Wyse Datacenter for Citrix XenDesktop Reference Architecture

A Reference Architecture for the design, configuration and implementation of a Citrix XenDesktop environment.

Dell Engineering
April 2017

A Dell Reference Architecture

Revisions

Date	Description
April 2014	Initial release
March 2016	Document overhaul, new Broadwell CPUs, networking, servers
September 2016	Minor updates to Endpoints section
December 2016	Updated Citrix software section with latest XenDesktop 7.12 info and other minor updates
January 2017	Replaced Compellent and EqualLogic storage recommendations with Dell EMC storage
April 2017	Added test results to document

Contents

1 Introduction

1.1 Purpose

This document addresses the architecture design, configuration and implementation considerations for the key components required to deliver virtual desktops or shared sessions via Citrix® XenDesktop® and XenApp® on VMware® vSphere® 6.x or Microsoft® Windows Server® Hyper-V® (2016 or 2012 R2). Proposed design choices include rack or blade servers, local disks or shared storage. Guidance contained within this document follows a building block methodology enabling the combination of several different components each with their own scaling capabilities.

1.2 Scope

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

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

1.3 What's New

- Incorporated test results into document.
- Updated specifications for Windows 2016 Hyper-V support.

1.4 External Publications & Industry Vertical Information

- Dell Wyse Datacenter for Mobile Clinical Computing: LINK
- Dell Wyse Datacenter for Government: LINK
- Dell Wyse Datacenter for Education: LINK

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 XenApp 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 Tier 1 (LT1) and Shared Tier 1 (ST1). “Tier 1” in the Dell Wyse Datacenter context defines from which disk source the VDI sessions execute. LT1 includes rack servers or blades with SSDs while ST1 can include rack or blade servers due to the usage of ST1 storage. Tier 2 (T2) 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 T2 storage for all solution models. Dell Wyse Datacenter is a 100% virtualized solution architecture.

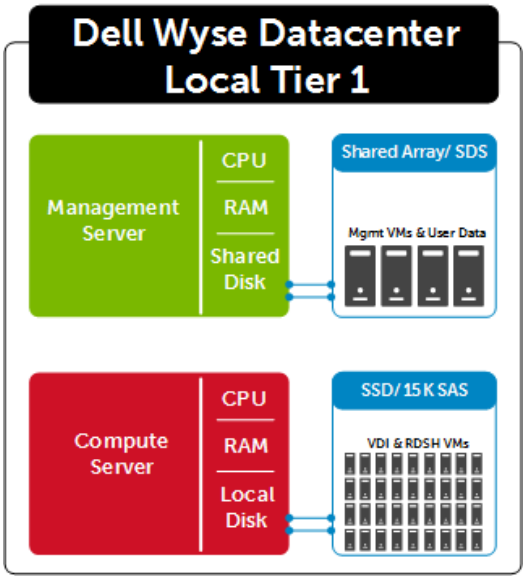

In the Shared Tier 1 solution model, an additional high-performance shared storage array is added to handle the execution of the VDI sessions. All compute and management layer hosts in this model are diskless where possible or utilize two mirrored disks if required for the operating system.

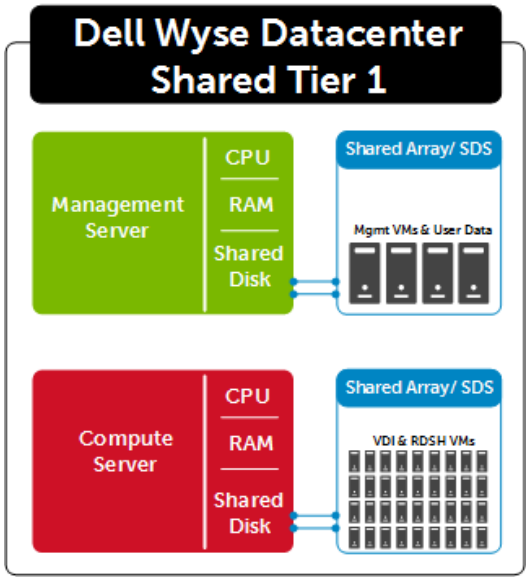

2.3 Solution Layers

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 the Dell Wyse Endpoints section)

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 S-Series 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 S-Series 48-port switch for 10Gb LAN traffic along with a single Brocade fiber channel switch for SAN connectivity. Additional switches are added and stacked as required to provide High Availability for the Network layer.

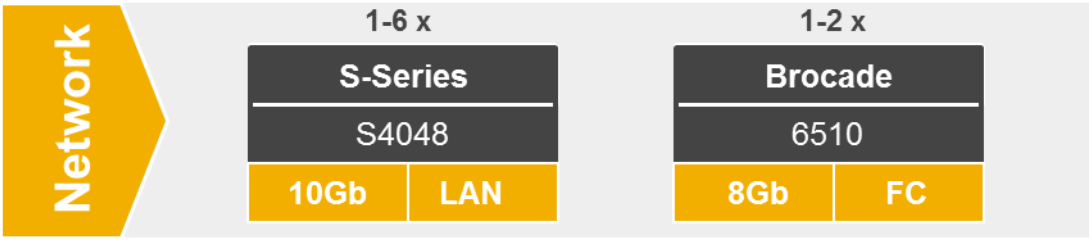

2.3.2 Compute

The compute layer consists of the server resources responsible for hosting the XenDesktop or XenApp user sessions with their respective VMs hosted either via VMware vSphere or Microsoft Hyper-V hypervisors, local or shared Tier 1 solution models (local Tier 1, all-flash, pictured below).

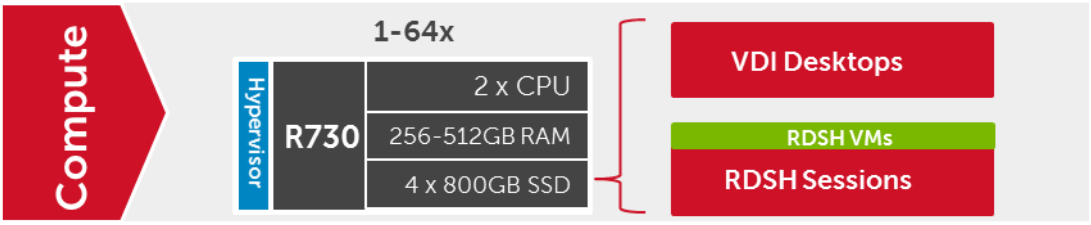

2.3.3 Management

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

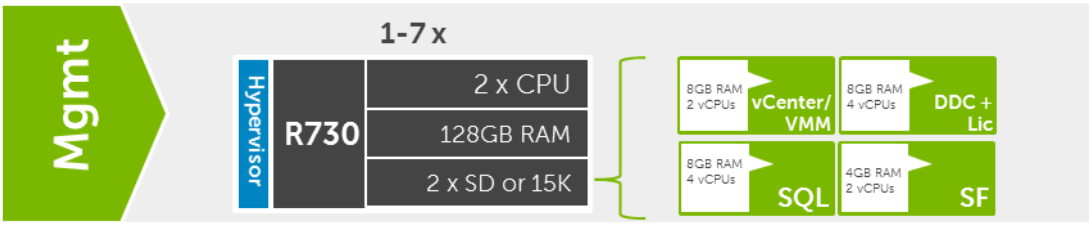

2.3.4 Storage

The storage layer consists of options provided by Dell EMC XtremIO arrays for combined shared T1, T2, and file storage (optionally and only for 500 users or less) 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. These configurations support 3,000 knowledge worker users with 5GB of user data (disk configurations can be adjusted to support large amounts of user data).

