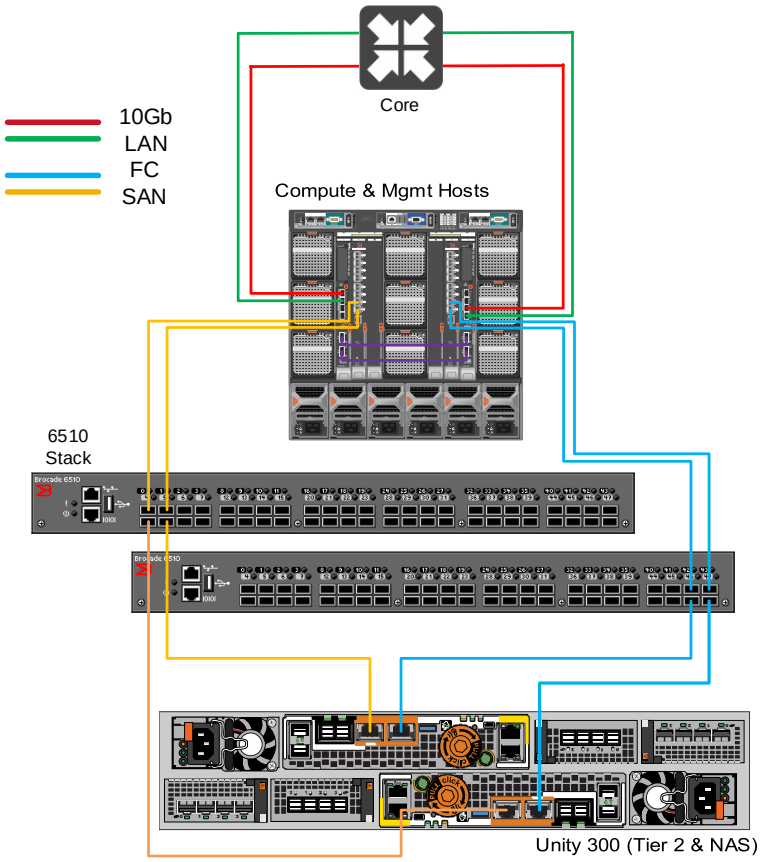

Refer to the Tier 2 – Unity 300 section for a diagram of disk enclosure connections via mini-SAS HD.

2.4.3.3 Local Tier 1 – Blade Scaling Guidance

Local Tier 1 HW Scaling – Blade Servers (FC)				
User Scale	NAS & T2	Blade LAN (A Fabric)	Blade FC (B Fabric)	ToR 8Gb FC
0 – 3,000	1 x Unity 300	IOA	M5424	6510
3,001 – 6,000	2 x Unity 300			
6,001 – 9,000	3 x Unity 300			
9,001 – 10,000	4 x Unity 300			

Note: For deployments over 10,000 users, create additional pods using sizing guidance contained herein.

2.5 Shared Tier 1 Rack

Shared Tier 1 for rack servers incorporates shared Tier 1 storage used for execution of VDI sessions.

2.5.1 Shared Tier 1 for Rack Servers

500 users or less

For small deployments of up to 500 users, Tier 1 and Tier 2 can be combined on a single XtremIO Starter X-Brick storage array if desired. In this configuration, Tier 2 storage is provided via the Unity Virtual Storage Adapter (VSA). VSA is implemented as a single (or clustered for HA) VM that utilizes disk resources on the XtremIO Starter X-Brick array.

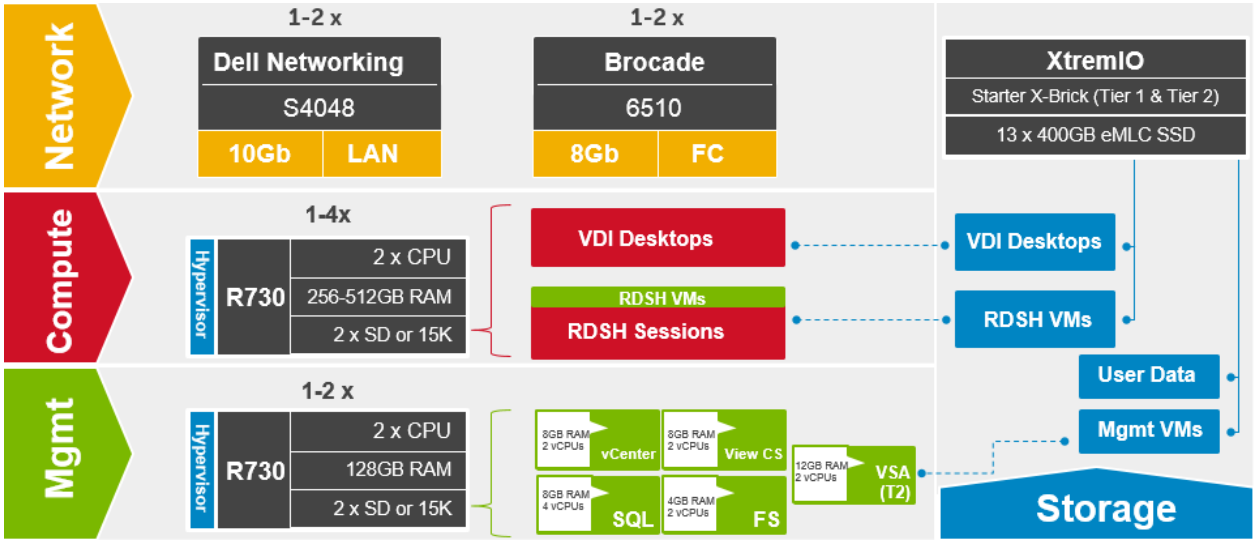

More than 500 users

For over 500 users, the storage layers are separated into discrete arrays, as depicted in the figure below. Additional arrays are added for Tier 1 or Tier 2 as the user count grows.

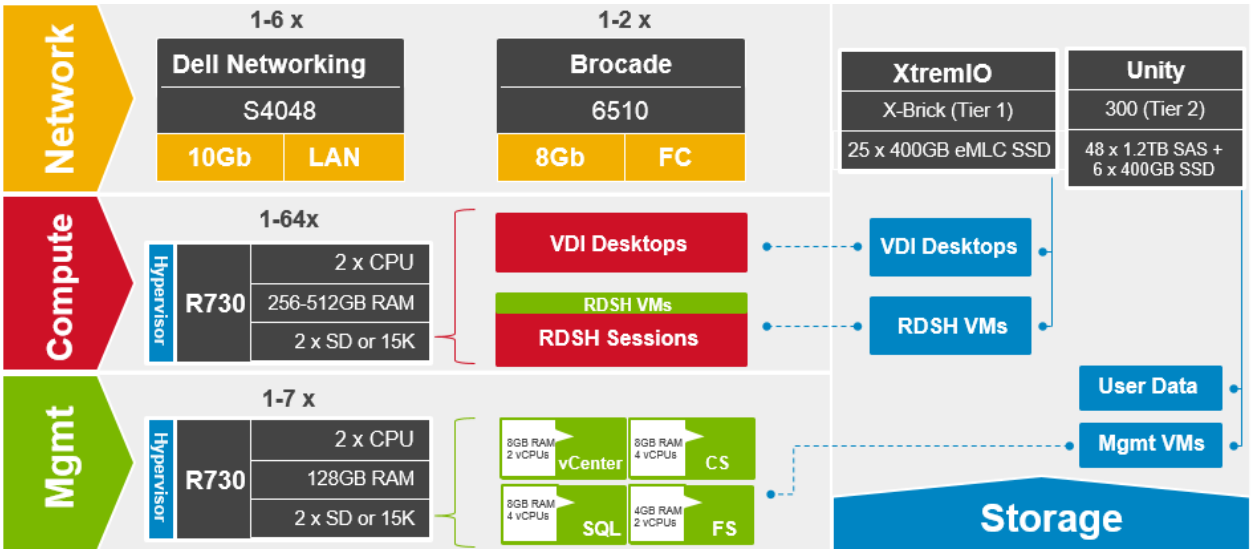

2.5.1.1 Shared Tier 1 – Network Architecture

In the Shared Tier 1 architecture for rack servers using FC, a separate switching infrastructure is required for FC. Management and compute servers both connect to shared storage using FC. Both management and compute servers connect to all network VLANs in this model. All ToR traffic has designed to be layer 2/ switched locally, with all layer 3/ routable VLANs routed through a core or distribution switch. The following diagrams illustrate the server NIC to ToR switch connections, vSwitch assignments, as well as logical VLAN flow in relation to the core switch.

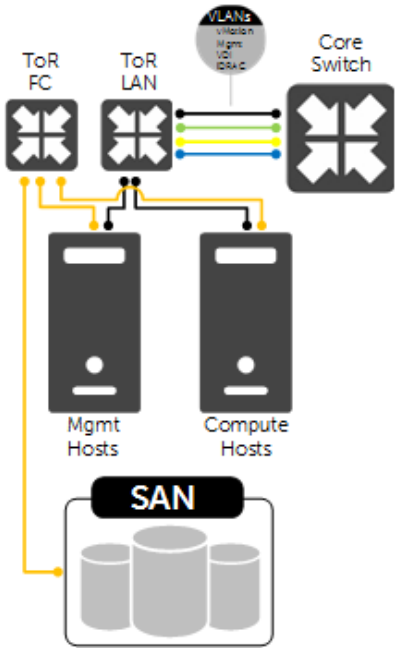

2.5.1.2 Shared Tier 1 – Cabling (HA)

Up to 500 users combined T1 and T2 option

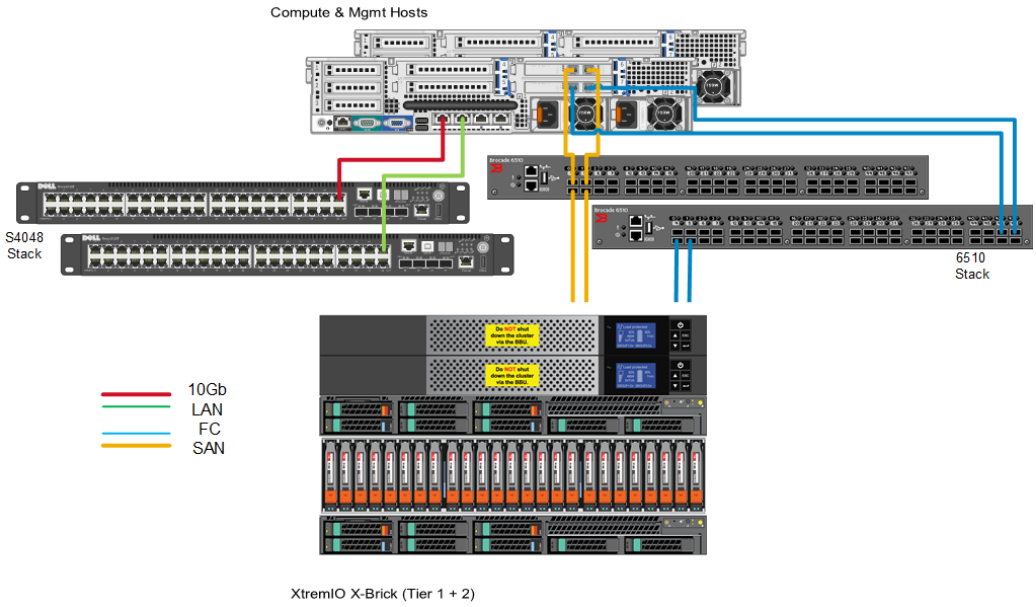

Discrete T1 and T2 storage arrays

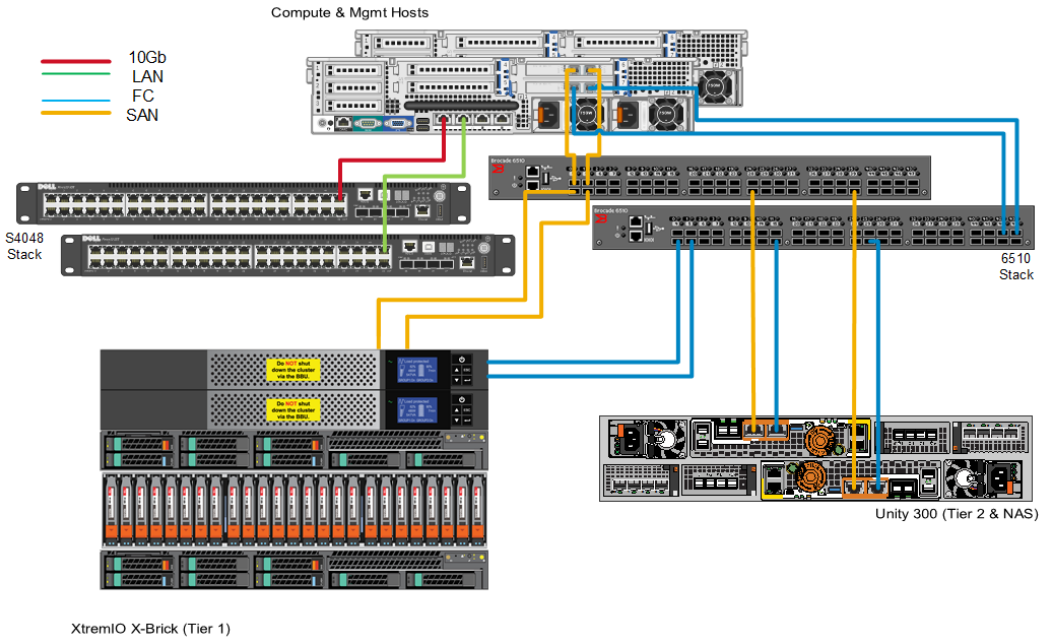

Refer to the Tier 2 – Unity 300 section for a diagram of disk enclosure connections via mini-SAS HD.

2.5.1.3 Shared Tier 1 – Rack Scaling Guidance

Shared Tier 1 HW scaling (Rack - FC)					
User Scale	**XtremIO T1**	**XtremIO T2 & NAS**	**Unity T2 & NAS**	**ToR LAN**	**ToR 8Gb FC**
0 - 500	Starter X-Brick SSD	SSD/VSA	-	S4048	6510
501 – 1,500	Starter X-Brick SSD	-	1 x Unity 300		
1,501 – 3,000	X-Brick SSD				
3,001 – 6,000	2 x X-Brick SSD		2 x Unity 300		
6,001 – 9,000	3 x X-Brick SSD		3 x Unity 300		
9,001 – 10,000	4 x X-Brick SSD		4 x Unity 300		

Note: For deployments over 10,000 users, create additional pods using sizing guidance contained herein.

2.6 Shared Tier 1 Blade

2.6.1 Shared Tier 1 for Blade Servers

As is the case in the Shared Tier 1 model using rack servers, blades can also be used in a 500-user bundle by combing Tier 1 and Tier 2 on an XtremIO Starter X-Brick storage array with Unity VSA. Above 500 users, separate Tier 1 and Tier 2 storage into discrete arrays. In the configurations shown below, ToR FC switching is optional if a suitable FC infrastructure is already in place.

500 users or less

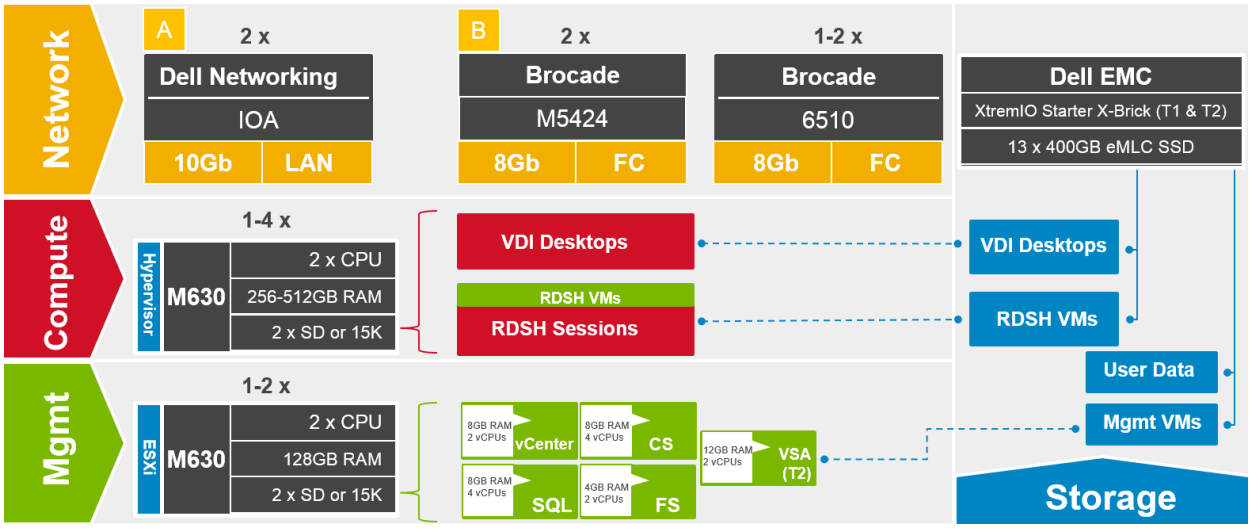

More than 500 users

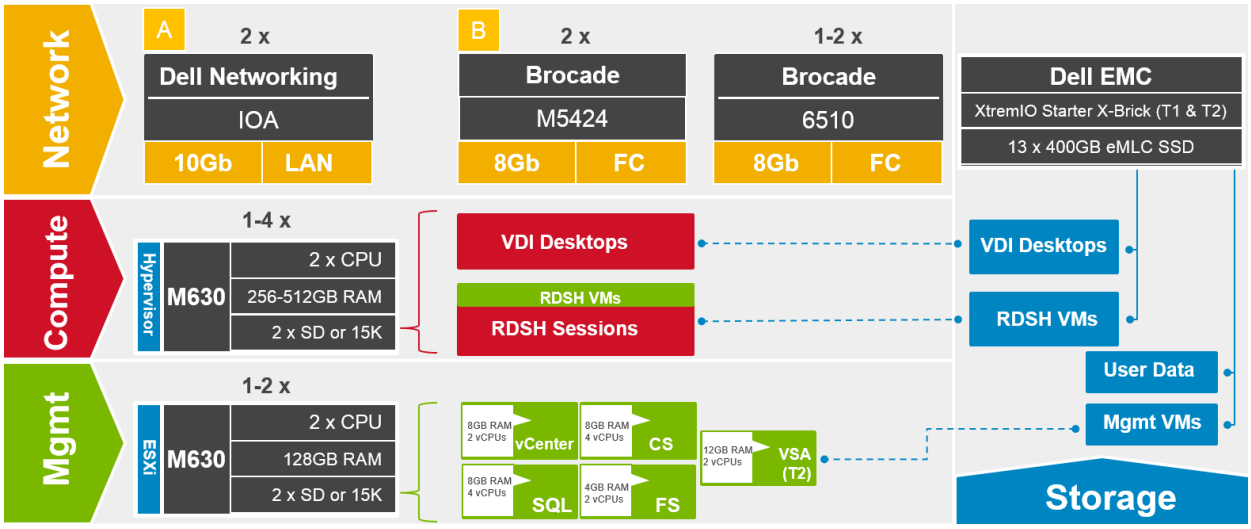

2.6.1.1 Shared Tier 1 – Network Architecture

In the Shared Tier 1 architecture for rack servers using FC, a separate switching infrastructure is required for FC. Management and compute servers both connect to shared storage using FC. Both management and compute servers connect to all network VLANs in this model. All ToR traffic has been designed to be layer 2/ switched locally, with all layer 3/ routable VLANs routed through a core or distribution switch. The following diagrams illustrate the server NIC to ToR switch connections, vSwitch assignments, as well as logical VLAN flow in relation to the core switch.