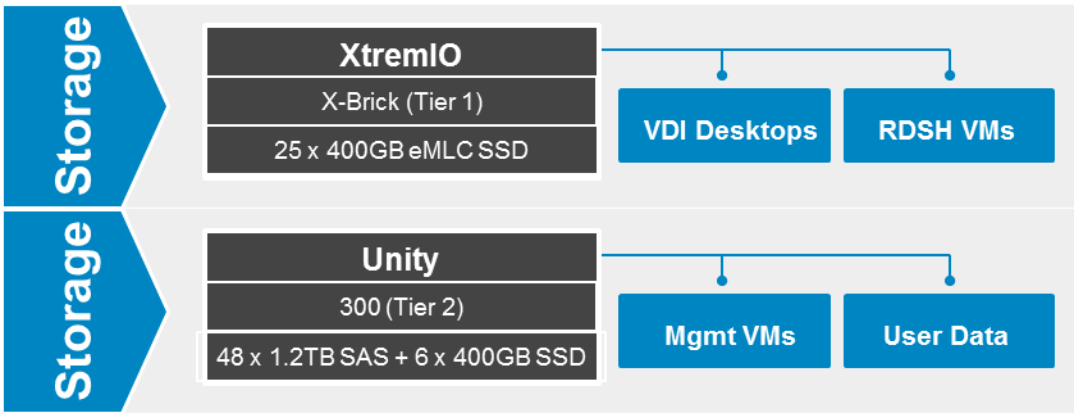

2.4 Local Tier 1

2.4.1 Base Offering

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.

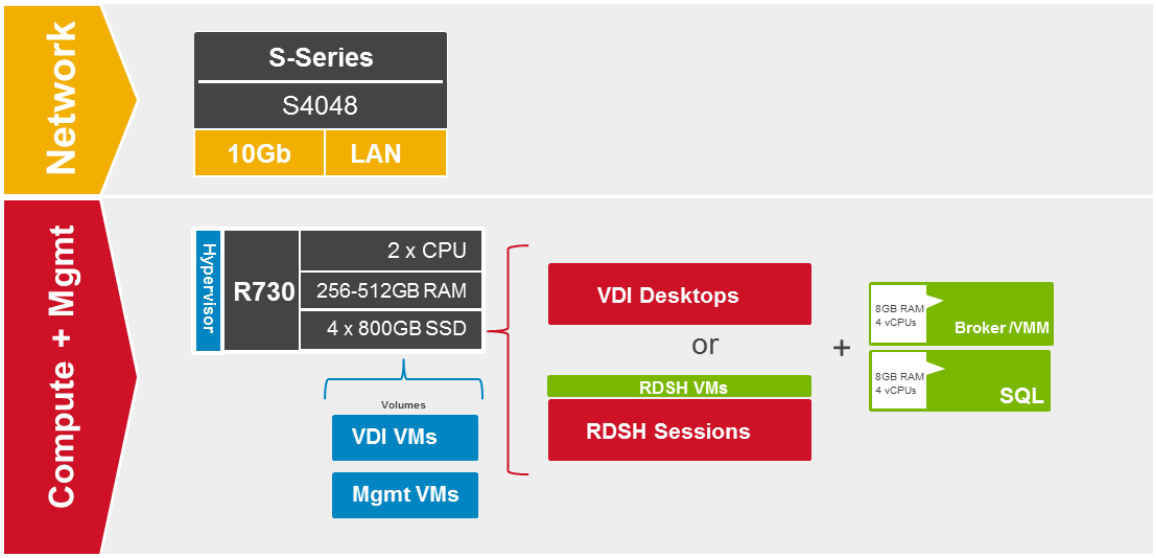

Note: 150 user density based on Task Worker workload.

2.4.2 Local Tier 1 for Rack Servers

The Local Tier 1 solution model provides a high-performance scalable rack-based configuration that hosts user VDI sessions on local SSD or spinning disk in the compute layer. vSphere or Hyper-V based solutions are available and scale based on the chosen hypervisor. All-flash pictured below, if spinning disk is desired substitute the SSDs with 10x 600GB 15K SAS HDDs for vSphere or 12 HDDs for Hyper-V. A Unity 300 array is added to provide shared storage (T2) for the management VMs as well as for user data.

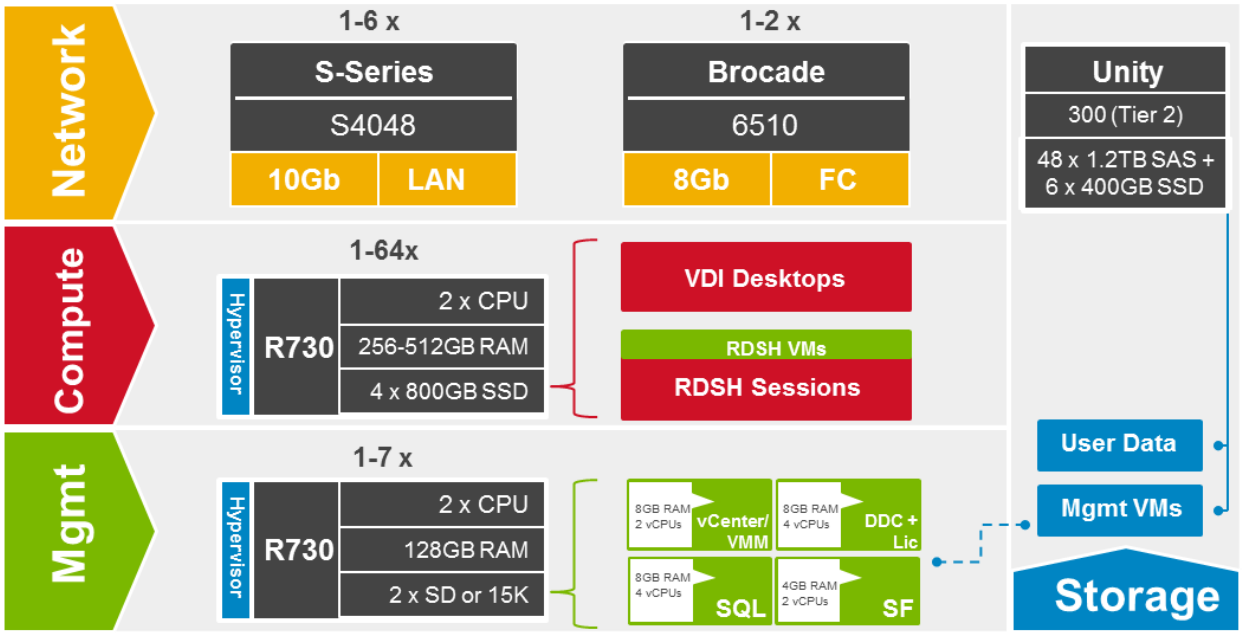

High-performance graphics capabilities compliment the solution and can easily be added at any time to any new or existing deployment. The Dell Precision Appliance for Wyse is readily orderable as a graphics-enabled compute host (ESXi hypervisor only). Simply add the appropriate number of appliances to your architecture and provide a superior user experience with NVIDIA GRID vGPU technology.