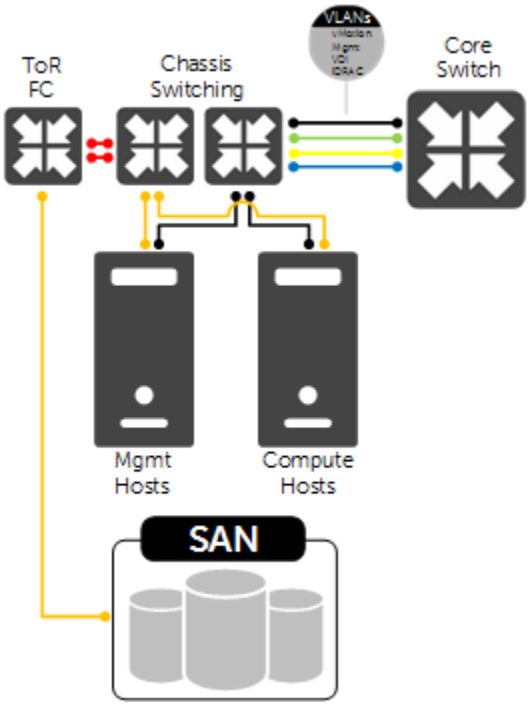

2.6.1.2 Shared Tier 1 – Cabling

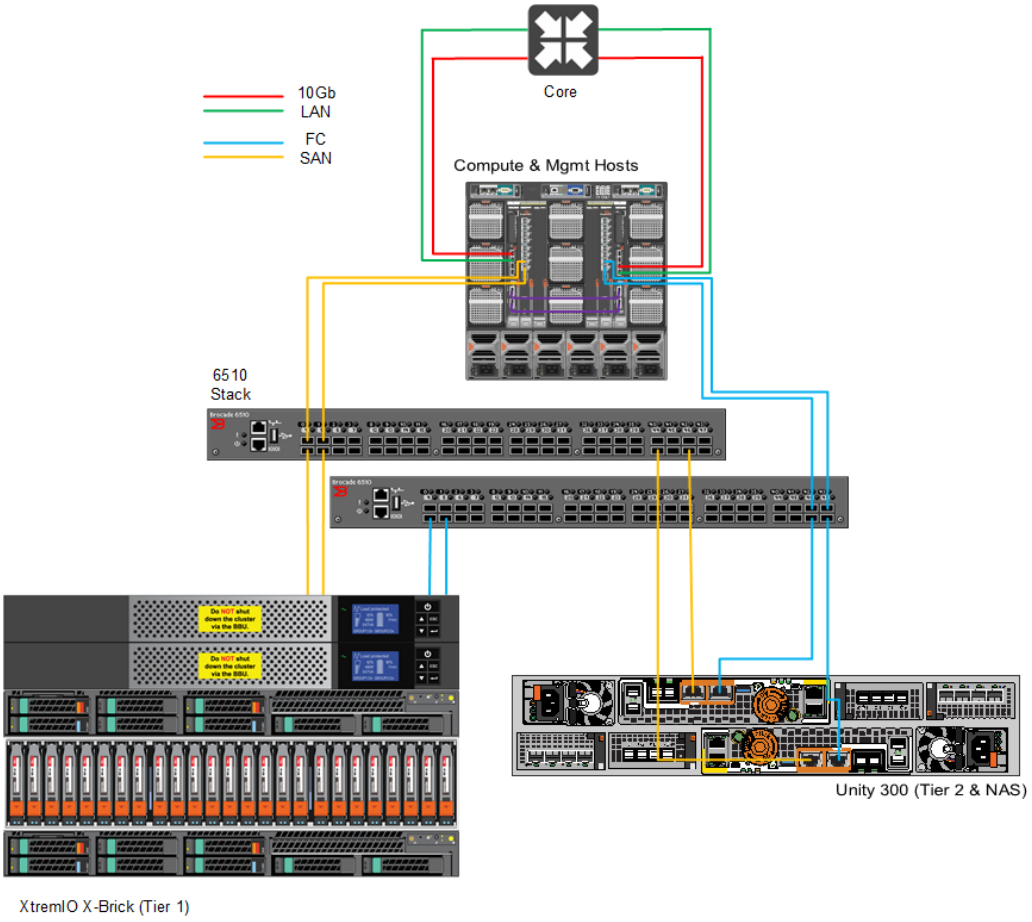

Refer to the Tier 2 – Unity 300 section for a diagram of disk enclosure connections via mini-SAS HD.

2.6.1.3 Shared Tier 1 – Scaling Guidance

Shared Tier 1 HW scaling (Blade - FC)						
User Scale	**XtremIO T1**	**XtremIO T2 & NAS**	**Unity T2 & NAS**	**Blade LAN (A Fabric)**	**Blade FC (B Fabric)**	**ToR 8Gb FC**
0-500	Starter X-Brick SSD	SSD/VSA	-	IOA	M5424	6510
501 – 1,500		-	1 x Unity 300			
1,501 – 3,000	X-Brick SSD					
3,001 – 6,000	2 x X-Brick SSD		2 x Unity 300			
6,001 – 9,000	3 x X-Brick SSD		3 x Unity 300			
9,001 – 10,000	4 x X-Brick SSD		4 x Unity 300			

Note: For deployments over 10,000 users, create additional pods using sizing guidance contained herein.

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)

For out-of-band management such as iDRAC or in environments where 1Gb networking is sufficient, Dell recommends the S3048 network switch. The S3048 is a low-latency top-of-rack (ToR) switch that features 48 x 1GbE and 4 x 10GbE ports, a dense 1U design, and up to 260Gbps performance. The S3048-ON also supports Open Network Installation Environment (ONIE) for zero-touch installation of alternate network operating systems.

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

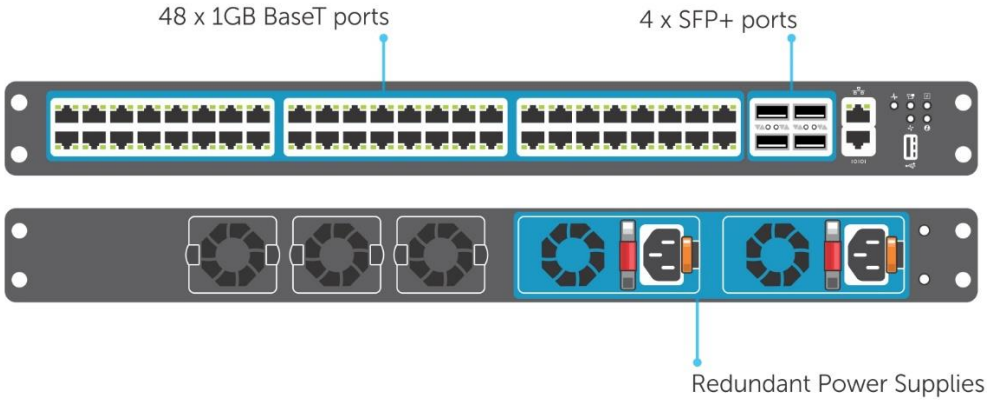

3.1.2 Dell Networking S4048 (10Gb ToR Switch)

Optimize your network for virtualization with a high-density, ultra-low-latency ToR switch that features 48 x 10GbE SFP+ and 6 x 40GbE ports (or 72 x 10GbE ports in breakout mode) and up to 720Gbps performance. The S4048-ON also supports ONIE for zero-touch installation of alternate network operating systems. For BaseT connectivity, the S4048T model is available.

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

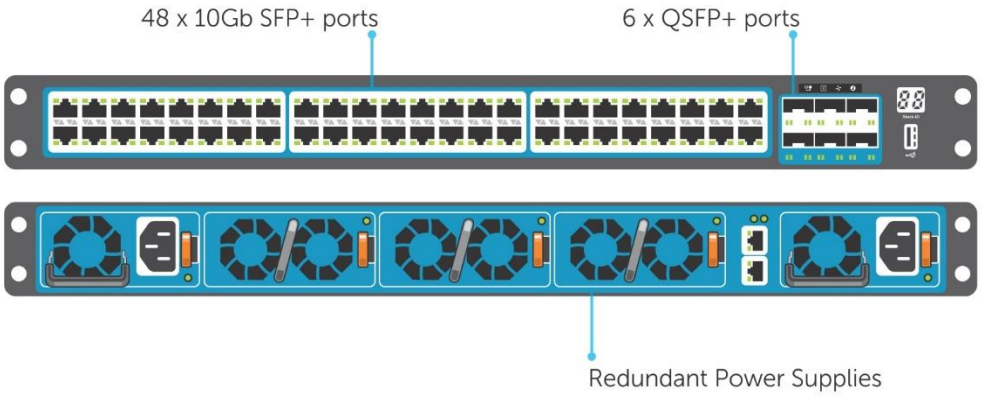

For more information on the S3048, S4048 switches and Dell Networking, please visit: LINK

3.1.3 Brocade 6510 (FC ToR Switch)

The Brocade 6510 Switch meets the demands of hyper-scale, private cloud storage environments by delivering market-leading speeds up to 16Gb Fibre Channel (FC) technology and capabilities that support highly virtualized environments. Designed to enable maximum flexibility and investment protection, the Brocade 6510 is configurable in 24, 36, or 48 ports and supports 2, 4, 8, or 16Gb speeds in an efficiently designed 1U package. It also provides a simplified deployment process and a point-and-click user interface—making it both powerful and easy to use. The Brocade 6510 offers low-cost access to industry-leading Storage Area Network (SAN) technology while providing “pay-as-you-grow” scalability to meet the needs of an evolving storage environment.

Model	Features	Options	Uses
Brocade 6510	48 x 2/4/8/16Gb Fiber Channel Additional (optional) FlexIO module Up to 24 total ports (internal + external)	Ports on demand from 24, 36, and 48 ports	FC ToR switches for all solutions. Optional for blades

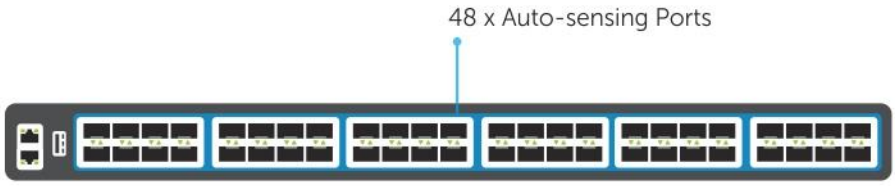

For more information on the Brocade 6510 switch, please visit: LINK

3.1.4 Brocade M5424 (FC Blade Interconnect)

The Brocade® M5424 switches and Dell™ PowerEdge™ M1000e Blade enclosures provide robust solutions for FC SAN deployments. Not only does this offering help simplify and reduce the amount of SAN hardware components required for a deployment, but it also maintains the scalability, performance, interoperability and management of traditional SAN environments. The M5424 can easily integrate FC technology into new or existing storage area network (SAN) environments using the PowerEdge™ M1000e Blade enclosure. The Brocade® M5424 is a flexible platform that delivers advanced functionality, performance, manageability, scalability with up to 16 internal Fabric ports and up to 8 2GB/4GB/8GB auto-sensing uplinks and is ideal for larger storage area networks. Integration of SAN switching capabilities with the M5424 also helps to reduce complexity and increase SAN manageability.

Model	Features	Options	Uses
Brocade M5424	16 x internal Fabric ports Up to 8 2/4/8Gb auto-sensing uplinks	Ports on demand from 12 to 24 ports	Blade switch for FC in Shared Tier 1 model

For more information on the Brocade M5424 switch, please visit: LINK

3.1.5 PowerEdge M I/O Aggregator (10Gb Blade Interconnect)

Simplify network management and increase server bandwidth with the PowerEdge™ M I/O Aggregator, enabling easy, plug-and-play data center convergence.

Model	Features	Options	Uses
PowerEdge M I/O Aggregator (IOA)	Up to 32 x 10Gb ports + 4 x external SFP+ 2 x line rate fixed QSFP+ ports 2 optional FlexIO modules	2-port QSFP+ module in 4x10Gb mode	Blade switch for iSCSI in Shared Tier 1 blade solution, LAN + iSCSI in Local Tier 1 blade solution
		4-port SFP+ 10Gb module	
		4-port 10GBASE-T copper module (one per IOA)	
		Stack up to 2 IOAs using QSFP+ ports	

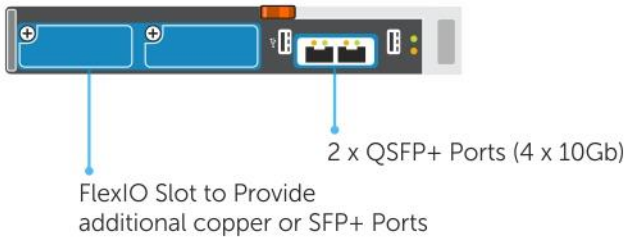

For more information on the Dell IOA switch, please visit: LINK

3.2 Servers

3.2.1 Dell EMC PowerEdge R730

The foundation of the Dell Wyse Datacenter solution portfolio is the best-in-class Dell PowerEdge 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 two double-wide GPU accelerators from AMD or NVIDIA. The Dell PowerEdge R730 offers uncompromising performance and scalability in a 2U form factor.

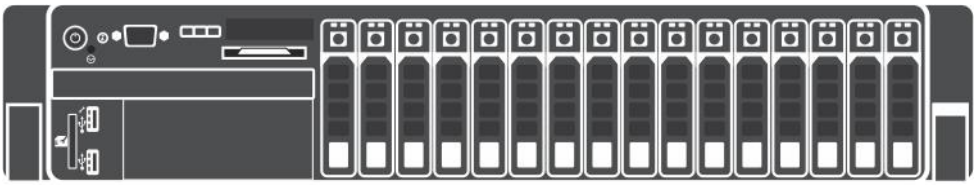

For more information on the R730, please visit: Link

3.2.2 Dell EMC PowerEdge M630

The blade server platform recommendation for the Dell Wyse Datacenter solution is the PowerEdge M630. This half-height blade server is a feature-rich, dual-processor platform that offers a blend of density, performance, efficiency and scalability. The M630 offers remarkable computational density, scaling up to 22 cores, 2 socket Intel Xeon processors(Broadwell) and 24 DIMMs (768GB RAM) of DDR4 memory in an extremely compact half-height blade form factor.

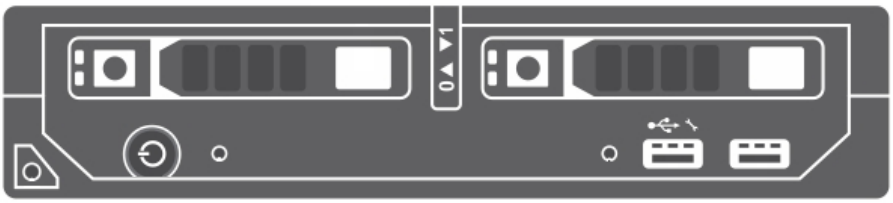

For more information, please visit: Link

3.3 Compute Server Infrastructure

3.3.1 Local Tier 1 Rack

In the Local Tier 1 model, VDI sessions execute from local storage on each Compute server. The hypervisor used in this solution is vSphere. In this model, only the Management server hosts access FC storage to support the solution's Management role VMs done with the 1 GB NICs on the embedded Network Daughter Card (NDC). Additional NICs can be added as required for increased bandwidth or resiliency requirements. Refer to section 2.4 for cabling implications. The Management server host has reduced RAM and CPU and does not require local disk to host the management VMs. All-flash shown for compute below, the SSDs can be optionally substituted for 10-12 15K SAS.

Local Tier 1 Compute Host Dell EMC PowerEdge R730	**Local Tier 1 Management Host Dell EMC PowerEdge R730**
2 x Intel Xeon E5-2698v4 Processor (2.2Ghz)	2 x Intel Xeon E5-2660v4 Processor (2Ghz)
512GB Memory (16 x 32GB RDIMMs, 2400MT/s)	128GB Memory (8 x 16GB RDIMMs, 2400MT/s)
VMware vSphere on internal 8GB Dual SD	VMware vSphere on internal 8GB Dual SD
4 x 800GB SSD	Embedded 4 x 10Gb NDC
PERC H730 Integrated RAID Controller –Raid 10	iDRAC8 Enterprise
Embedded 4 x 10Gb NDC	2 x 750W PSUs
iDRAC8 Enterprise	
2 x 750W PSUs	

3.3.2 Local Tier 1 Blade

In the Local Tier 1 model for blades, VDI sessions execute on local high-performance SSDs on each compute host. vSphere is the supported hypervisor in this solution due to its ability to run from integrated SD freeing the pair of SSDs for VDI execution only. In this model, shared storage is not required for Tier 2 unless management host-level HA is required. All management and desktop VMs are hosted locally on their respective blades.

Local Tier 1 Compute Host – Dell EMC PowerEdge M630
2 x Intel Xeon E5-2698v4 Processor (2.2GHz)
512GB Memory (16 x 32GB RDIMMs, 2400Mhz)
VMware vSphere on 2 x 8GB internal SD
2 x SSD
QLogic 57810S-k 10Gb DP KR NDC (LAN)
iDRAC8 Enterprise w/ vFlash, 8GB SD

Local Tier 1 Management Host – Dell EMC PowerEdge M630
2 x Intel Xeon E5-2660v4 Processor (2GHz)
128GB Memory (8 x 16GB RDIMMs, 2400Mhz)
VMware vSphere on 2 x 8GB internal SD
QLogic 57810S-k 10Gb DP KR NDC (iSCSI-HA)
iDRAC8 Enterprise w/ vFlash, 8GB SD

3.3.3 Shared Tier 1 Rack

Fiber Channel is optionally leveraged as the block storage protocol for Compute and Management hosts with Dell EMC Tier 1 and Tier 2 storage. All configurations are identical except for CPU and RAM which have been reduced on the management hosts.

Shared Tier 1 Compute Host Dell EMC PowerEdge R730	Shared Tier 1 Management Host Dell EMC PowerEdge R730
2 x Intel Xeon E5-2698v4 Processor (2.2Ghz)	2 x Intel Xeon E5-2660v4 Processor (2Ghz)
512GB Memory (16 x 32GB RDIMMs, 2400MT/s)	128GB Memory (8 x 16GB RDIMMs, 2400MT/s)
VMware vSphere on internal 8GB Dual SD	VMware vSphere on internal 8GB Dual SD
Embedded 4 x 10Gb NDC	Embedded 4 x 10Gb NDC
2 x QLogic 2562 8Gb DP FC	2 x QLogic 2562 8Gb DP FC HBA
iDRAC8 Enterprise	iDRAC8 Enterprise
2 x 750W PSUs	2 x 750W PSUs

3.3.4 Shared Tier 1 Blade

The Shared Tier 1 blade server is configured in line with its rack server equivalent. Two network interconnect Fabrics are configured for the blades: the A-Fabric dedicated to 10Gb LAN traffic and the B-Fabric dedicated to 8Gb FC. Please note that FC is only currently supported using vSphere.

Shared Tier 1 Compute Host Dell EMC PowerEdge M630	Shared Tier 1 Management Host Dell EMC PowerEdge M630
2 x Intel Xeon E5-2698v4 Processor (2.2Ghz)	2 x Intel Xeon E5-2660v4 Processor (2Ghz)
512GB Memory (16 x 32GB RDIMMs, 2400MT/s)	128GB Memory (8 x 16GB RDIMMs, 2400MT/s)
VMware vSphere on internal 8GB Dual SD	VMware vSphere on internal 8GB Dual SD
QLogic 57810S-k 10Gb DP KR NDC (LAN)	QLogic 57810S-k 10Gb DP KR NDC (LAN)
1 x QLogic QME2572 8Gb FC mezz (FC)	1 x QLogic QME2572 8Gb FC mezz (FC)
iDRAC8 Enterprise w/ vFlash, 8GB SD	iDRAC8 Enterprise w/ vFlash, 8GB SD

3.4 Storage

3.4.1 XtremIO X-Brick

XtremIO is an all-flash storage solution with built in inline data reduction, innovative data protection and load balancing, VAAI integration, and excellent performance for random I/O requests. It enables enterprises to provision virtual desktops that provide user experiences similar to tablets, ultrabooks, and physical desktops containing solid-state drives (SSDs) (as opposed to typical VDI that tries to mimic the experience of a desktop PC with a hard-disk drive). By ensuring the best user experience for VDI end users, simplifying the management of virtual machines for administrators, and providing an attractive cost per desktop, Dell EMC XtremIO provides customers with a great return on their VDI investment.

A VMware Horizon environment backed by the XtremIO X-Brick for desktop storage and the Unity® unified storage platform for user data provides a high-performance desktop experience with an easy-to-use storage environment. The combination of the XtremIO array and the Unity platform delivers the right storage environment for the complex storage needs of a virtual desktop environment. This system meets or exceeds the requirements of the Citrix Ready VDI Capacity Program.

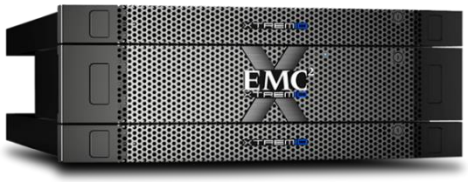

The XtremIO all-flash storage array has a revolutionary architecture with the following elements to enable the agile data center: linear scale-out, inline all-the-time data services, and ample data center services for the workloads.

The basic hardware building block for the XtremIO array is the X-Brick. Each X-Brick is made up of two active-active controller nodes and a disk array enclosure packaged together, presenting no single point of failure.

An X-Brick by itself is a high-availability, high-performance SAN appliance available in 5 TB, 10 TB, 20 TB, and 40 TB capacity configurations that can drive incredible database loads, handle thousands of virtual machines, and support thousands of virtual desktops.

The scale-out, flash-optimized, global data-reduction architecture of XtremIO allows for a number of multiplying effects across many aspects of the array, which in turn leads to a number of key benefits. These benefits include extending the effective capacity of the array as well as minimizing the required writes to media. This improves the XtremIO hosted application performance and increases the usable lifespan of the purchased flash.

The XtremIO all-flash storage array:

- Supports simultaneous iSCSI, Fibre Channel (FC) front-end interconnects
- Scales up to 200 SSD drives via an 8-node cluster, monitored by a single console
- Includes inline compression & data-deduplication features which can dramatically reduce the storage footprint

- Seamlessly integrates with the Unity VSA NAS appliance for fully interoperable block and file storage solutions

XtremIO Tier 1 storage consists of a standard dual controller configuration and scales upward by adding additional discrete arrays, which are managed as a single cluster. XtremIO will support Tier 1 for up to 3000 knowledge worker users, as depicted below, utilizing all eMLC SSD disks. Scaling above this number, additional X-Bricks will need to be implemented. Additional capacity and performance capability is achieved by adding larger disks or X-Bricks, as appropriate, up to the controller's performance limits. Disk sparing is automatically taken into account. RAID is virtualized across all disks in an array (RAIDXP). Please refer to the test methodology and results for specific workload characteristics in section 7.

Controller	Front-End IO	Back-End IO	Disk Shelf	Disks	XIOS (min)
1 x dual-controller X-Brick (256GB)	4 x 8Gb FC cards, 4 x 10Gb iSCSI (2 per controller	Dual redundant SAS interconnect modules	2.5" SAS shelf (25 disks each)	2.5" 400GB SSD (~2100 IOPS each)	4.0

In order to ensure optimal performance for EUC deployments when implementing XtremIO in vSphere environments, we recommend setting several ESXi storage parameters such as DataMover.HardwareAcceleratedInit. For a complete list of parameters and values, please refer to the XtremIO Host Configuration Guide.

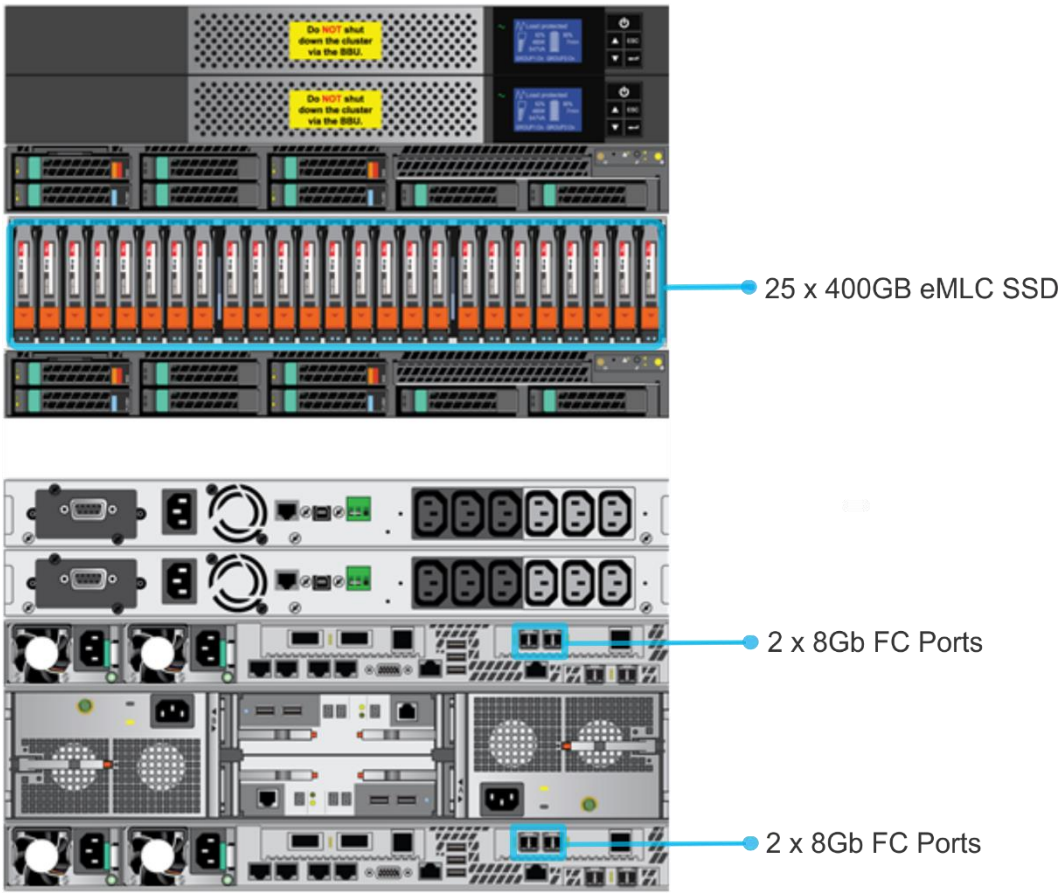

3.4.2 Tier 2 – Unity 300

Dell EMC Unity™ is the only storage system that successfully meets all four requirements of today's IT professionals.

- **Unity is Simple**: Unity Hybrid solutions sets new standards for storage systems with compelling simplicity, modern design, affordable prices, and flexible deployments – to meet the needs of resource-constrained IT professionals in large or small companies.
- **Unity is Modern**: Unity has a modern 2U architecture designed for all-flash, designed to support the high density SSD's including 3D NAND TLC (triple level cell) drives. Unity includes automated data lifecycle management to lower costs, inline compression*, built-in encryption, local point-in-time copies and remote replication, data-in-place conversions, and deep ecosystem integration with VMware and Microsoft.
- **Unity is Affordable**: Our dual-active controller system was designed to optimize the performance, density, and cost of your storage to deliver all-flash or hybrid configurations for much less than you thought possible.
- **Unity is Flexible**: Unity is available as a virtual storage appliance, purpose-built all flash or hybrid configurations, or as converged systems – with one Unity operating environment that connects them all together.

Based on the powerful new family of Intel E5-2600 processors, Unity Hybrid storage systems implement an integrated architecture for block, file, and VMware VVols with concurrent support for native NAS, iSCSI, and Fibre Channel protocols. Each system leverages dual storage processors, full 12 Gb SAS back end connectivity and EMC's patented multi-core architected operating environment to deliver unparalleled performance & efficiency. Additional storage capacity is added via Disk Array Enclosures (DAEs).

*All-Flash pools, block only

Model	Features	Options	Uses
Unity 300	Dual active-active controllers, 24GB cache per controller (cache mirroring), SMB & NFS support, AD-integration.	Each array can support 3,000 concurrent users.	Provides a discrete array for Tier 2 (Management VM and user data storage).

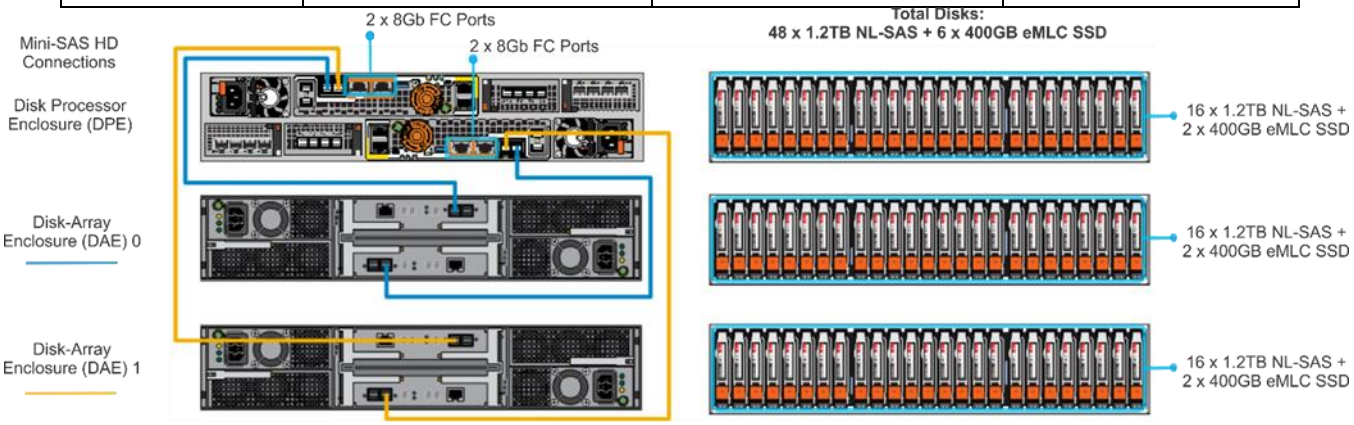

3.4.3 Storage Configuration

Each tier of storage is managed as a separate pool or group to isolate specific workloads. Manage shared Tier 1 arrays used for hosting VDI sessions grouped, while managing shared Tier 2 arrays used for hosting Management server role VMs and user data grouped separately.

3.4.4 XtremIO Starter X-Brick

Unity Tier 2 storage is optional if a customer has 500 users or less and wishes to deploy both data tiers on a single array. For this purpose, we recommend implementing NAS service on the XtremIO Starter X-Brick using the Unity Virtual Storage Adapter (VSA). This solution is ideal for smaller deployments that wish to avoid additional expense of managing discrete arrays for each Tier. For larger deployments, we recommend utilizing discrete Unity arrays to provide Tier 2 services. VSA is implemented as a single (or clustered for HA) VM that utilizes disk resources on the XtremIO X-Brick to provide file services for user profiles, and user data.

The recommended Starter X-Brick configuration for this purpose is listed in the table below.

Controller	Front-End IO	Back-End IO	Disk Shelf	Disks
1 x dual-controller X-Brick (256GB)	4 x 8Gb FC cards, 4 x 10Gb iSCSI (2 per controller	Dual redundant SAS interconnect modules	2.5” SAS shelf (13 disks each)	2.5” 400GB SSD (~2100 IOPS each)

3.4.5 Network Attached Storage (NAS)

3.4.5.1 Unity Virtual Storage Appliance

The UnityVSA™ (Unity™ Virtual Storage Appliance) is ideal for customers with 500 users or less using the Starter X-Brick for combined Tier 1 and Tier 2 storage. The UnityVSA allows the advanced unified storage and data management features of the Dell EMC Unity family to be easily deployed on a VMware ESXi server, enabling customers to implement an affordable software-defined solution. Benefits of this approach are a low acquisition cost option for hardware consolidation, multi-tenant storage instances, remote/branch office storage environment and easier to build/maintain/destroy environment for staging and testing.

With all-inclusive software, the UnityVSA allows you to:

- Set up for NAS or SAN in just a few minutes using the Unisphere wizards
- Allow VMware administrators to manage storage from within VMware vCenter
- Protect data locally with unified point-in-time snapshots
- Replicate data remotely to other instances of UnityVSA or to other Unity purpose-built storage arrays
- Optimize performance, efficiency, and simplify storage management with automated-tiering through EMC Fully-Automated Storage Tiering Virtual Pools (FAST VP)
- Administrate the storage array using the same HTML-5 based Unisphere as Unity purpose-built storage arrays.
- Provide advanced file services to block storage for your customers

The Unity VSA has the following characteristics:

- 2 vCPU
- 12 GB memory
- Supports 500 users based upon the Login VSI Knowledge Worker4.1.5+ workload

3.4.5.2 Unity 300

In addition to providing Tier 2 storage for the Management layer, the Unity 300 array also provides NAS services for user data storage for configurations with over 500 users.

3.5 GPUs

3.5.1 NVIDIA Tesla GPUs

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

3.5.1.1 NVIDIA Tesla M10

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

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

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

3.5.1.2 NVIDIA Tesla M60

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

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

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

3.6 Dell Wyse Thin Clients

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

3.6.1 Wyse 5030 PCoIP Zero Client

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

3.6.2 Wyse 5040 AIO Thin Client with PCoIP

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

3.6.3 Wyse 5050 AIO PCoIP Zero Client

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

3.6.4 Wyse 7030 PCoIP Zero Client

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

3.6.5 Wyse 5060 Thin Client (ThinOS) with PCoIP

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

3.6.6 Wyse 7040 Thin Client with Windows Embedded Standard 7P

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

3.6.7 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.

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

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

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

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

4 Software Components

4.1 VMware

4.1.1 VMware vSphere 6.x

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

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

Throughout the Dell Wyse Datacenter solution, all VMware and Microsoft best practices and prerequisites for core services are adhered to (NTP, DNS, Active Directory, etc.). The vCenter used in the solution is a vCenter Server Appliance (VCSA) residing on a host in the management Tier. Horizon Composer is installed on a standalone Windows Server 2012 R2 VM when using the VCSA.

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

Improved Appliance Management

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

VMware vCenter High Availability

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

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

Backup and Restore

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

vSphere HA Support for NVIDIA GRID vGPU Configured VMs

vSphere HA now protects VMs with the NVIDIA GRID vGPU shared pass-through device. In the event of a failure, vSphere HA attempts to restart the VMs on another host that has an identical NVIDIA GRID vGPU profile. If there is no available healthy host that meets this criterion, the VM fails to power on.

For more information on VMware vSphere and what's new in this release, visit link.

4.1.2 VMware Horizon

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

VMware Horizon provides a complete virtual desktop delivery system by integrating several distributed components with advanced configuration tools that simplify the creation and real-time management of the virtual desktop infrastructure. For the complete set of details, please see the Horizon resources page at http://www.vmware.com/products/horizon-view/resources.html.