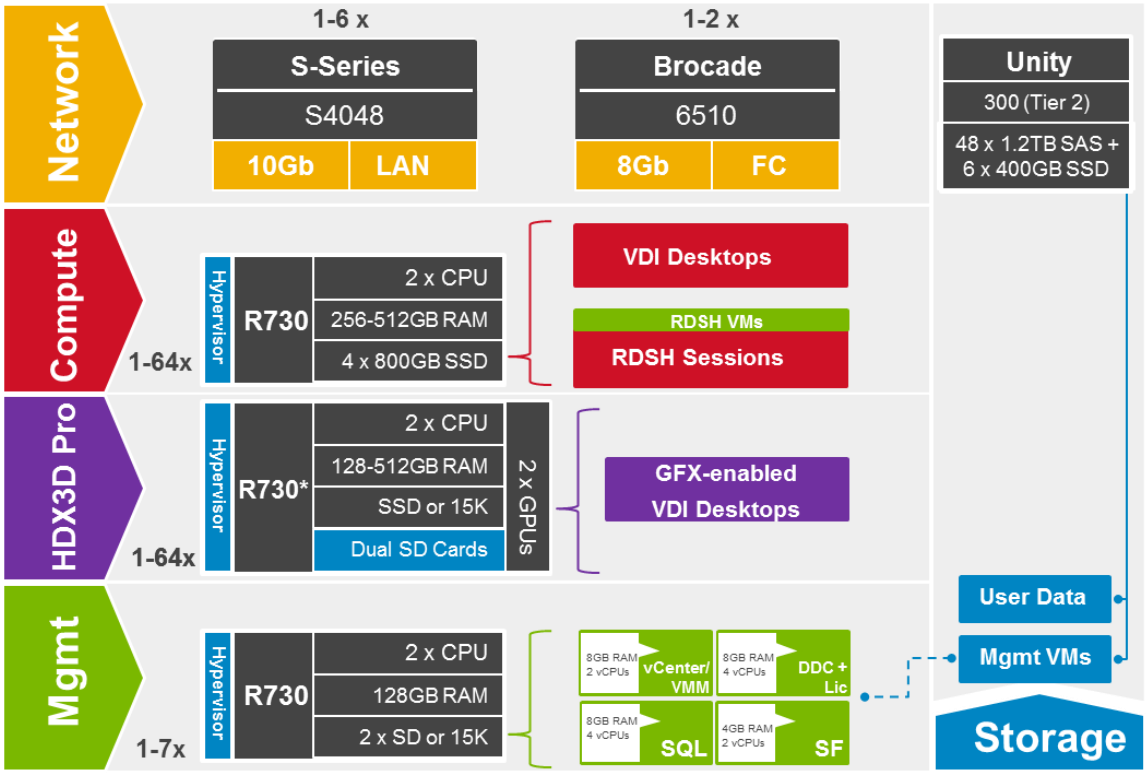

* The Dell Precision Appliance for Wyse is a R730 server based solution

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.

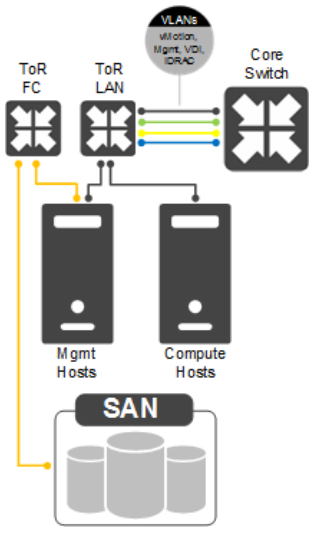

2.4.2.2 Local Tier 1 – Rack Cabling (HA)

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

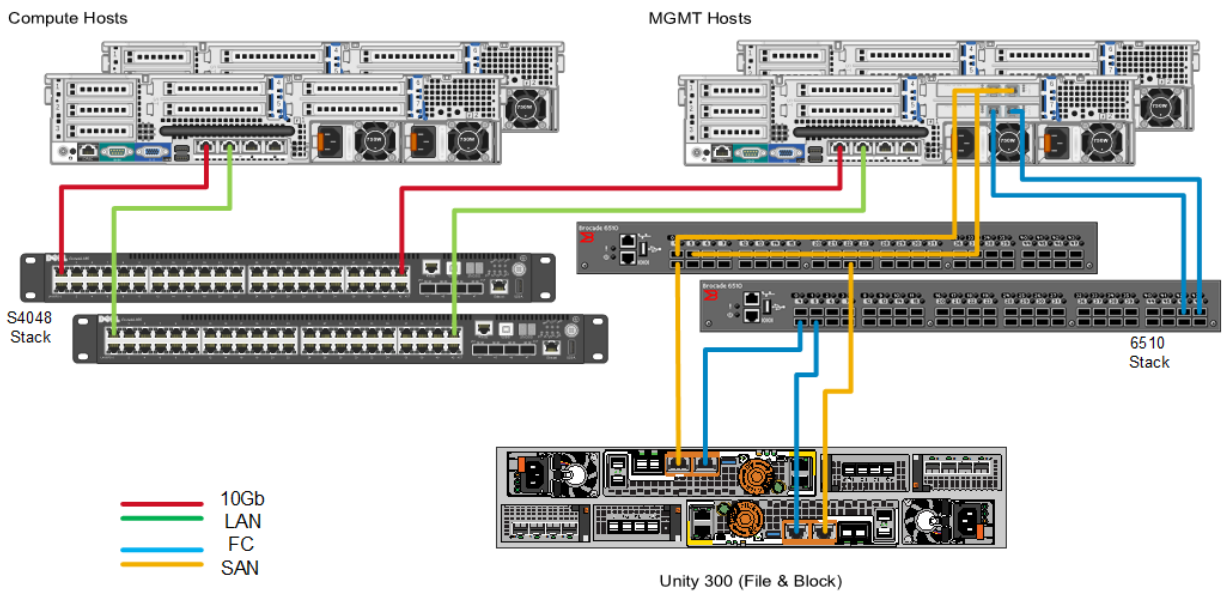

Refer to the Unity 300 Tier 2 section for a diagram of disk enclosure connections via mini-SAS HD.

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 as

well as storage for user data. User VDI sessions are hosted locally on SSDs in each blade server using Citrix MCS for desktop delivery. A pair of PowerEdge M I/O Aggregator switches are 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.

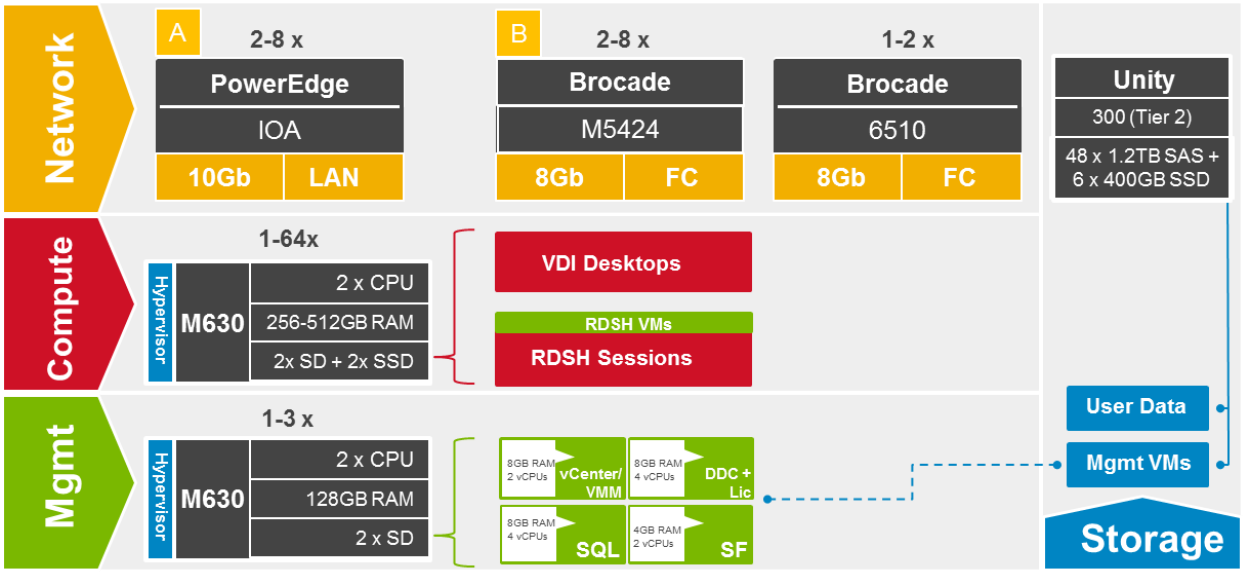

Note: Hypervisor installation to SD cards is recommended for vSphere but not supported for Hyper-V.