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

The core Horizon components include:

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

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

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

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

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

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

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

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

4.1.2.1 What's new in this release of Horizon

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

4.1.4.1.1 Just in time delivery with Instant Clone Technology

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

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

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

4.1.4.1.2 Transformational user experience with Blast Extreme

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

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

4.1.4.1.3 Modernize application lifecycle management with App Volumes

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

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

4.1.4.1.4 Smart policies with streamlined access

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

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

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

4.1.3 VMware Horizon Apps

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

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

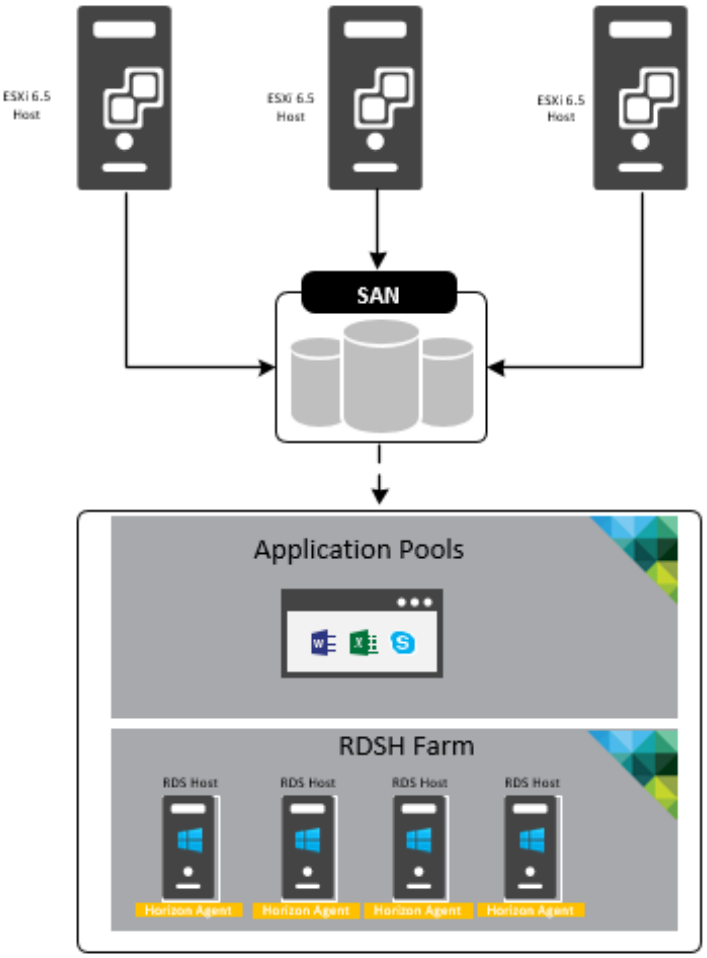

VMware Horizon features and components such as the Blast Extreme display protocol, instant-clone provisioning, App Volumes application delivery, and User Environment Manager are heavily integrated into RDSH to provide a seamless user experience and an easy-to-manage, scalable solution.

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

Horizon JMP Apps are composed of:

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

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

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

4.2 NVIDIA GRID vGPU

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

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

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

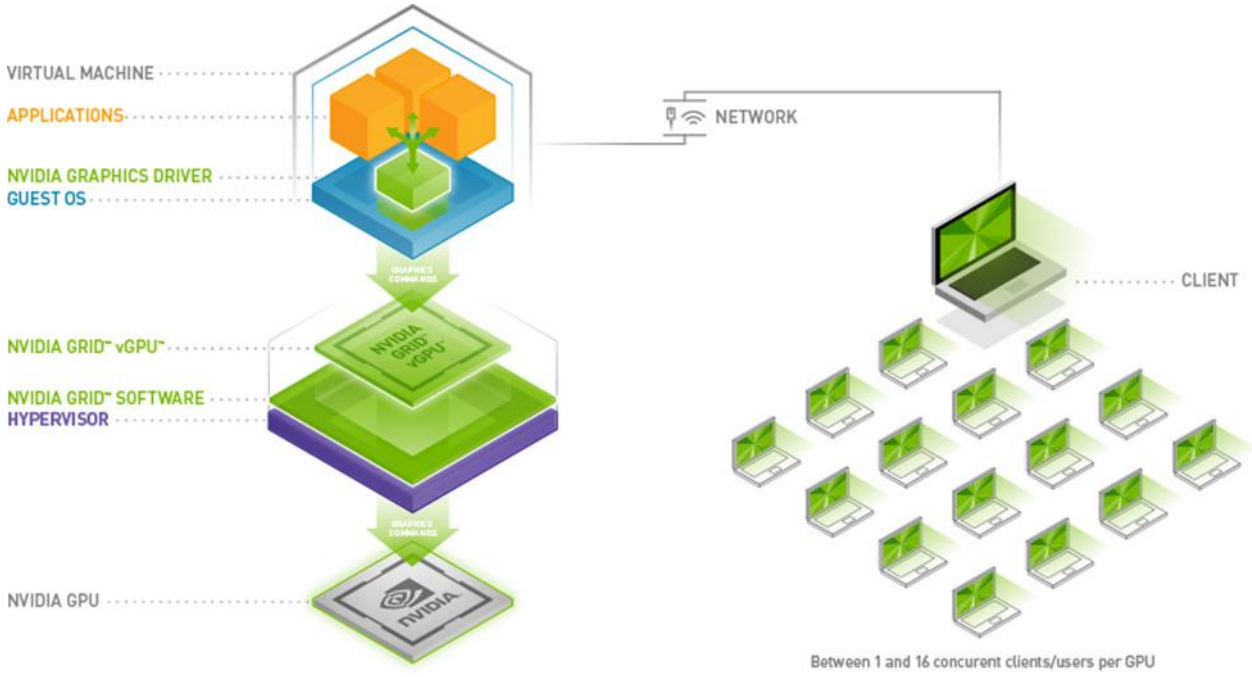

(Image provided by NVIDIA Corporation. Copyright NVIDIA Corporation)

4.2.1 vGPU Profiles

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

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

For more information about NVIDIA GRID vGPU, please visit: LINK

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

NVIDIA® Tesla® M10 GRID vGPU Profiles:

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

	M10-0B	512MB	2	2560x1600	16	64	128
	M10-8A	8GB	1	1280x1024	1	4	8
	M10-4A	4GB			2	8	16
	M10-2A	2GB			4	16	32
	M10-1A	1GB			8	32	64

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

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

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

NVIDIA® Tesla® M60 GRID vGPU Profiles:

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

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

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

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

4.2.1.1 GRID vGPU Licensing and Architecture

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

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

**0Q profiles only support up to 2560x1600 resolution

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

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

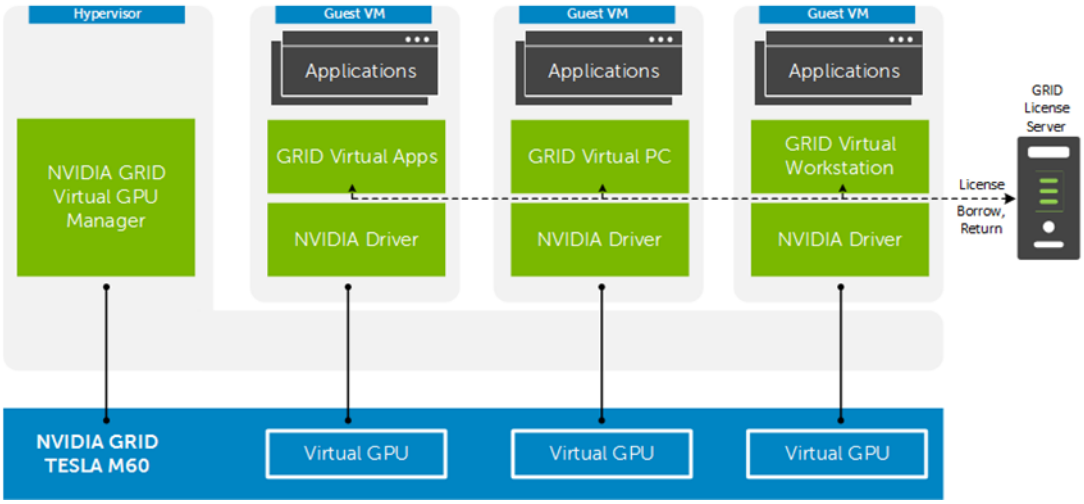

(Image provided courtesy of NVIDIA Corporation, Copyright NVIDIA Corporation)

5 Solution Architecture for Horizon 7

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 volume. Present Tier 2 volumes with a special purpose (called out above) in the format specified below:

Role	vCPU	RAM (GB)	NIC	OS + Data vDisk (GB)	Tier 2 Volume (GB)
VMware vCenter	2	8	1	290	-
Horizon Connection Server	2	8	1	60	-
SQL Server	5	8	1	60	210 (VMDK)
File Server	1	4	1	60	2048 (RDM)
VSA(Optional)	2	12	1	30(3 x 10)	
Total	12	40	5	500	2358

5.1.1 RDSH VM Configuration

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

RDSH VM configuration on ESXi

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

5.1.2 NVIDIA GRID License Server Requirements

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

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

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

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

Additional license server requirements:

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

5.1.3 SQL Databases

The VMware databases will be hosted by a single dedicated SQL 2016 (or higher) Server VM (check DB compatibility at Link . Use caution during database setup to ensure that SQL data, logs and TempDB are properly separated onto their respective volumes. Create all Databases that will be required for:

- Horizon Connection Server
- vCenter
- Horizon Composer

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

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

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

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

5.1.4 DNS

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

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

5.1.4.1 DNS for SQL

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

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

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

5.2 Storage Architecture Overview

The Dell Wyse Datacenter solution has a wide variety of Tier 1 and Tier 2 storage options to provide maximum flexibility to suit any use case. Customers have the choice to leverage best-of-breed Dell EMC storage solutions using Fibre Channel or iSCSI while being assured the storage Tiers of the Dell Wyse Datacenter solution will consistently meet or outperform user needs and expectations.

5.2.1 Local Tier 1 Storage

Selecting the local Tier 1 storage model means that the compute host servers use 10 locally installed hard drives to house the user desktop VMs. In this model, Tier 1 storage exists as local hard disks or SSDs on the Compute hosts themselves. To achieve the required performance level, RAID 10 is recommended for use across all local disks. A single volume per local Tier 1 Compute host is sufficient to host the provisioned desktop VMs along with their respective write caches.

5.2.2 Shared Tier 1 Storage

Selecting the Shared Tier 1 model means that the virtualization compute hosts are deployed without Tier 1 local storage and leverage shared storage hosted on a high performance Dell Storage array. In this model, shared storage is leveraged for Tier 1 and used for VDI execution and write cache. Based on the heavy performance requirements of Tier 1 for VDI, it is recommended to use separate arrays for Tier 1 and Tier 2 when possible. We recommend using 500GB LUNs for VDI and running no more than 125 VMs per volume along with their respective write caches. Sizing to 500 basic users will require 4 x 500GB volumes.

Volumes	Size (GB)	Storage Array	Purpose	File System
VDI-1	500	Tier 1	125 x desktop VMs + WC	VMFS or NTFS
VDI-2	500	Tier 1	125 x desktop VMs + WC	VMFS or NTFS
VDI-3	500	Tier 1	125 x desktop VMs + WC	VMFS or NTFS
VDI-4	500	Tier 1	125 x desktop VMs + WC	VMFS or NTFS

5.2.3 Shared Tier 2 Storage

Tier 2 is shared file storage (CIFS and/or NFS) used to host the Management server VMs and user data. For smaller deployments (<=500 users), the XtremIO Starter X-Brick can host both block & file services via the Unity Virtualized Storage Appliance (VSA). Unity 300 arrays are used for Shared Tier 2 (Local Tier 1 only) for configurations using over 500 users up to 10k users. The table below outlines the volume requirements for Tier 2. Larger disk sizes are chosen to meet the capacity needs of the customer. The user data are presented either via a file server VM using VSA for small-scale deployments or via NAS for large scale or HA deployments. The solution as designed presents all SQL disks using VMDK or VHDX formats. RAID XP is used in smaller deployments hosted on XtremIO. The recommendation for larger scale and deployments is to use an external Unity array with RAID 6 to maximize performance and recoverability. The following depicts the component volumes required to support a 500-user environment. Additional Management volumes are created as needed along with size adjustments as applicable for user data and profiles.

Volumes	Size (GB)	Storage Array	Purpose	File System
Management	350	Tier 2	vCenter, Horizon Connection Server, File and SQL	VMFS
User Data	2048	Tier 2	File Server/ NAS	RDM/NTFS
User Profiles	20	Tier 2	User profiles	VMFS
SQL DATA	100	Tier 2	SQL	VMFS
SQL LOGS	100	Tier 2	SQL	VMFS
TempDB Data	5	Tier 2	SQL	VMFS
TempDB Logs	5	Tier 2	SQL	VMFS
SQL Witness	1	Tier 2	SQL (optional)	VMFS
Templates/ ISO	200	Tier 2	ISO storage (optional)	VMFS

5.2.4 Storage Networking – XtremIO Fibre Channel (FC)

The XtremIO all-flash array provides built-in intelligence and automation to dynamically manage enterprise data throughout its lifecycle. Together, block-level intelligence, storage virtualization, integrated software and modular, platform-independent hardware enable exceptional efficiency, simplicity and security.

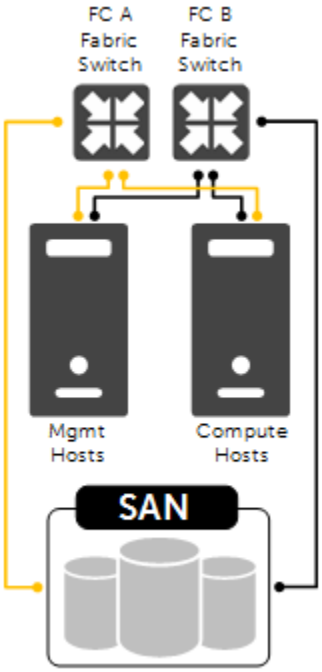

XtremIO actively manages data at a block level using real-time intelligence, providing fully virtualized storage at the disk level. Resources are pooled across the entire storage array. All virtual volumes are thin-provisioned. With inline data compression and dedupe, physical storage requirements can be vastly reduced.

If Fiber Channel is the selected block storage protocol, then the XtremIO Integration for VMware vSphere client plug-in is installed on all hosts. This plugin enables all newly created data stores to be automatically aligned at the recommended 4MB offset. Although a single Fabric are configured to begin with to reduce costs, as a best practice recommendation, the environment is configured with two Fabrics to provide multi-pathing and end-to-end redundancy.

The following QLogic HBA BIOS settings are used:

- Set the “connection options” field to 1 for point to point only
- Set the “login retry count” field to 60 attempts
- Set the “port down retry” count field to 60 attempts
- Set the “link down timeout” field to 30 seconds
- Set the “queue depth” (or “Execution Throttle”) field to 255
- This queue depth are set to 255 because the ESXi VMkernel driver module and DSNRO can more conveniently control the queue depth.

5.2.4.1 FC Zoning

Zone at least one port from each server HBA to communicate with each XtremIO controller. The result of this is two distinct FC Fabrics and four redundant paths per server as shown in the diagram below. Round Robin or Fixed Paths are supported. You can leverage Dell EMC ViPR software to ease storage management in a heterogeneous environment.

5.3 Virtual Networking

5.3.1 Local Tier 1

The network configuration in this model will vary between the Compute and Management hosts. The Compute hosts do not need access to FC storage since they are hosting VDI VMs on local disk. Since the Management VMs are hosted on shared storage, they can take advantage of HA including Live Migration.

The following outlines the VLAN requirements for the Compute and Management hosts in this solution model, applicable to rack or blade servers with HA:

- Compute hosts (Local Tier 1)
  - Management VLAN: Configured for hypervisor infrastructure traffic – L3 routed via core switch
  - VDI VLAN: Configured for VDI session traffic – L3 routed via core switch
- Management hosts (Local Tier 1)
  - Management VLAN: Configured for hypervisor Management traffic – L3 routed via core switch
  - Live Migration VLAN: Configured for Live Migration traffic – L2 switched only, trunked from Core (HA only)
  - Failover Cluster VLAN: Configured for Cluster and Cluster Shared Volume traffic – L2 switched only, trunked from core (Hyper-V only)
  - VDI Management VLAN: Configured for VDI infrastructure traffic – L3 routed via core switch
- A VLAN for iDRAC is configured for all hardware management traffic – L3 routed via core switch

Following best practices, LAN and block storage traffic is separated in solutions >500 users. This traffic is combined within a single switch in smaller stacks to minimize the initial investment; however, VLANs are required for each traffic type to enable traffic separation. Configure the LAN traffic from the server to the ToR switch as a LAG.

vDSwitches should be used as desired for VM traffic especially in larger deployments to ease the management burden across numerous hosts. In the Local Tier 1 rack model the MGMT hosts connect to shared storage and require additional VMK ports. Network share values should be configured equally among the VMKernel port groups that share a physical set of network adapters.

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

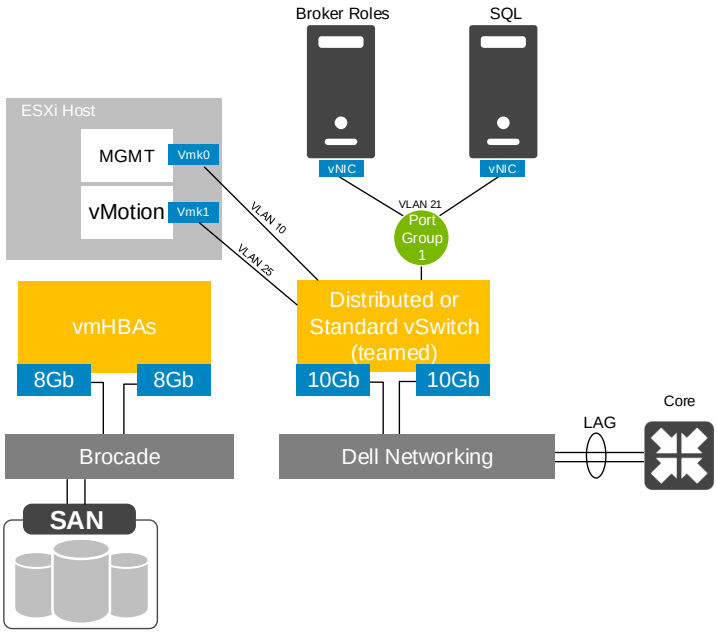

The Compute hosts are configured in the same basic manner, minus the shared storage, with the desktop VMs connecting to the primary port group on the external vSwitch.