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 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.

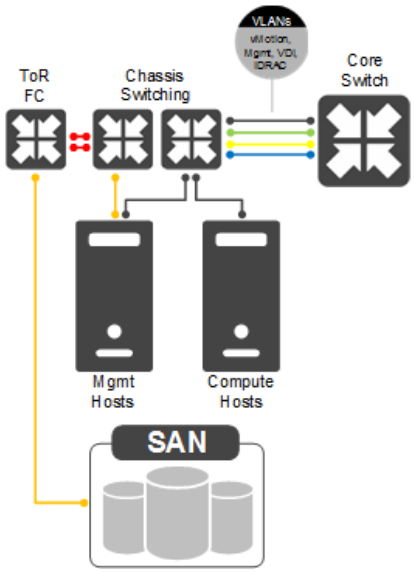

2.4.3.2 Local Tier 1 – Blade Chassis Cabling (HA)

The following diagram depicts the LT1 blade solution including optional components for HA. The A Fabric, B Fabric and ToR switches are stacked, respectively.

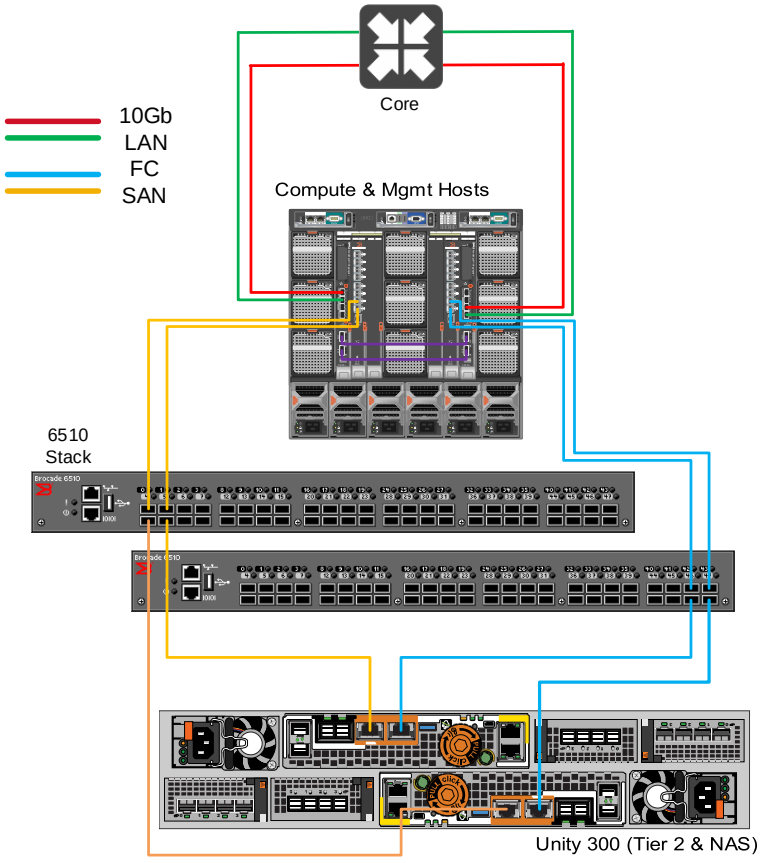

Refer to the Unity 300 Tier 2 section for a diagram of disk enclosure connections via mini-SAS HD.

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. Please refer to the Storage section for detailed information.

2.5.1 Shared Tier 1 for Rack Servers

500 Users or Less

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

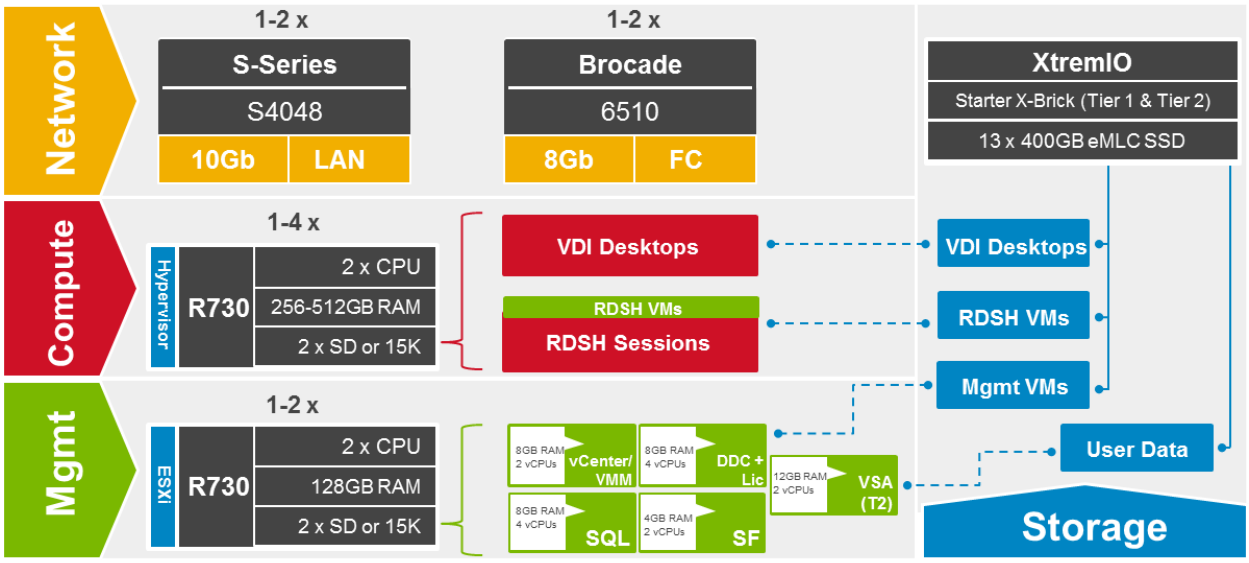

Note: The Unity VSA VM must be deployed on vSphere.

As with the Local Tier 1 for Rack Servers architecture, the Dell Precision Appliance for Wyse can be added in as a graphics-enabled compute host in a Shared Tier 1 architecture as well.