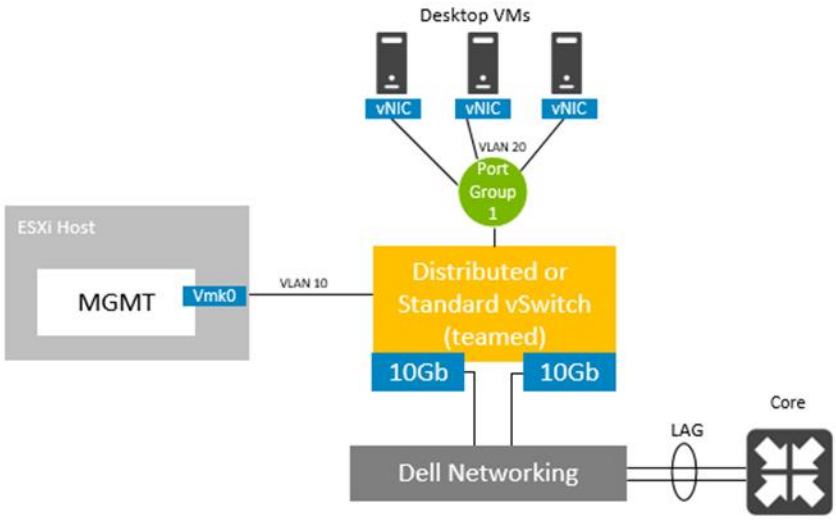

5.3.2 Shared Tier 1

Using Fiber Channel based storage requires additional storage fabrics to be built out in the network stack. The network configuration in this model is identical between the Compute and Management hosts. The benefits of shared storage are available to all hosts such as Live Migration and HA. The following outlines the VLAN requirements for the Compute and Management hosts in this solution model:

- Compute hosts (Shared Tier 1)
  - Management VLAN: Configured for hypervisor Management traffic – L3 routed via core switch
  - Live Migration VLAN: Configured for Live Migration traffic – L2 switched only, trunked from Core
  - Failover Cluster VLAN: Configured for Cluster and Cluster Shared Volume traffic – L2 switched only, trunked from core (Hyper-V only)
  - VDI VLAN: Configured for VDI session traffic – L3 routed via core switch
- Management hosts (Shared Tier 1)
  - Management VLAN: Configured for hypervisor Management traffic – L3 routed via core switch
  - Live Migration VLAN: Configured for Live Migration traffic – L2 switched only, trunked from Core
  - Failover Cluster VLAN: Configured for Cluster and Cluster Shared Volume traffic – L2 switched only, trunked from core (Hyper-V only)
  - VDI Management VLAN: Configured for VDI infrastructure traffic – L3 routed via core switch
- A VLAN for iDRAC is configured for all hardware management traffic – L3 routed via core switch

FC and LAN traffic are physically separated into discrete switching Fabrics. Each Shared Tier 1 Compute and Management host have a quad port NDC (4 x 10Gb) as well as 2 x 8Gb dual port FC HBAs. LAN traffic from the server to the ToR switch is configured as a LAG.

The same basic configuration applies to rack or blade servers although the physical NIC and switching components differ. Network share values should be configured equally among the VMkernel port groups that share a physical set of network adapters.

Management Servers

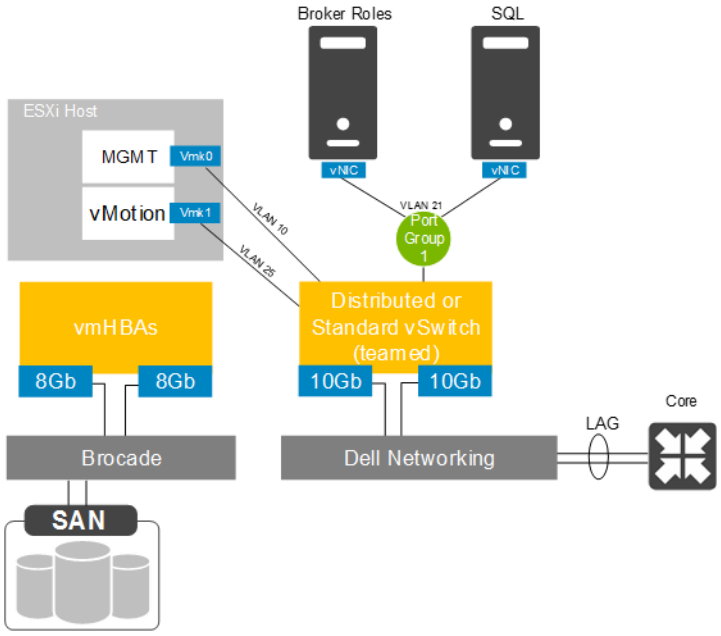

Compute Servers

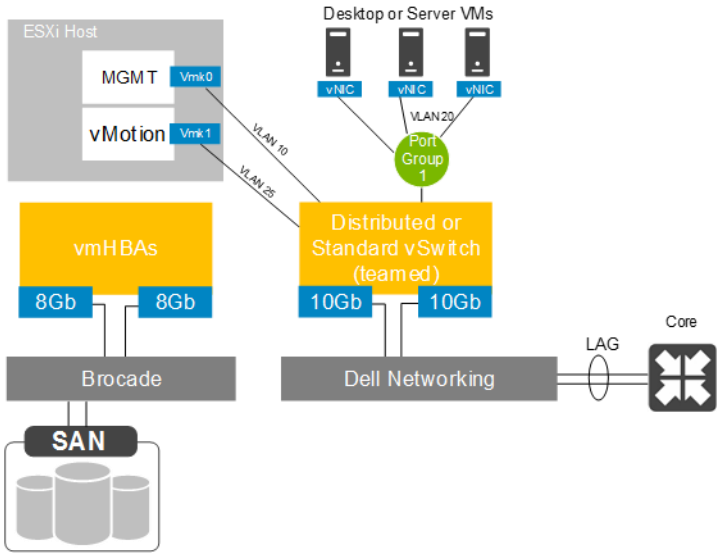

5.3.3 VMware NSX

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

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

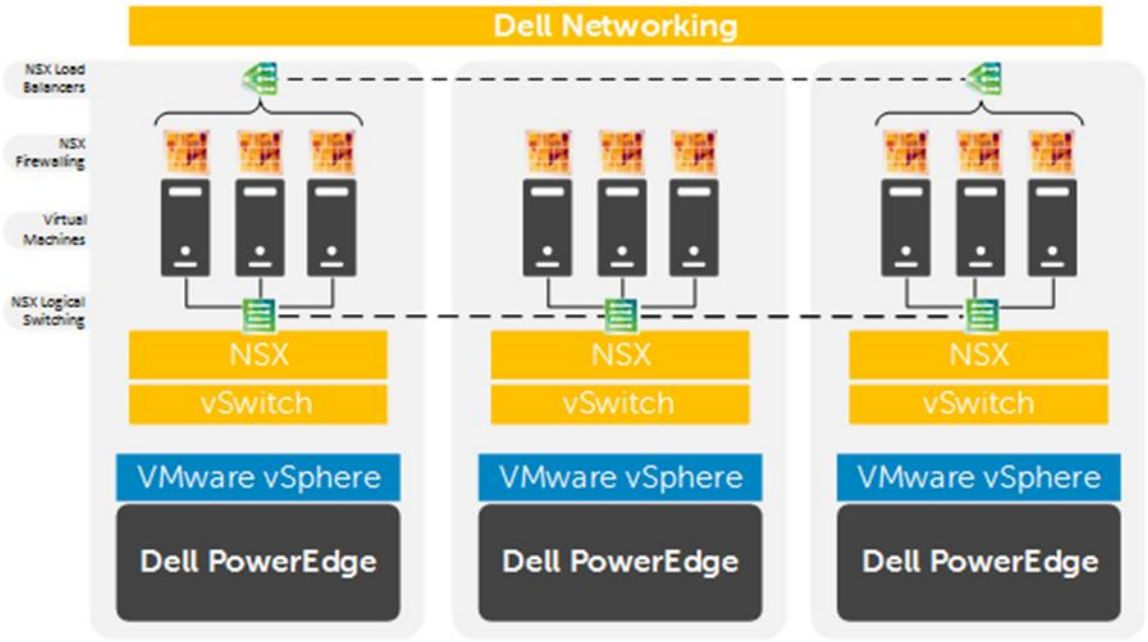

Key Features of Dell Open Networking and VMware NSX	
Power of Choice	Choose from best-of-breed open networking platforms, operating systems and applications.
Accelerated Innovation	Take advantage of open networking with open source standards-based tools and expertise to help accelerate innovation.
Open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

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

5.4 Scaling Guidance

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

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

5.5 Solution High Availability

High availability (HA) is offered to protect each architecture solution layer, individually if desired. Following the N+1 model, additional ToR switches are added to the Network layer and stacked to provide redundancy as required, additional compute and management hosts are added to their respective layers, vSphere or Hyper-V clustering is introduced in both the management and compute layers, SQL is configured with AlwaysOn and NetScaler is leveraged for load balancing. Storage protocol switch stacks and NAS selection will vary based on chosen solution architecture.

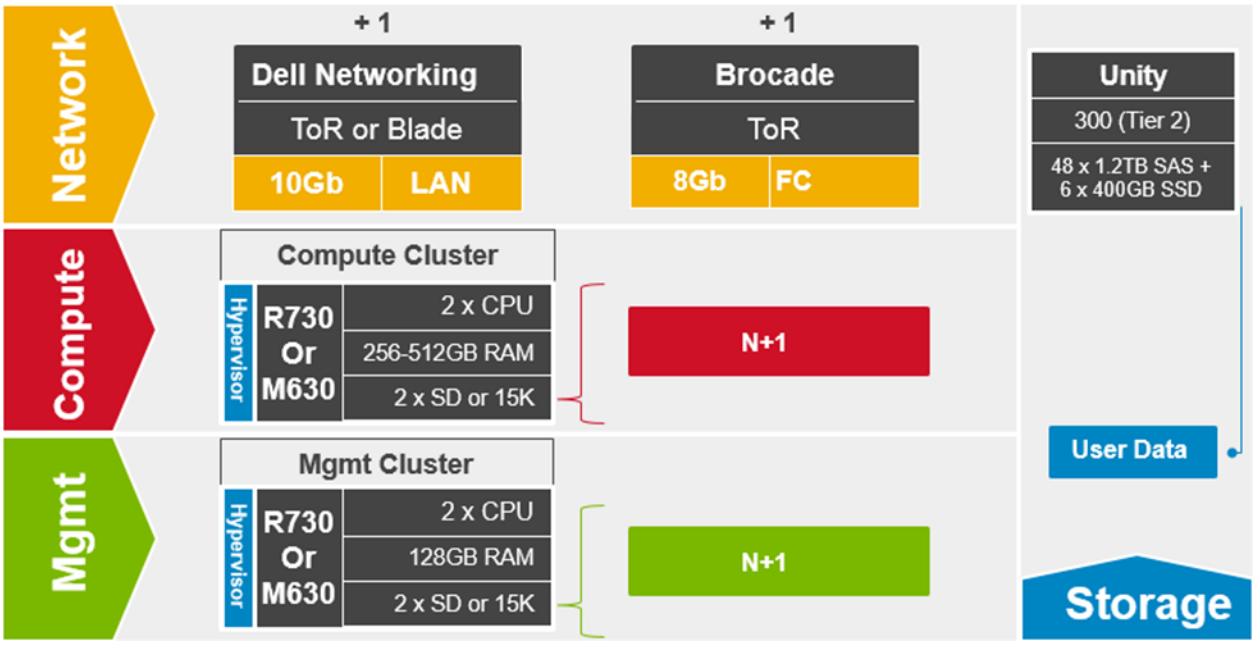

The HA options provide redundancy for all critical components in the stack while improving the performance and efficiency of the solution as a whole.

- Additional switches added to the existing thereby equally spreading each host's network connections across multiple switches.
- Additional ESXi hosts added in the compute or mgmt layers to provide N+1 protection.
- A number of enhancements occur at the Management tier, the first of which is the addition of another host. The Management hosts will then be configured in an HA cluster. All applicable Horizon server roles can then be duplicated on the new host where connections to each will be load balanced via the addition of a F5 Load Balancer. SQL will also receive greater protection through the addition and configuration of a SQL mirror with a witness.

5.5.1 Compute layer HA (Local Tier 1)

The optional HA bundle adds an additional host in the Compute and Management layers to provide redundancy and additional processing power to spread out the load. The Compute layer in this model does not leverage shared storage so hypervisor HA does not provide a benefit here. If a single host fails, another will need to be spun up in the cluster or extra server capacity can be pre-configured and running in active status to handle the reconnection/startup of new desktops to accommodate the users from failed host.

Because only the Management hosts have access to shared storage, in this model, only these hosts need to leverage the full benefits of hypervisor HA. The Management hosts can be configured in an HA cluster with or without the HA bundle. An extra server in the Management layer will provide protection should a host fail.

vSphere HA Admission control can be configured one of three ways to protect the cluster. This will vary largely by customer preference but the most manageable and predictable options are percentage reservations or a specified hot standby. Reserving by percentage will reduce the overall per host density capabilities but will make some use of all hardware in the cluster. Additions and subtractions of hosts will require the cluster to be manually rebalanced. Specifying a failover host, on the other hand, will ensure maximum per host density numbers but will result in hardware sitting idle.

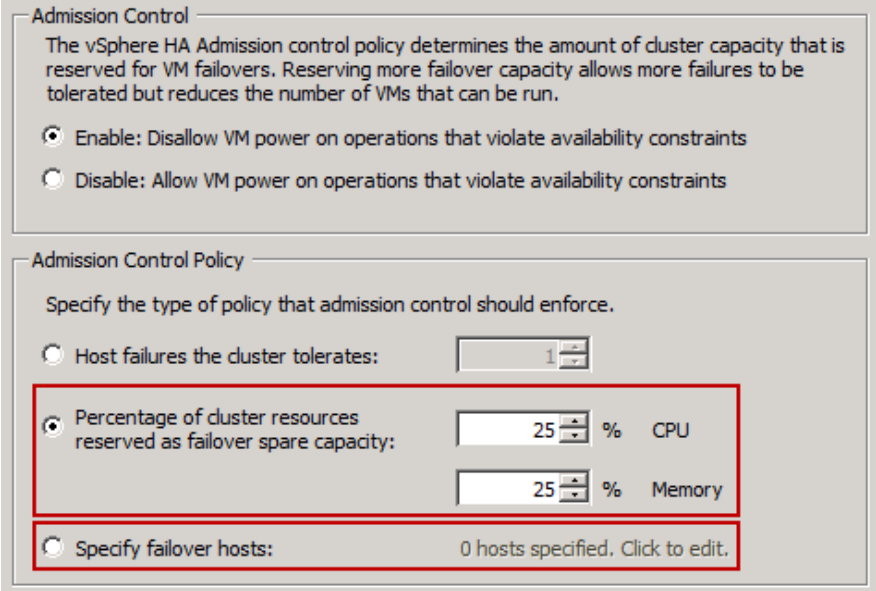

5.5.2 vSphere HA (Shared Tier 1)

Both compute and management hosts are identically configured, within their respective tiers and leverage shared storage so can make full use of vSphere HA. The Compute hosts can be configured in an HA cluster following the boundaries of vCenter with respect to limits imposed by VMware (2000 VMs per vCenter). This will result in multiple HA clusters managed by multiple vCenter servers.

A single HA cluster will be sufficient to support the Management layer up to 10K users. An additional host can be used as a hot standby or to thin the load across all hosts in the cluster.

5.5.3 Management server high availability

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

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

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

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

5.5.4 Horizon CS high availability

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

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

5.5.5 SQL Server High Availability

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

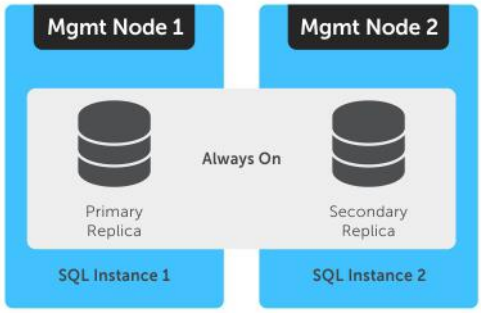

Please refer to these links for more information: LINK1 and LINK2

5.6 VMware Horizon communication flow

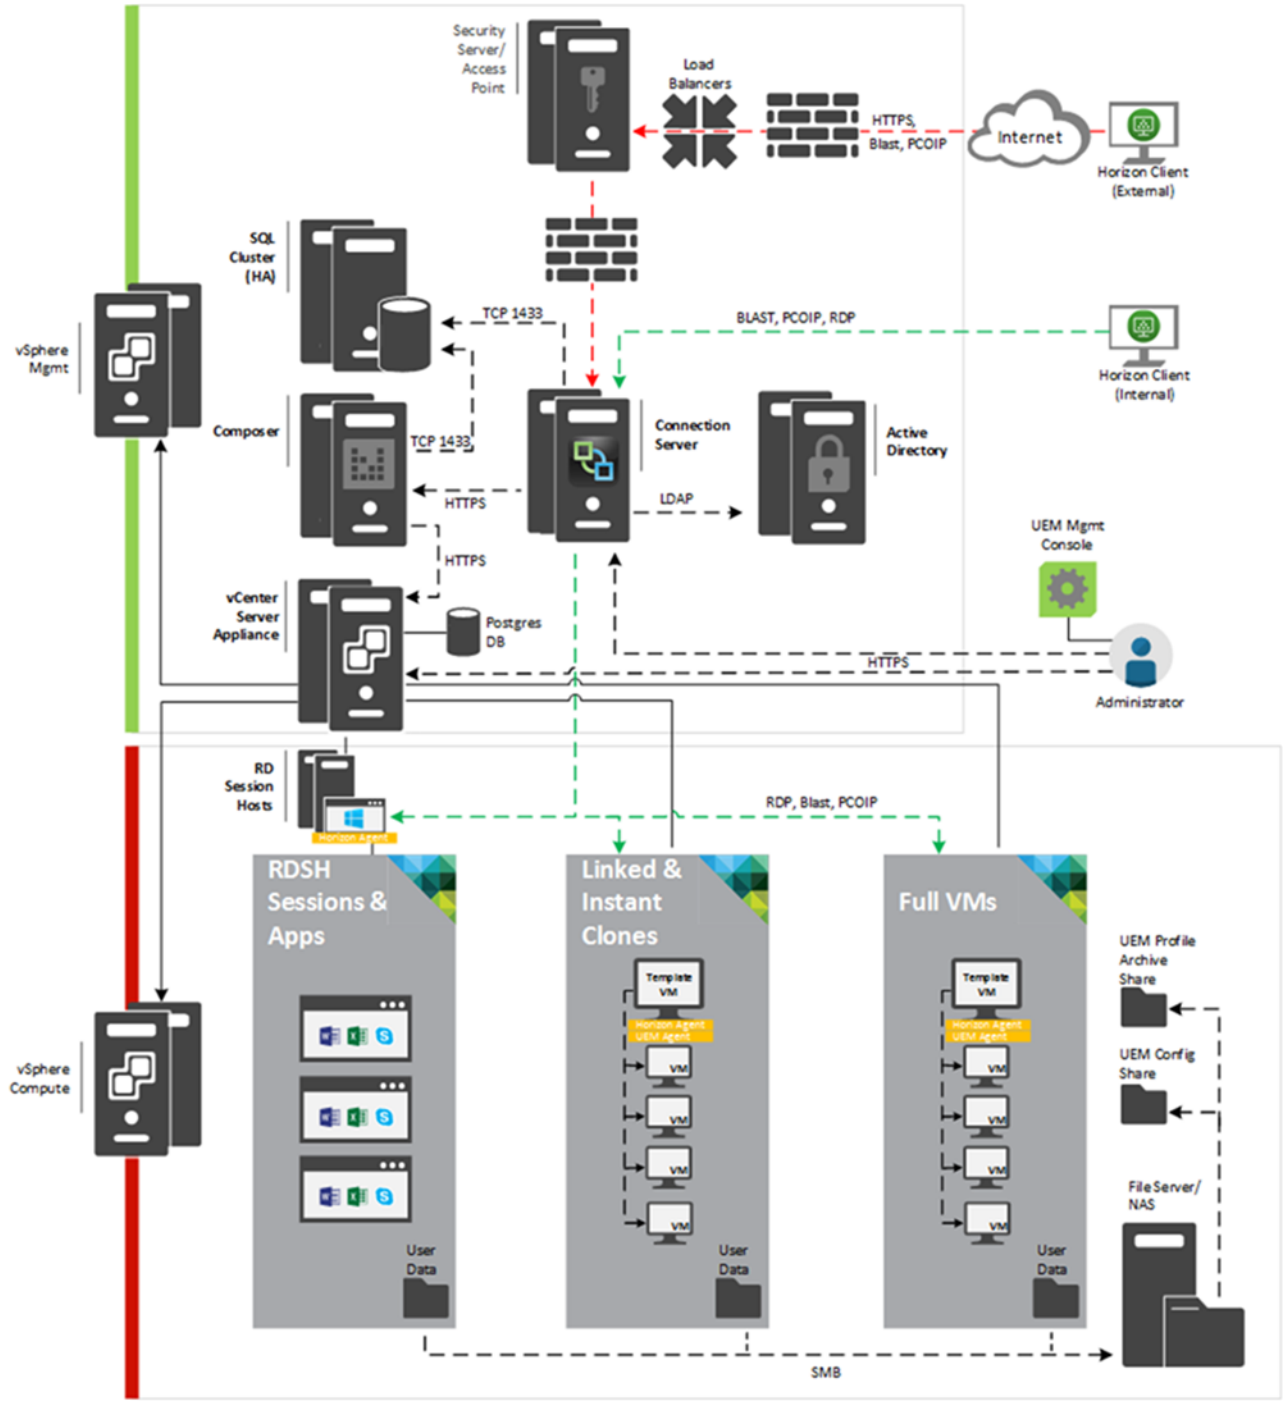

6 Customer-provided solution components

6.1 Customer-provided storage requirements

In the event that a customer wishes to provide their own storage array solution for a Dell Wyse Datacenter solution, the following minimum hardware requirements must be met:

Feature	Minimum Requirement	Notes
Total Tier 2 Storage Space (Mgmt VMs + User Data)	User count and workload dependent	1Gb/ 10Gb iSCSI or FC storage required on NL SAS disks minimally.
Tier 1 IOPS Requirement	(Total Users) x workload IOPS	6-30 IOPS per user may be required depending on workload. T1 storage should be capable of providing user IOPS requirement concurrently to all hosted users.
Tier 2 IOPS Requirement (Mgmt VMs + User Data)	(Total Users) x User Data IOPS	1 – 4 IOPS per user depending on user data needs. File share usage and size of deployment may shift this requirement.
Data Networking	10GbE Ethernet for LAN/T2 iSCSI 8Gb FC for T1/T2 FC	Data networking traffic should be isolated on dedicated NICs and HBAs in each applicable host.

6.2 Customer-provided switching requirements

Feature	Minimum Requirement	Notes
Switching Capacity	Line rate switch	10Gb switching pertinent to solution being implemented.
Fiber channel	8Gbps Enterprise-class	
10Gbps Ports	Uplink to Core	10Gbps Ports
VLAN Support	IEEE 802.1Q tagging and port-based VLAN support.	
Stacking Capability	Yes	The ability to stack switches into a consolidated management framework is preferred to minimize disruption and planning when up linking to core networks.

7 Solution Performance and Testing

At the time of publication, here are the available density recommendations per compute server. Please refer to Section 3.2 for hardware specifications

Table 1 User density summary

Host Config	Hypervisor	Broker & Provisioning	Workload	Template	User Density
R730	ESXi 6.5	Horizon 7 & Linked Clone	Task Worker	Windows 10 x64 & Office 2016	250
R730	ESXi 6.5	Horizon 7 & Linked Clone	Knowledge Worker	Windows 10 x64 & Office 2016	170
R730	ESXi 6.5	Horizon 7 & Linked Clone	Power Worker	Windows 10 x64 & Office 2016	140

All tests above were performed with LoginVSI version 4.1, on VMware Horizon 7 .For detailed validation results and analysis of these reference designs are in the following sections.

7.1 Test and performance analysis methodology

7.1.1 Testing process

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

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

At different times during testing, the testing team will complete some manual “User Experience” Testing while the environment is under load. This will involve a team member logging into a session during the run and completing tasks similar to the User Workload description. While this experience will be subjective, it will help provide a better understanding of the end user experience of the desktop sessions, particularly under high load, and ensure that the data gathered is reliable.

7.1.1.1 Load generation

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

desktops). The launchers and Login VSI environment are configured and managed by a centralized management console.

Additionally, the following login and boot paradigm is used:

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

7.1.1.2 Profiles and workloads

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

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

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

Profile to workload mapping

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

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

Login VSI Task Worker Workload

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

Login VSI Knowledge Worker Workload

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

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

Login VSI Power Worker Workload

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

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

Graphics - Login VSI Power Worker with ProLibrary workload

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

Graphics – LVSI Custom workload

This is a custom Login VSI workload specifically for higher performance, intensive graphics testing. For this workload, SPECwpc benchmark application is installed to the client VMs. During testing, a script is started that launches SPECwpc, which executes the Maya and sw-03 modules for high performance tests and

module sw-03 only for high density tests. The usual activities such as Office application execution are not performed with this workload. This type of workload is typically used for lower density/high performance pass-through, vGPU, and other dedicated, multi-user GPU configurations.

7.1.2 Resource monitoring

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

7.1.2.1 GPU resources

ESXi hosts

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

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

7.1.2.2 VMware vCenter

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

7.1.3 Resource utilization

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

Resource utilization thresholds

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

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

7.2 Test configuration details

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

Table 2 Hardware and software test components

Component	Description/Version
Hardware platform(s)	PowerEdge R730
Hypervisor(s)	ESXi 6.5
Broker technology	Horizon 7
Broker database	Microsoft SQL 2016
Management VM OS	Windows Server 2012 R2 (Connection Server & Database)
Virtual desktop OS	Windows 10 Enterprise
Office application suite	Office Professional 2016
Login VSI test suite	Version 4.1

7.2.1 Compute VM configurations

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

Desktop VM specifications

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

7.2.2 Platform Configuration

Please refer to Section 3.2 for hardware specifications

7.3 Test results and analysis

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

Platform Config	Hypervisor	Broker & Provisioning	Login VSI Workload	Density Per Host	Avg CPU	Avg Mem Consumed	Avg Mem Active	Avg IOPS / User
R730	ESXi 6.5	Horizon 7 & Linked Clone	Task Worker	250	85%	504GB	200GB	5
R730	ESXi 6.5	Horizon 7 & Linked Clone	Knowledge Worker	180	75%	504GB	200GB	8
R730	ESXi 6.5	Horizon 7 & Linked Clone	Power Worker	150	75%	500GB	210GB	12

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

Avg CPU: This is the average CPU usage over the steady state period. For clusters, this represents the combined average CPU usage of all compute hosts. On the latest Intel series processors, the ESXi host CPU metrics will exceed the rated 100% for the host if Turbo Boost is enabled (by default). An additional 35% of CPU is available from the Turbo Boost feature but this additional CPU headroom is not reflected in the VMware vSphere metrics where the performance data is gathered. Therefore, CPU usage for ESXi hosts is adjusted and a line indicating the potential performance headroom provided by Turbo boost is included in each CPU graph.

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

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

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

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

7.3.1 R730 Compute

Please refer to Section 3.2 for hardware specifications

7.3.1.1 Task Worker, 250 Users, ESXi 6.5, Horizon 7 linked clone

The below graph shows the CPU usage data for 250 user sessions on the host The CPU reaches a steady state average of 84% during the test when all 250 users are logged on.

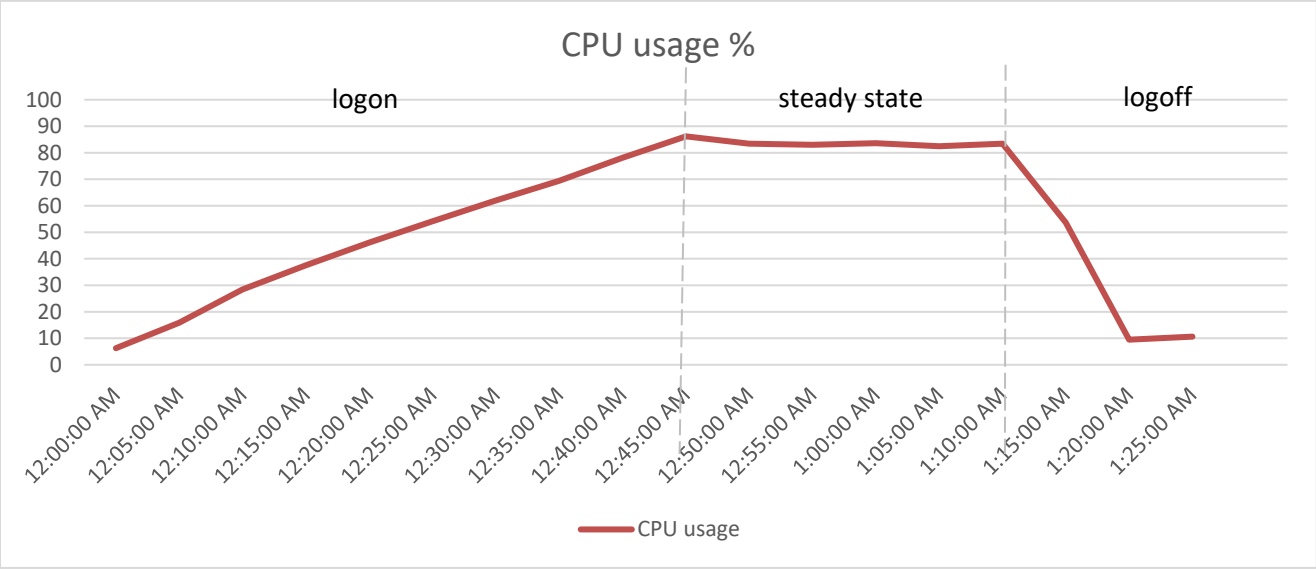

The VDI VMs consumed almost all the available physical memory on host, but active memory usage only reaching a max just above 200GB.

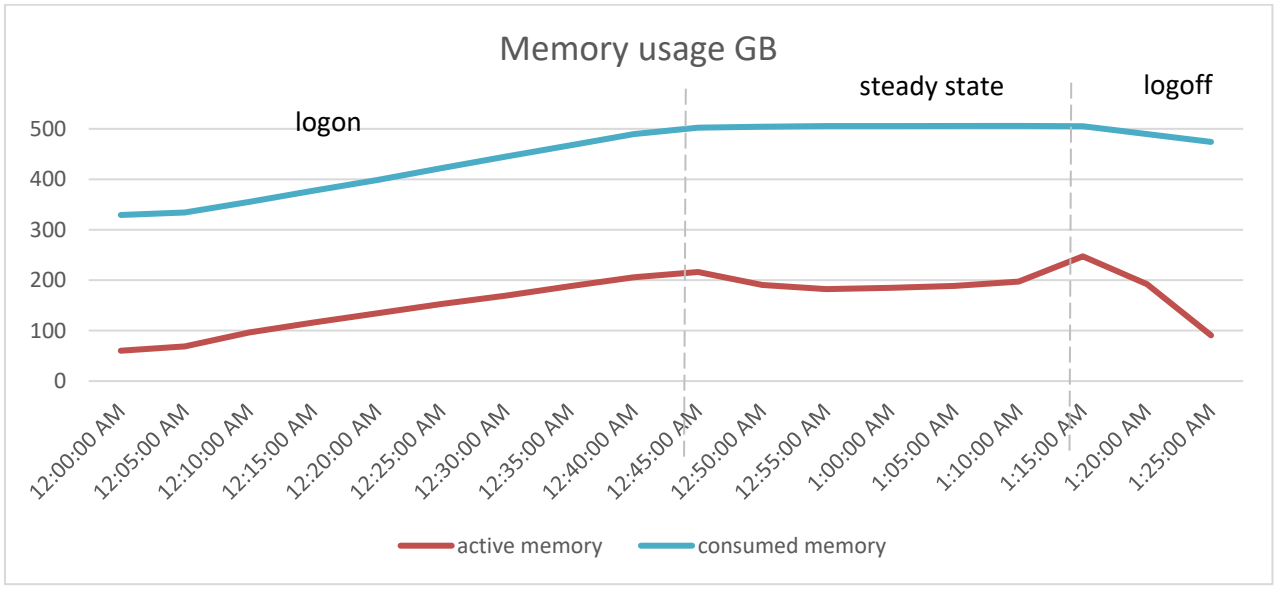

Network usage in the test runs with a peak at just over 1Mbps per user, less than 300Mbps in total on the compute host. The host has a pair of teamed 1 GB NICs, so the bandwidth is not under pressure of saturation.

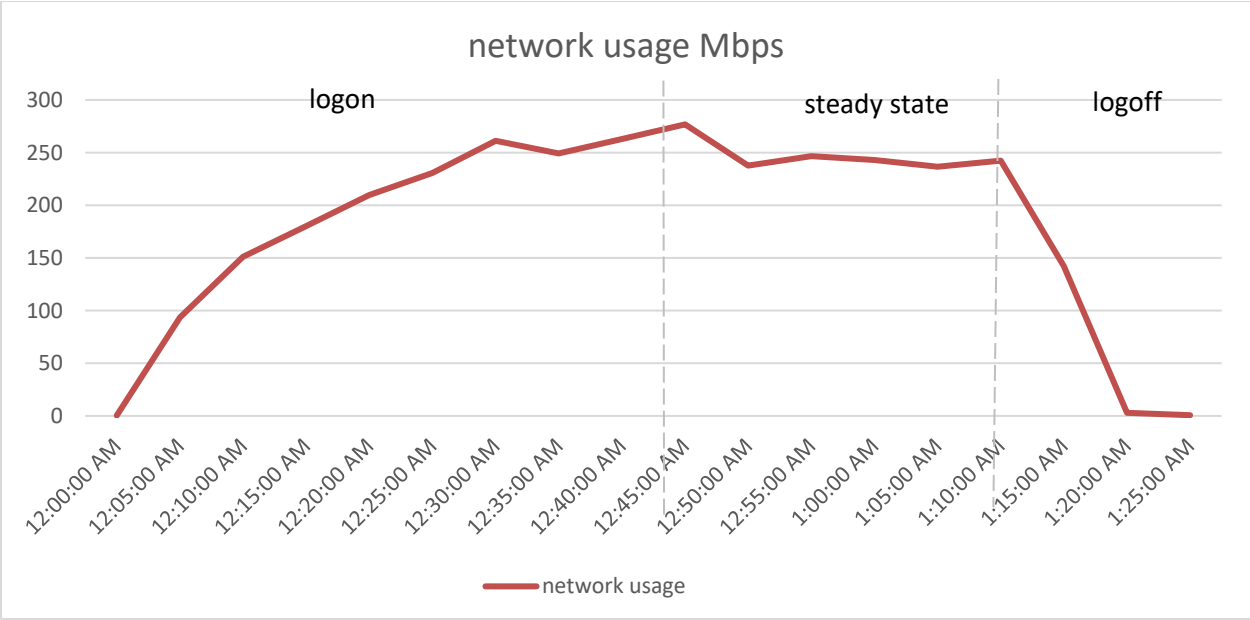

The data store IOPS is shown in below graph, the read IO is higher than write as expected and peaked when all users were logged on. Average IO Max around 6 IOPS and is in line with LoginVSI expectation. The read and write latency remained under 2ms.

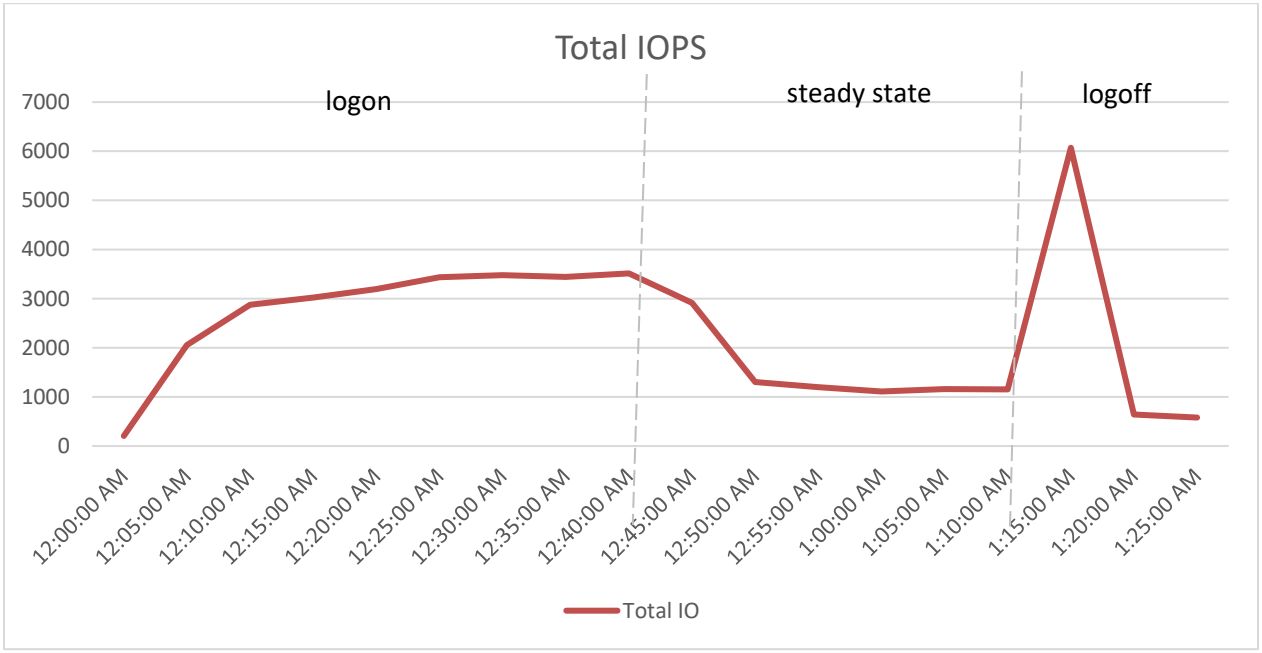

The Login VSI Max user experience score shown below for this test was not reached indicating there was little deterioration of user experience during testing and manually interacting with the test sessions backed this up, mouse and window responses were fast and video play back was of good quality.

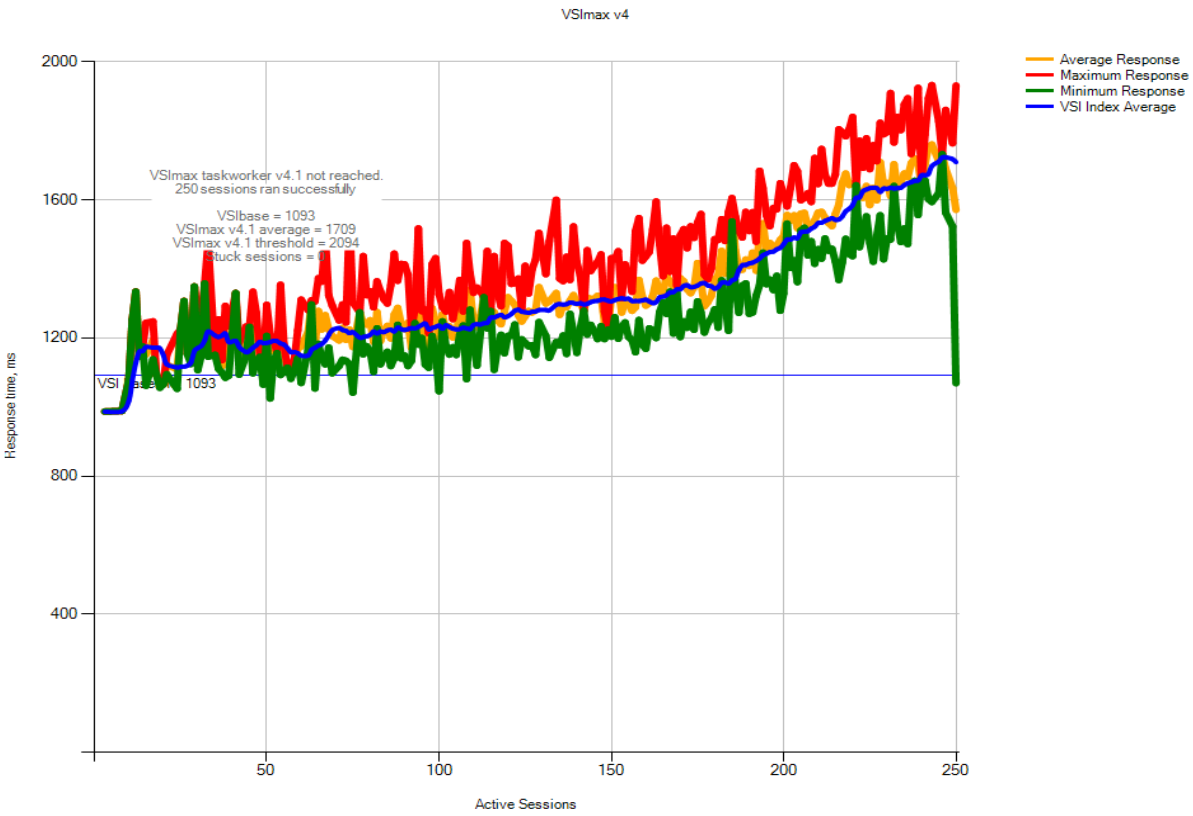

7.3.1.2 Knowledge Worker, 180 Users, ESXi 6.5, Horizon 7 linked clone

The below graph shows the CPU usage during the test. The CPU usage reaches just over 80% when all users logged on and settled just below 80% during steady state.

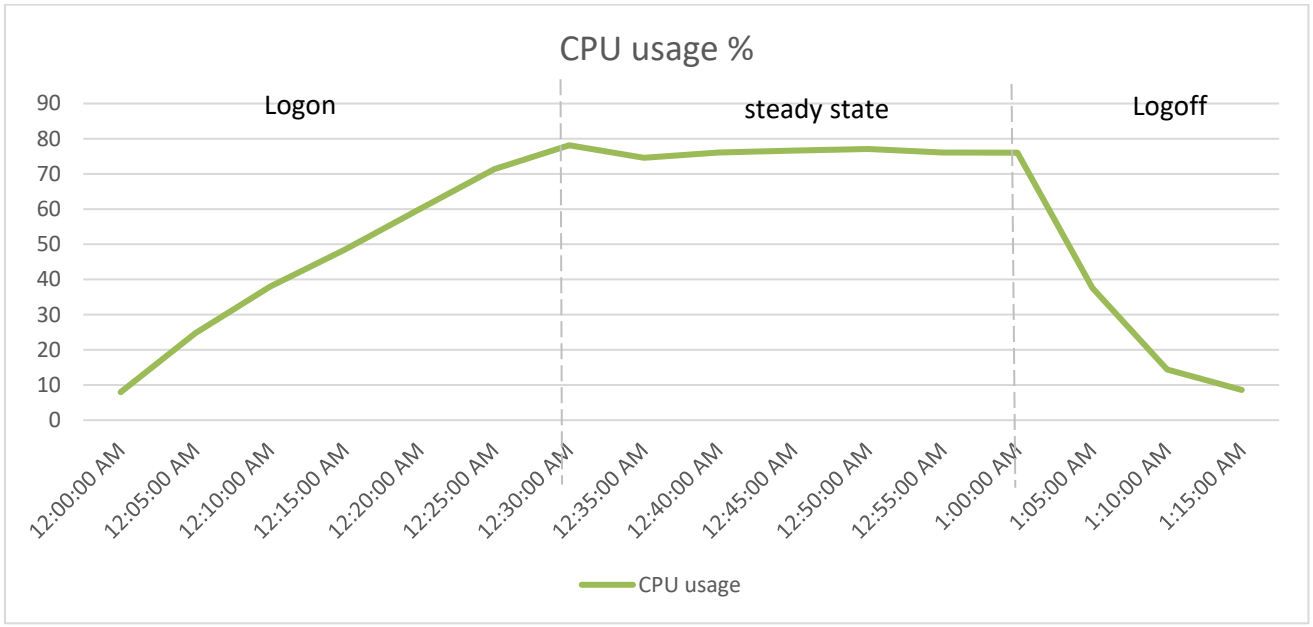

In terms of host memory utilization. Almost all 512GB available host memory was consumed.

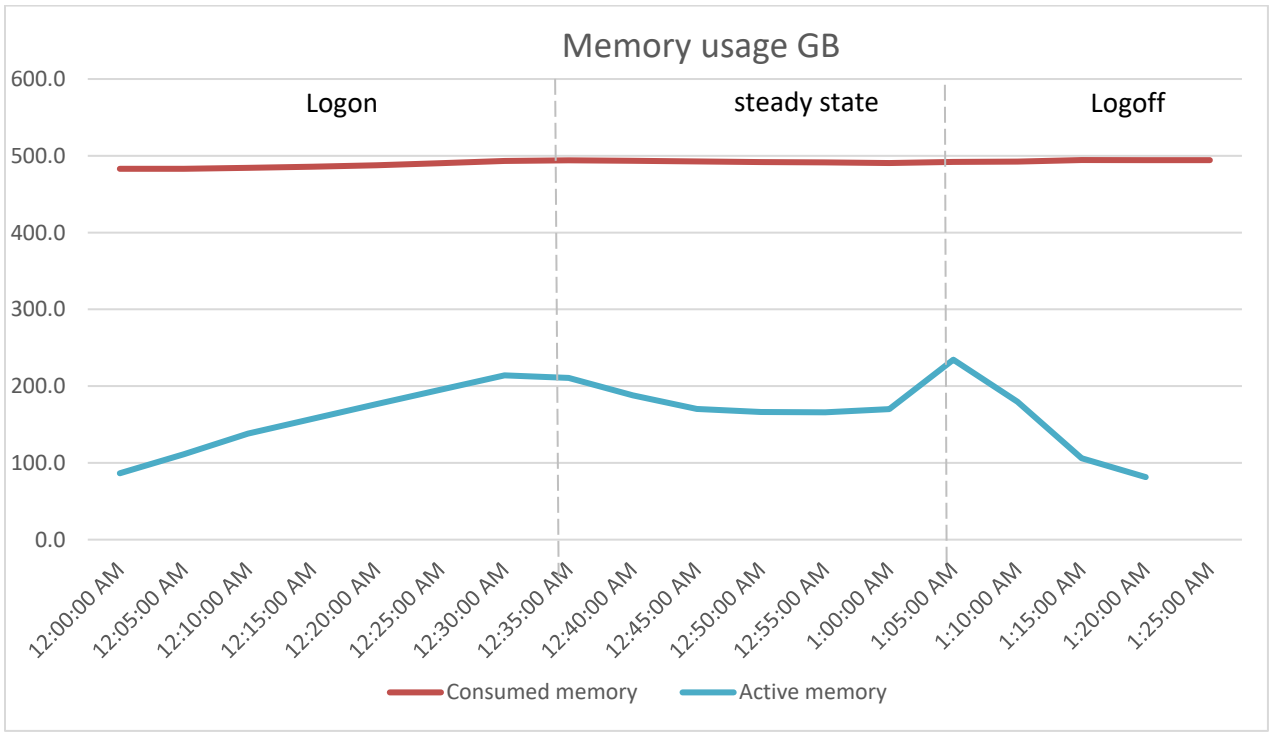

Network bandwidth is not an issue on this test run with a steady state peak of approximately 2Mbps per user which was just over 400Mbps overall on the host.

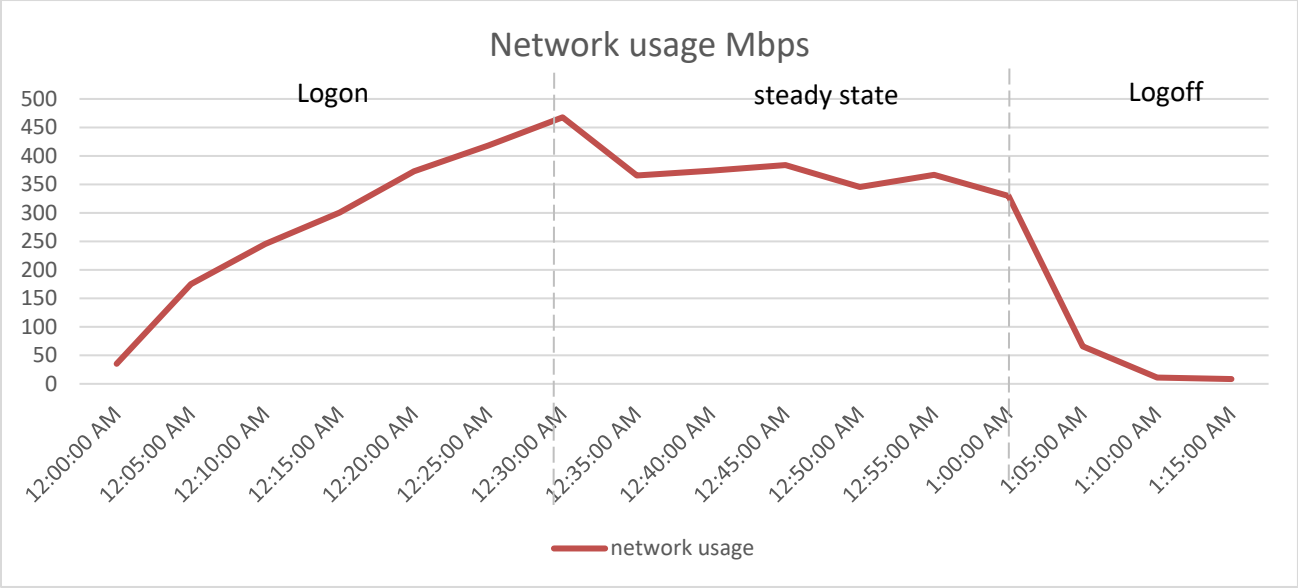

The graph below shows the datastore read and write IOPS on the host's local storage. The average IO per user maximum at 7.5 IOPS during steady. It is in line with LoginVSI's expectation.

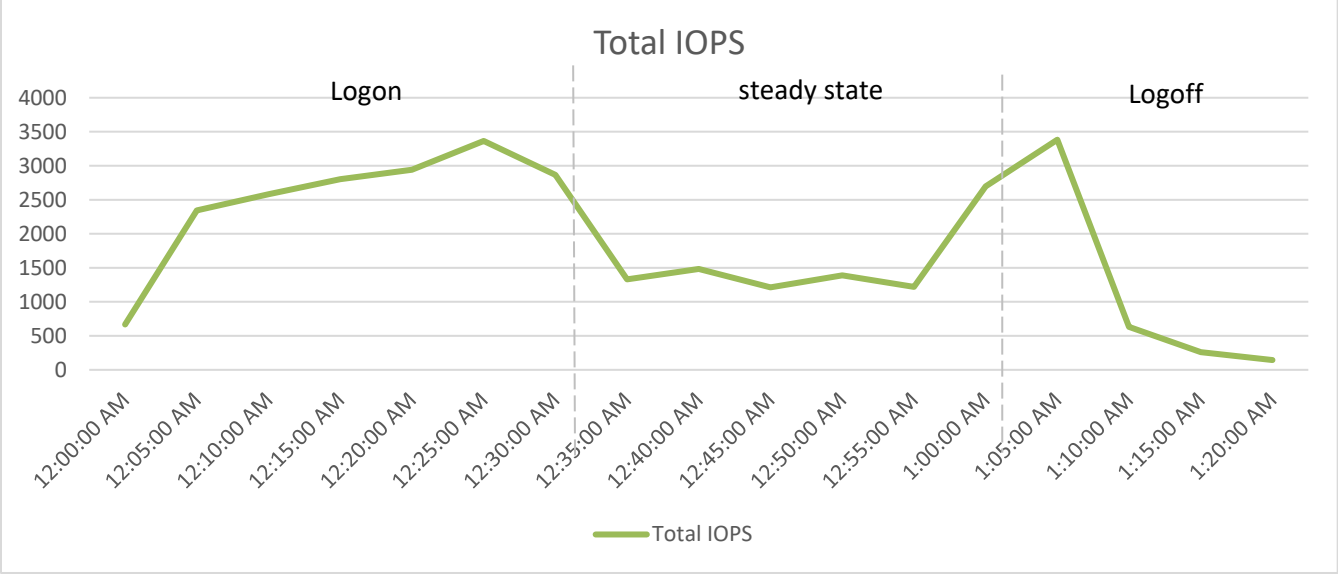

The Login VSI Max user experience score for this test was not reached indicating there was little deterioration of user experience during testing and manually interacting with the test sessions backed this up, mouse and window responses were fast and video play back was of good quality.

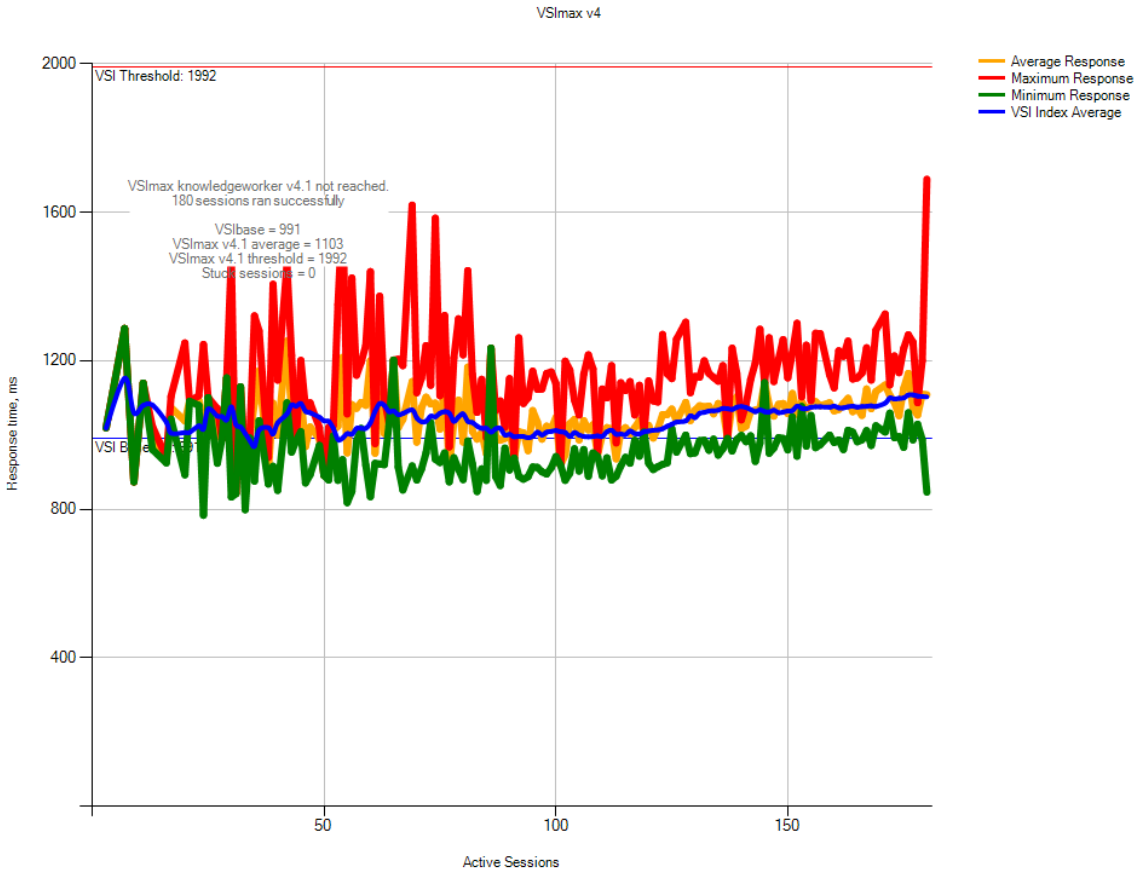

7.3.1.3 Power Worker, 150 Users, ESXi 6.5, Horizon 7 linked clone

The CPU usage peaked at around 85% utilization and this is in line with above task worker workload and knowledge worker workload tests.