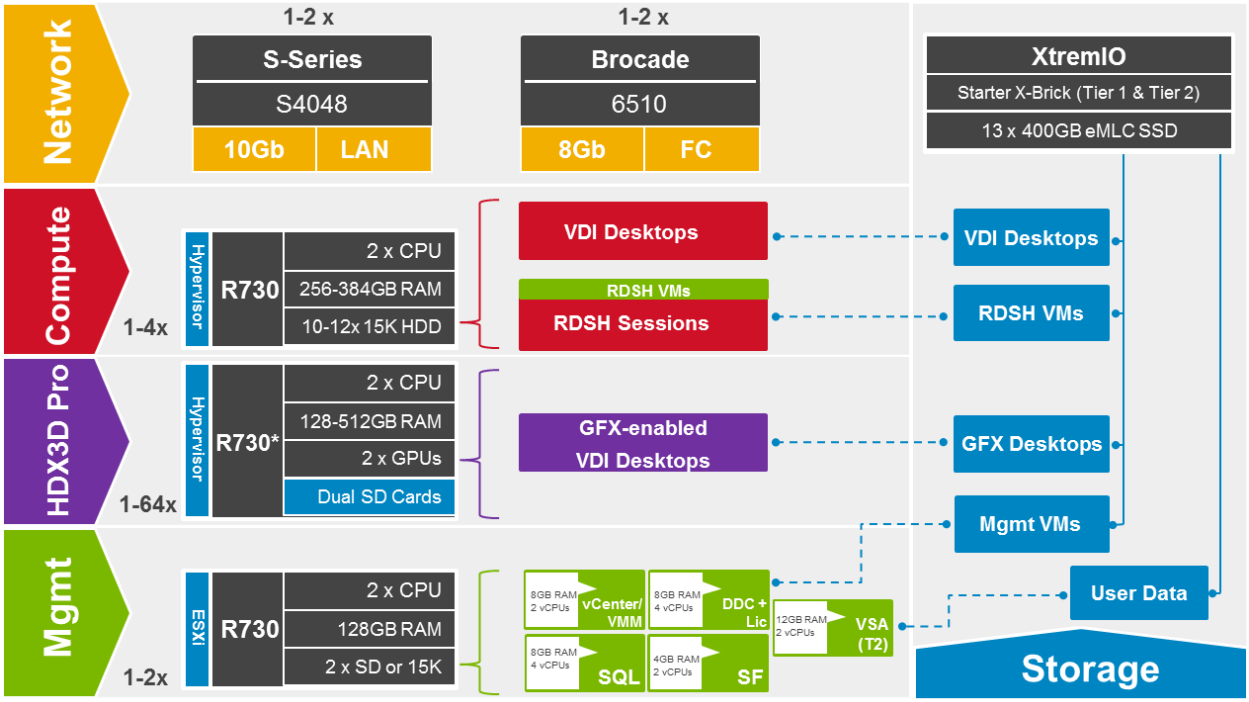

* The Dell Precision Appliance for Wyse is a R730 server based solution

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 and Tier 2 as the user count grows. Additional arrays and/or larger disk sizes can also be added to Tier 2 if necessary due to larger user data capacity needs.

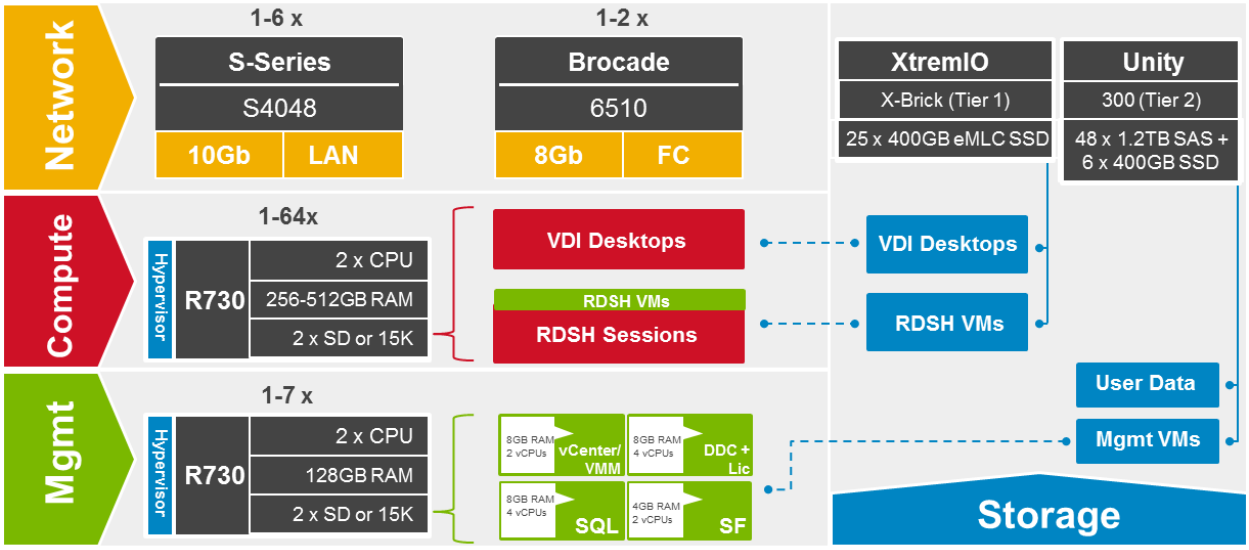

The diagram below shows the discrete T1 and T2 arrays with a graphics-enabled compute configuration.

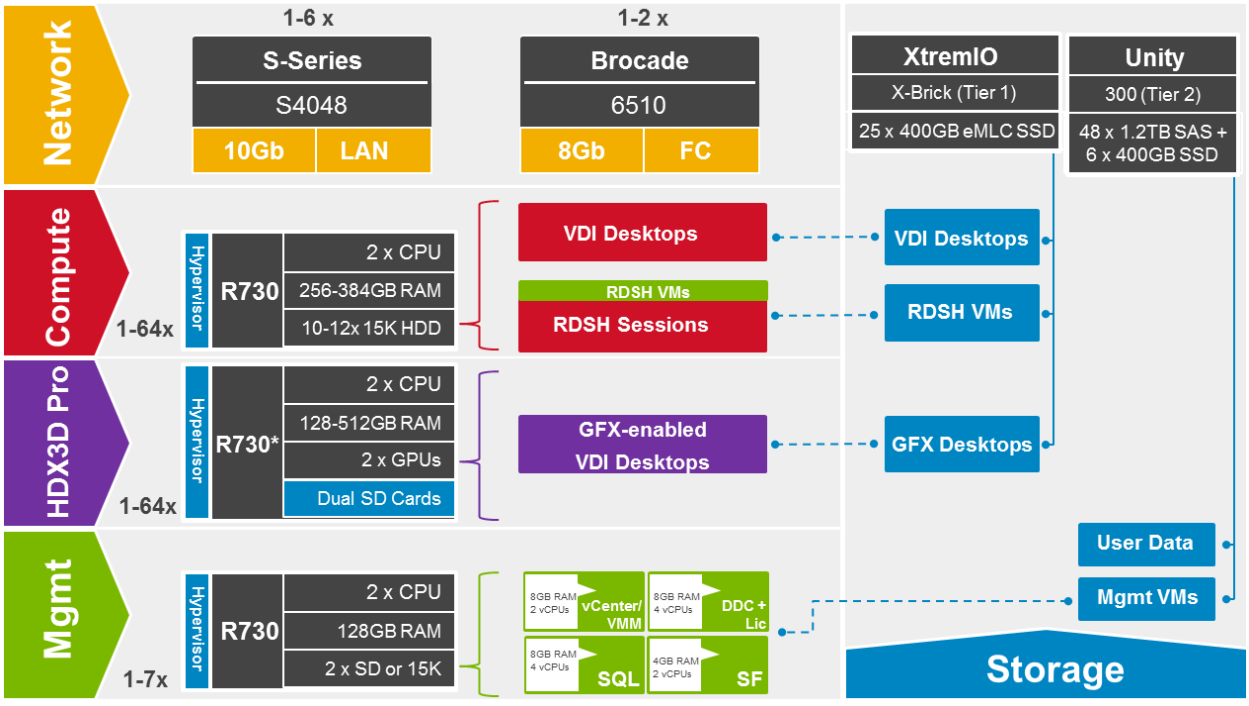

* The Dell Precision Appliance for Wyse is a R730 server based solution

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 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.

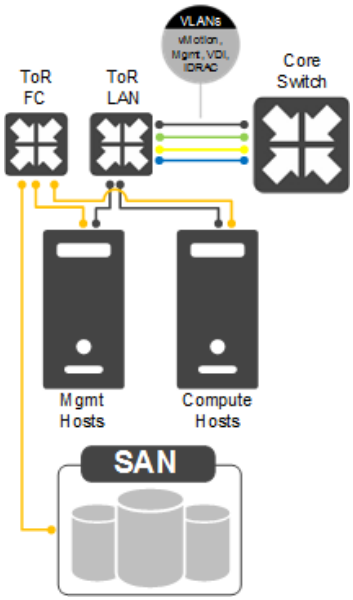

2.5.1.2 Shared Tier 1 – Rack Cabling (HA)

500 Users or Less - Combined T1 and T2 Option

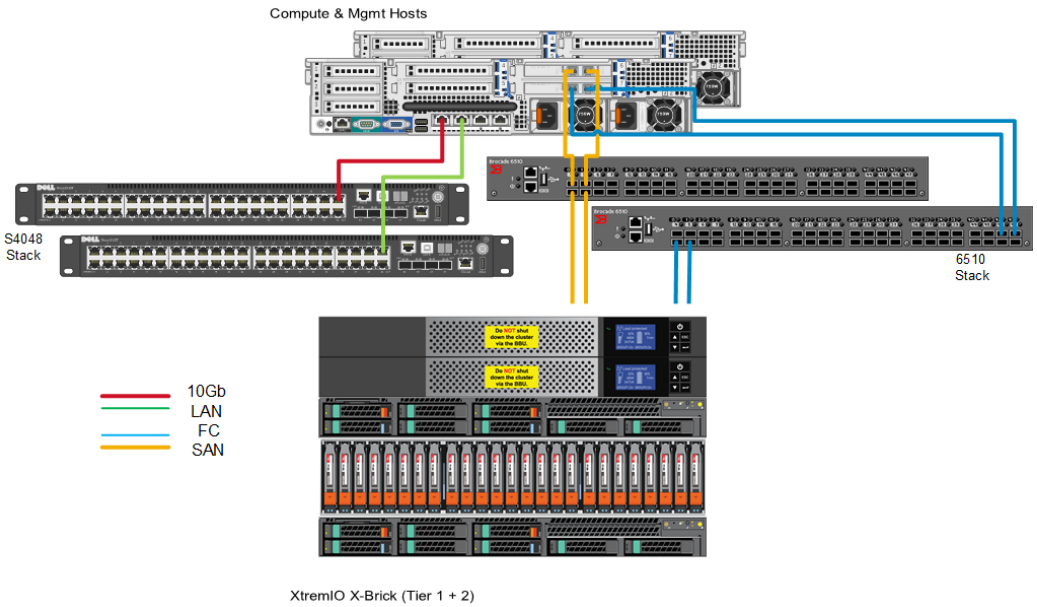

Discrete T1 and T2 Storage Arrays

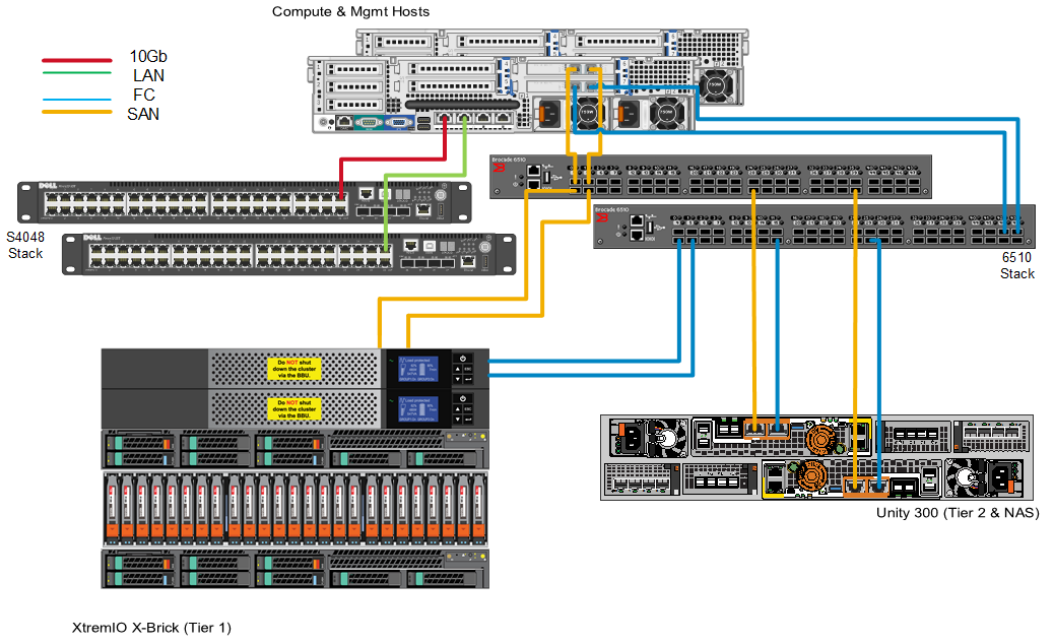

Refer to the Unity 300 Tier 2 section for a diagram of disk enclosure connections via mini-SAS HD.

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 ST1 model using rack servers, blades can also be used in a 500 user bundle by combing T1 and T2 on an XtremIO Starter X-Brick storage array with Unity VSA. Above 500 users, separate T1 and T2 storage into discrete arrays and add an appropriate number of arrays based on user scaling. In the configurations shown below, ToR FC switching is optional if a suitable infrastructure is already in place.

500 Users or Less

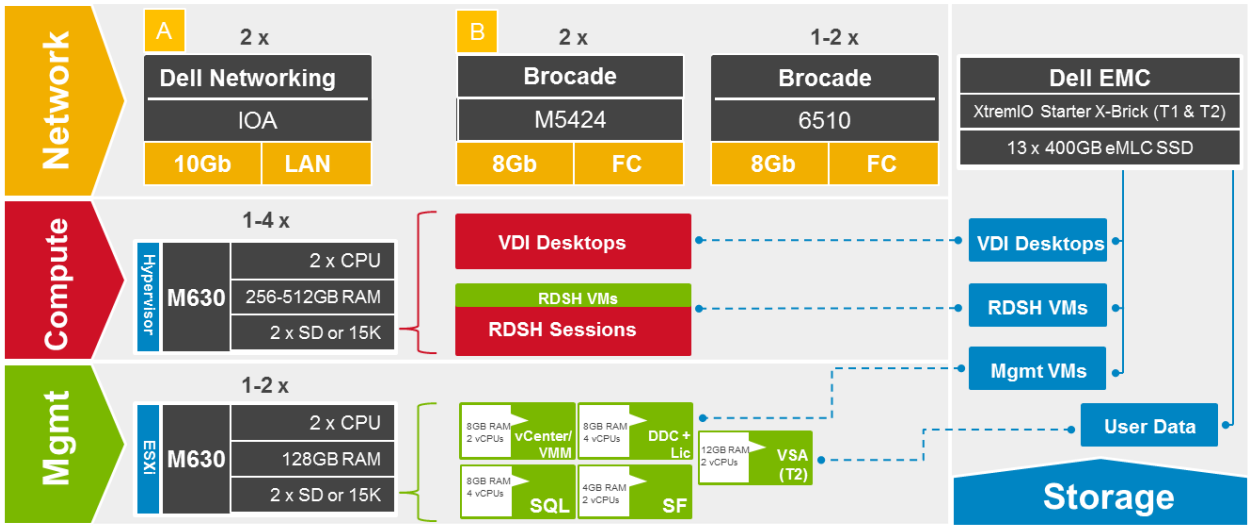

Note: The Unity VSA VM must be deployed on vSphere.