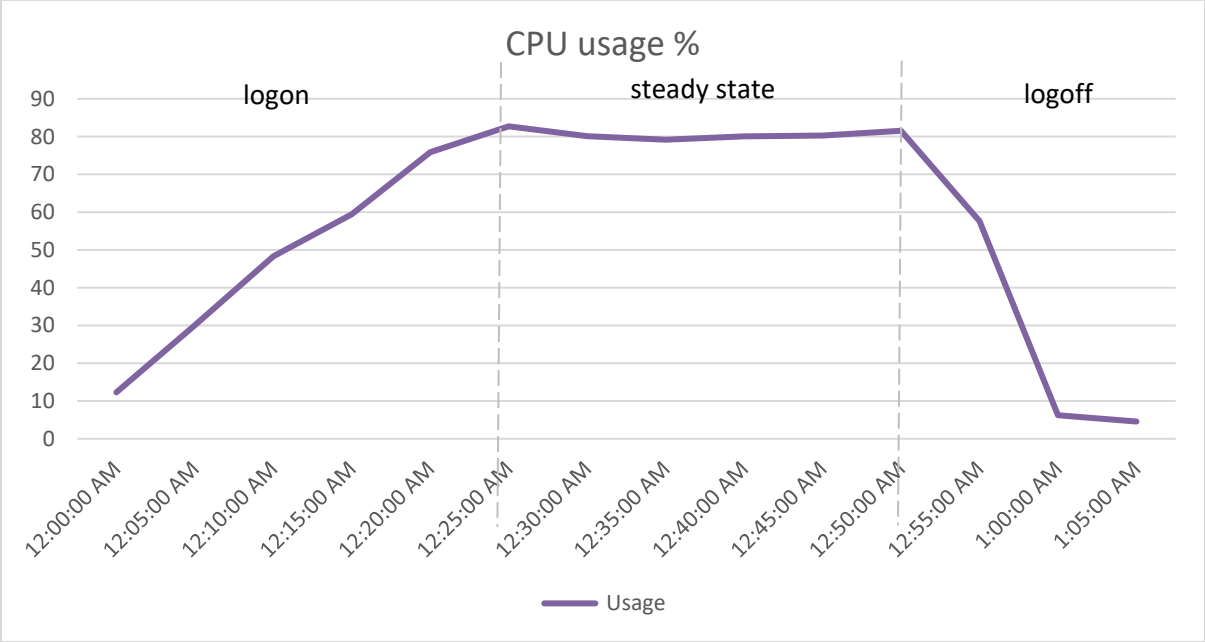

In regards to memory consumption for the host, out of a total of 512GB available memory, the VMs consumed almost all. However active memory usage is only slightly over half, peaked at 265GB. Like the other two workload tests, there was a small amount of memory ballooning at no more than 5GB. There was no memory swapping occurring at any time during tests.

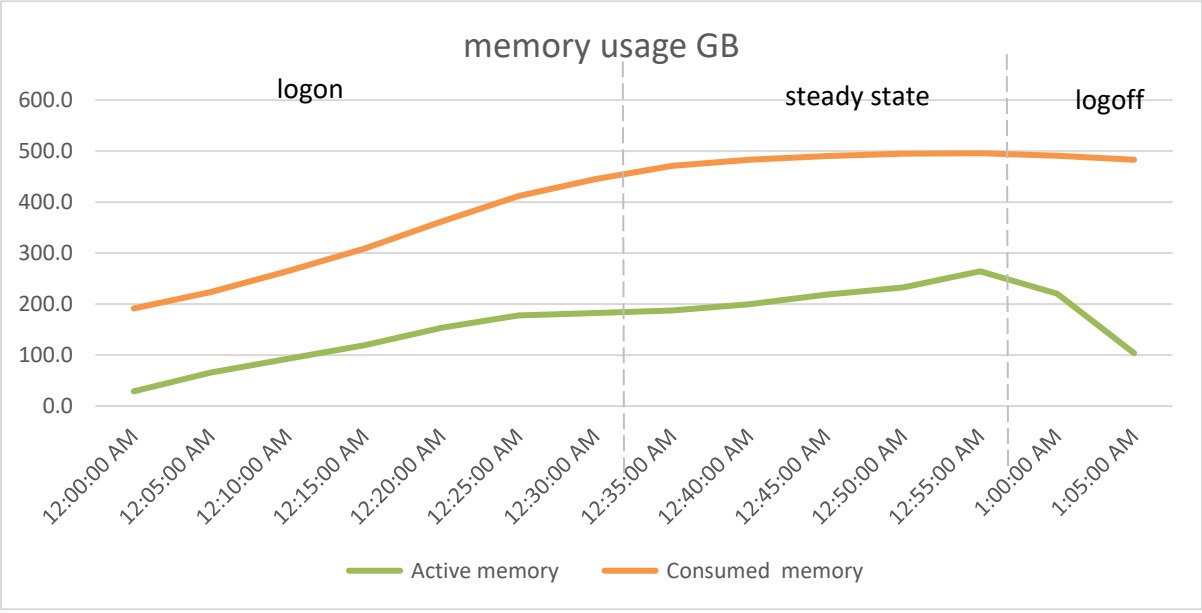

Network usage of the Power worker workload tests showing heavier usage than other two workloads. Peaked at around 4.5 Mbps per user, and under 500Mbps on the host. It is still well under the 1Gbps NIC capacity.

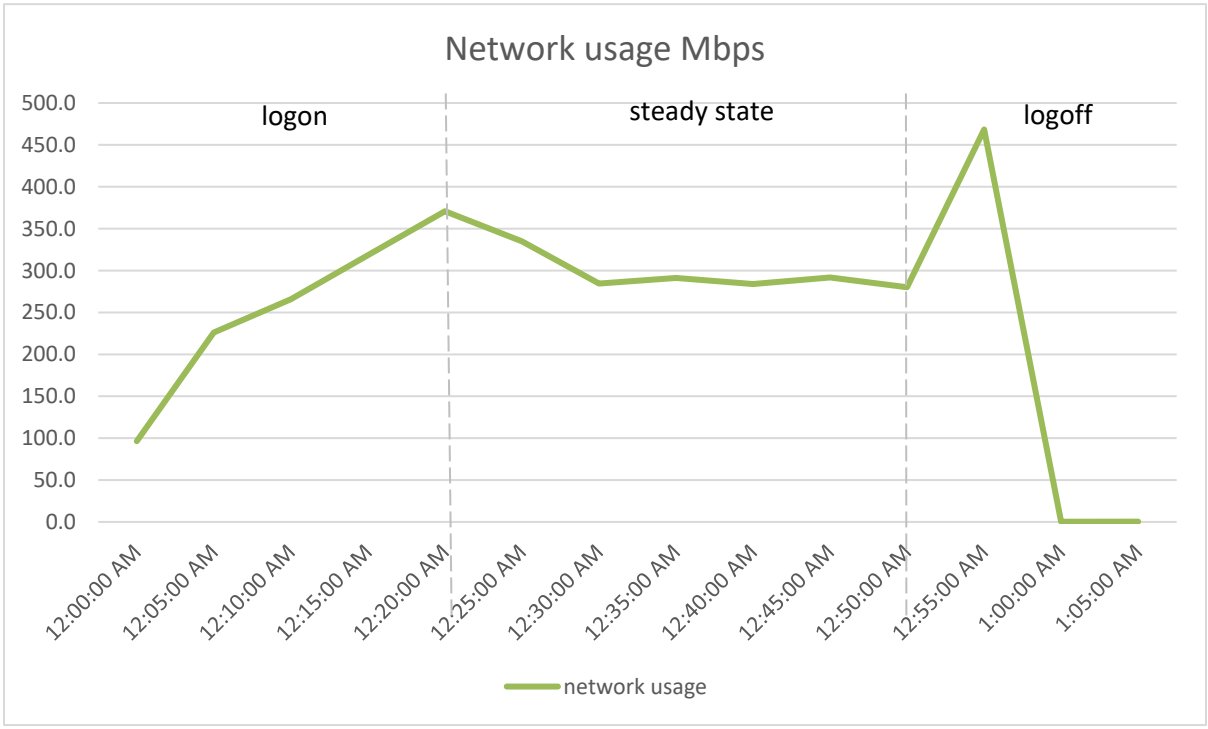

The datastore IOPS reached maximum total IOPS around 4,500, during the testing period, the data store latency remains under 2ms. Well below the 20ms latency threshold. Average IO per user is approx. 12 per user during steady state and this is in line with LoginVSI's recommendation

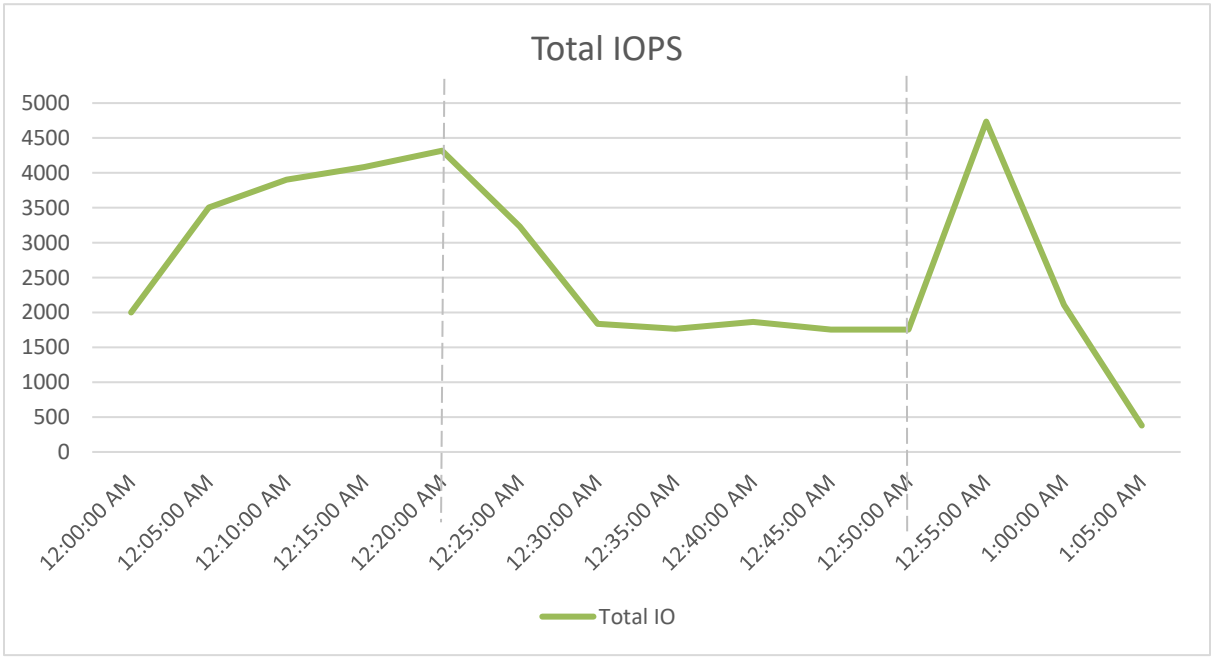

The Login VSI Max user experience score for this test was not reached indicating there was little deterioration of user experience during testing and manually interacting with the test sessions backed this up, mouse and window responses were fast and video play back was of good quality.

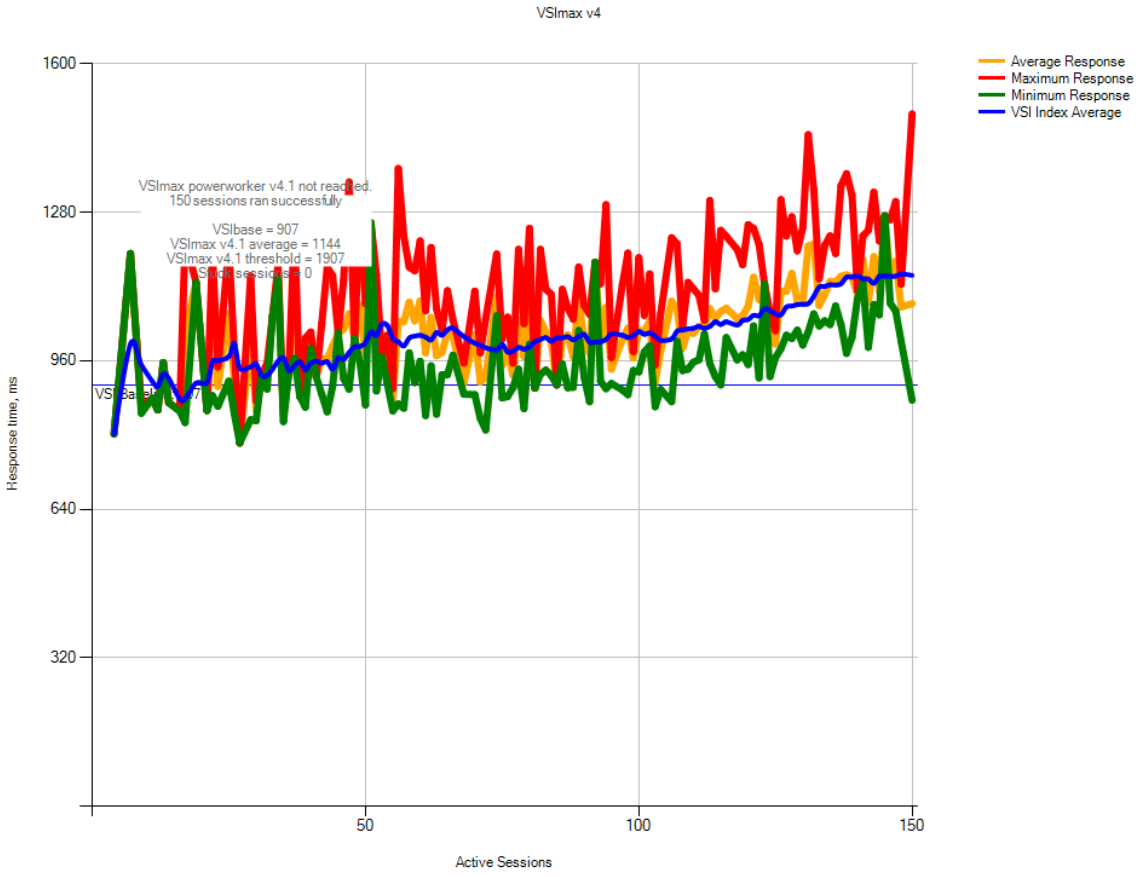

7.3.2 Shared Storage

7.3.2.1 XtremIO

At the time of publication, here are the available shared storage maximum density recommendations.

User Density	T1 Storage	T2 Storage	Workload	Template OS
<= 500	1 x XtremIO Starter X-Brick	1 x VSA using T1 X-Brick	All	All
501 – 1,500	1 x XtremIO Starter X-Brick	1 x Unity 300		
501 – 3,000	1 x XtremIO X-Brick			
3,000 – 6,000	2 x XtremIO X-Bricks	2 x Unity 300		
6,001 – 9,000	3 x XtremIO X-Bricks	3 x Unity 300		
9,001, 10,000[i]	4 x XtremIO X-Bricks	4 x Unity 300		

For detailed up-to-date validation results and analysis of the XtremIO reference designs and more, please visit: LINK

Tier 2 storage sizing is based upon 5GB allocated per user for profile and user data. Recommendations are based upon IOPS, not capacity. 6 x 400GB SSD drives were used in a RAID 1 configured as FAST Cache in combination of 48 x 2TB NL-SAS drives in a RAID 6 configuration. Each Unity 300 supports 3,000 Knowledge Workers' profile and user data, as well as the management VM. If additional space is required, the following table can help in determining the correct drive types to select for a given capacity.

Unity 300 Drive Size	Maximum User Profile & Data Per User @ 3,000 Users Per Array
2TB NL-SAS RAID 6	68GB
4TB NL-SAS RAID 6	136GB
6TB NL-SAS RAID 6	204GB

Please refer to the Sizing Advisor for additional sizing guidance.

7.3.2.2 Unity VSA – Combined Tier 1 and Tier 2 Storage

The Unity VSA supports up to 500 users on an XtremIO Starter X-Brick with the observed IOPS and disk usage observed during testing different Login VSI workloads as follows:

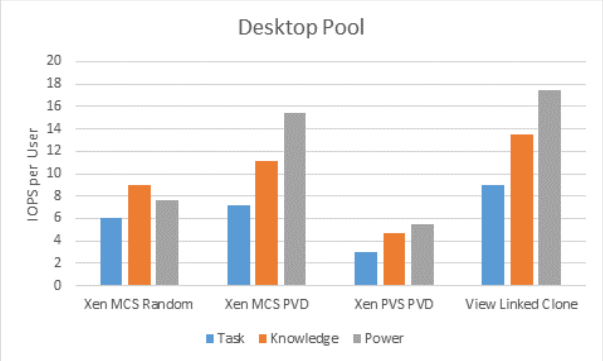

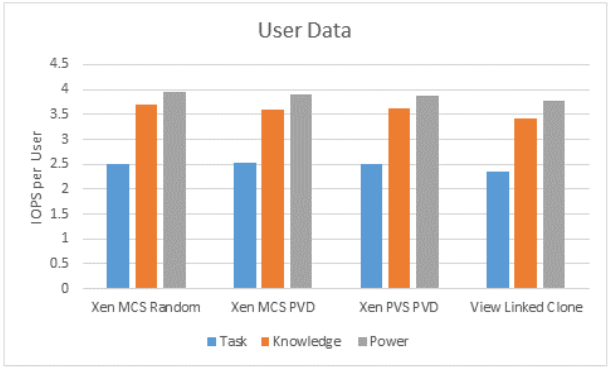

If your workload requirements vary significantly in either IOPS or capacity, we recommend placing T2 data on a discrete array. Our synthetic workload resulted in an overall data reduction rate of 3.7:1, as observed below:

UNITY VSA WITH XTREMIO FOR COMBINED T1/T2

DATA REDUCTION RATIO

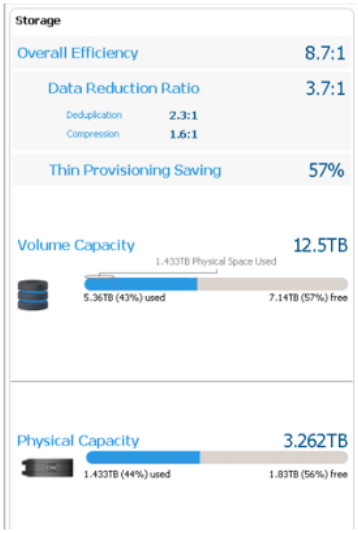

- Observed data reduction ratio across Login VSI workloads was 3.7:1
- 3.262TB physical capacity can support 3.262 * 3.7 = 12.0694TB data
- About 24GB capacity consumed (including desktop and user data) for each user in this 500-user test

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 David Hartman, a Consulting Solutions Engineer in the Global Solutions Engineering Group at Dell EMC, for contributing Dell EMC storage information and scaling guidance. His team produces EUC solutions across the Dell EMC product stack based upon Citrix XenDesktop and VMware Horizon. David is responsible for the creation, execution, and delivery of EUC Reference Architectures, Solution Guides, and legacy Design and Implementation Guides.

Thanks to Yinglong Jiang and Cormac Woods for their help in the validation of the solution in this RA.

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.

Cormac Woods is a Sr. Systems Engineer in the Desktop Virtualization solutions Group at Dell. Cormac has over 20 years of experience in IT and Telecoms product engineering. In addition, Cormac has a deep background in IT systems support in financial and manufacturing environments. In his current role, he has developed many of the solutions provided by the Dell Wyse Datacenter team using VMware Horizon as well as other VDI software environments on multiple Dell server platforms.

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

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

About the Authors

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

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