More Than 500 Users

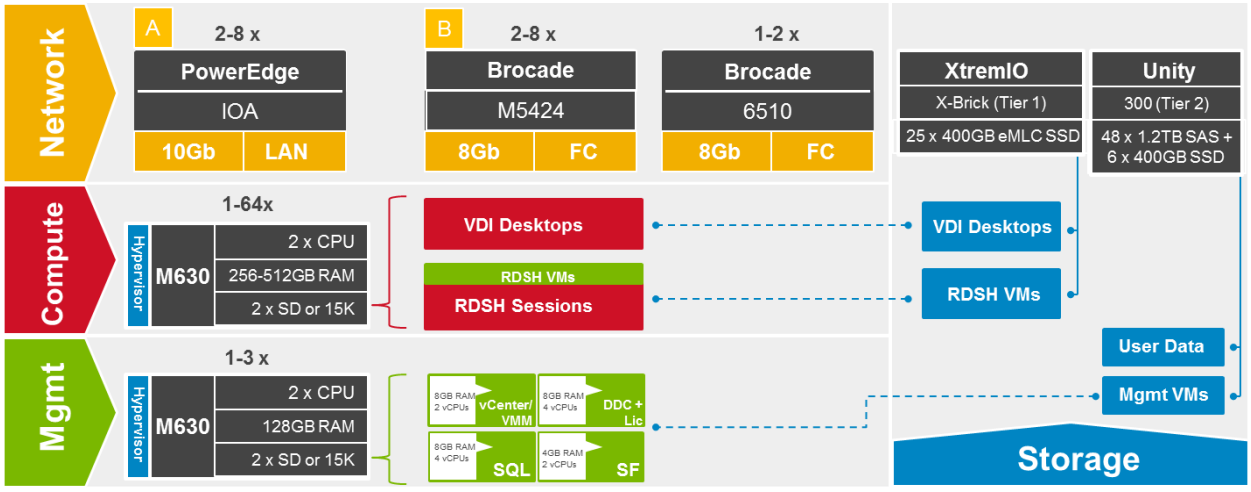

2.6.1.1 Shared Tier 1 – Network Architecture

In the Shared Tier 1 architecture for blade servers using FC, there is no need to switch LAN ToR since the IOAs in the chassis supports LAN to the blades and are uplinked to the core or distribution layers directly. However, a separate switching infrastructure is required for FC. Management and compute servers both connect to shared storage using FC switched via chassis interconnects. Both management and compute servers connect to all network VLANs in this model. For greater redundancy, a ToR switch is used to support iDRAC used outside of the chassis. 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.

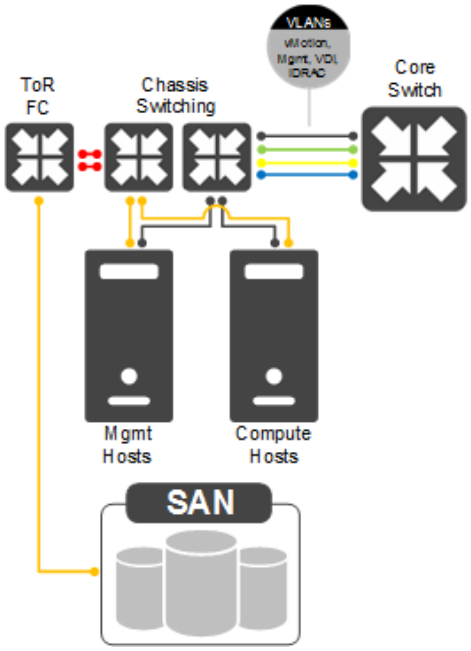

2.6.1.2 Shared Tier 1 – Rack Cabling (HA)

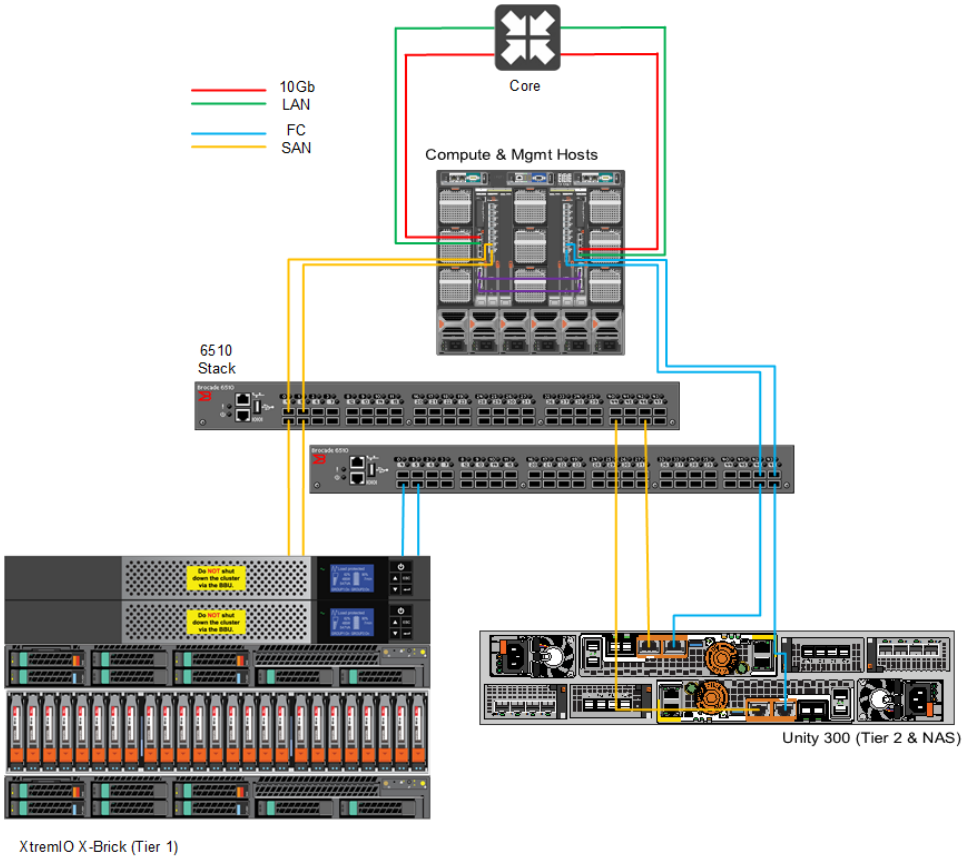

Refer to the Unity 300 Tier 2 section for a diagram of disk enclosure connections via mini-SAS HD.

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 S-Series 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 VRF-lite, Routed VLT, VLT Proxy Gateway User port stacking (up to 6 switches) Open Networking Install Environment (ONIE)	1Gb connectivity

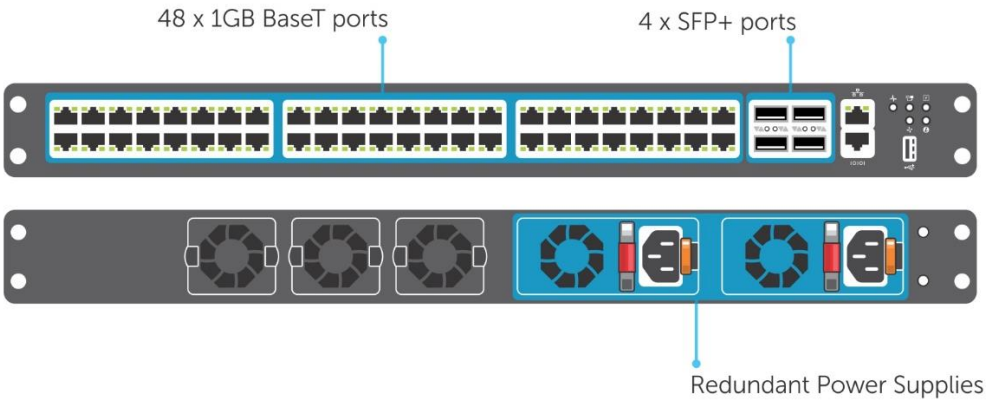

3.1.2 Dell Networking S-Series S4048 (10Gb ToR Switch)

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)	

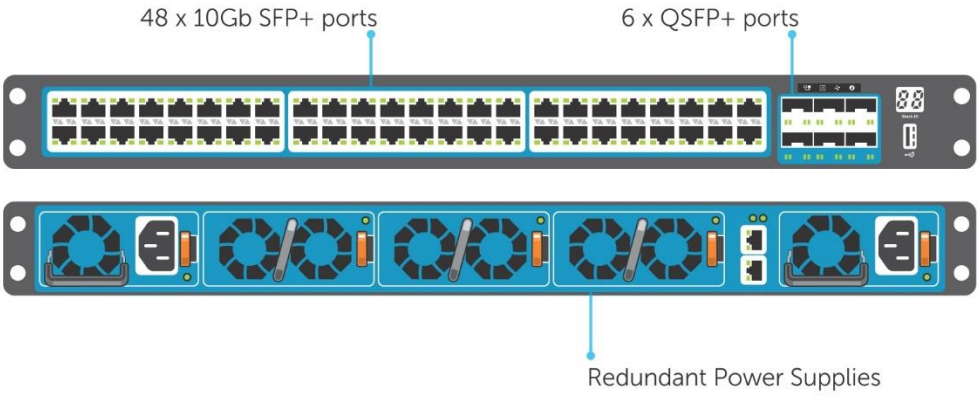

For more information on the S3048, S4048 switches and Dell Networking, please visit: LINK

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 Fiber 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

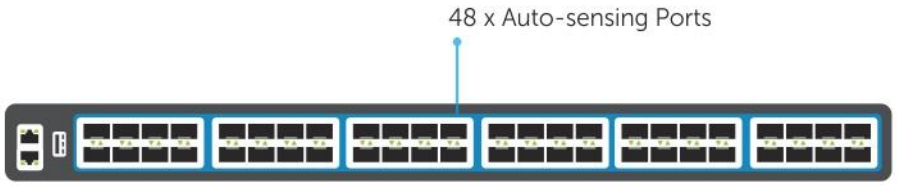

For more information on the Brocade 6510 switch, please visit: LINK

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	

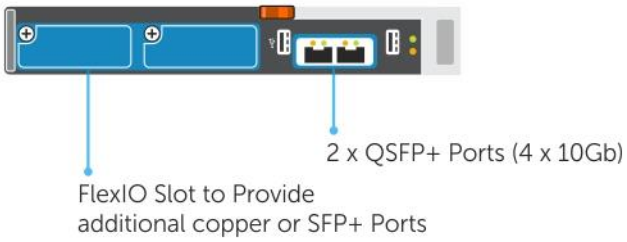

For more information on the PowerEdge IOA switch, please visit: LINK

3.2 Servers

3.2.1 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 CPUs, can host up to 24 DIMMs of DDR4-2400 RAM, supports up to 16 x 2.5" SAS disks and can be outfitted with 2 double-wide GPU accelerators from AMD or NVIDIA. The Dell PowerEdge R730 offers uncompromising performance and scalability in a 2U form factor.

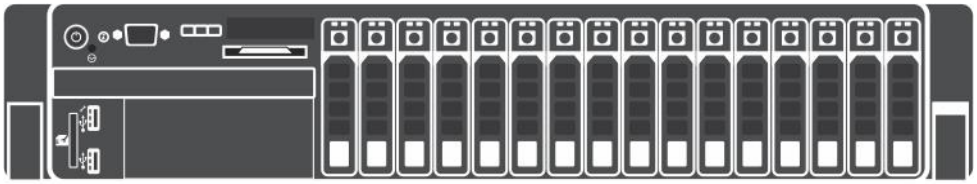

For more information on the R730, please visit: Link

3.2.1.1 Local Tier 1 Rack

In the Local Tier 1 model, VDI desktops or XenApp sessions execute from local storage on each Compute server. The hypervisors used in this solution are vSphere or Hyper-V. In this model, only the Management server hosts access FC storage to support the solution's Management role VMs. Additional NICs can be added as required for increased bandwidth or resiliency requirements. Refer to the Local Tier 1 section for cabling implications. The Management server host has reduced RAM and CPU and does not require local disk to host the management VMs. Any of the compute options below can also be used for XenApp while optionally reducing the amount of RAM. All-flash shown for compute below, the SSDs can be optionally substituted for 10-12 15K SAS.

Local Tier 1 Compute Host – PowerEdge R730	
2 x Intel Xeon E5-2698v4 20C CPU (2.2Ghz)	2 x Intel Xeon E5-2698v4 20C CPU (2.2Ghz)
512GB Memory (16 x 32GB RDIMMs, 2400MT/s)	512GB Memory (16 x 32GB RDIMMs, 2400MT/s)
VMware vSphere on internal 8GB Dual SD	**Microsoft Hyper-V** on local disks
4 x 800GB SSD or 10 x 15K SAS	4 x 800GB SSD or 12 x 15K SAS
PERC H730 Integrated RAID Controller – RAID10	PERC H730 Integrated RAID Controller – RAID10
Embedded 4 x 10Gb NDC	Embedded 4 x 10Gb NDC
iDRAC8 Enterprise	iDRAC8 Enterprise
2 x 750W PSUs	2 x 750W PSUs

Local Tier 1 Management Host PowerEdge R730	
2 x Intel Xeon E5-2660v4 14C CPU (2Ghz)	2 x Intel Xeon E5-2660v4 14C CPU (2Ghz)
128GB Memory (8 x 16GB RDIMMs, 2400MT/s)	128GB Memory (8 x 16GB RDIMMs, 2400MT/s)
VMware vSphere on internal 8GB Dual SD	**Microsoft Hyper-V** on 2 x 300GB 15K SAS
Embedded 4 x 10Gb NDC	Embedded 4 x 10Gb NDC
2 x QLogic 2562 8Gb DP FC HBA	2 x QLogic 2562 8Gb DP FC HBA
iDRAC8 Enterprise	iDRAC8 Enterprise
2 x 750W PSUs	2 x 750W PSUs

3.2.1.2 Graphics Compute Server

A Dell Precision Appliance for Wyse (R730) can be added as a graphics compute host in either Local Tier 1 or Shared Tier 1 architecture models when ESXi is used as the hypervisor. Recommended system component are shown in the table below. Due to the reduced densities on GPU-enabled compute hosts, recommended RAM has been reduced. However, the memory, CPUs, disks, and NICs can be adjusted to suit customer requirements.

Local or Shared Tier 1 GFX Compute Host – Dell Precision Appliance for Wyse (PowerEdge R730)
2 x Intel Xeon E5-2698v4 20C CPU (2.2Ghz) or Intel Xeon E5-2667v4 8C CPU (3.2GHz)
256GB Memory (16 x 16GB RDIMMs, (2400MT/s)
VMware vSphere on internal 8GB Dual SD
4 x 800GB SSD or 10 x 15K SAS
PERC H730 Integrated RAID Controller – RAID10
Embedded 4 x 10Gb NDC
2 x NVIDIA Tesla M60 or M10 GPUs
iDRAC8 Enterprise
2 x 1100W PSUs

3.2.1.2.1 Shared Tier 1 Rack (FC)

In the Shared Tier 1 model, VDI desktops or XenApp sessions execute on shared storage so there is no need for local disks on each server to host VMs. Fibre Channel is leveraged as the block storage protocol for Compute and Management hosts with Tier 1 and Tier 2 storage. All configuration options (beyond the hypervisor differences noted in the table below) are identical except for CPU and RAM which are reduced on the Management host.

Shared Tier 1 Compute Host PowerEdge R730	
2 x Intel Xeon E5-2698v4 20C CPU (2.2Ghz)	2 x Intel Xeon E5-2698v4 20C CPU (2.2Ghz)
512GB Memory (16 x 32GB RDIMMs, 2400MT/s)	512GB Memory (16 x 32GB RDIMMs, 2400MT/s)
VMware vSphere on internal 8GB Dual SD	**Microsoft Hyper-V** on 2 x 300GB 15K SAS
Embedded 4 x 10Gb NDC	Embedded 4 x 10Gb NDC
2 x QLogic 2562 8Gb DP FC HBA	2 x QLogic 2562 8Gb DP FC HBA
iDRAC8 Enterprise	iDRAC8 Enterprise
2 x 750W PSUs	2 x 750W PSUs

Shared Tier 1 Management Host PowerEdge R730	
2 x Intel Xeon E5-2660v4 14C CPU (2Ghz)	2 x Intel Xeon E5-2660v4 14C CPU (2Ghz)
128GB Memory (8 x 16GB RDIMMs, 2400MT/s)	128GB Memory (8 x 16GB RDIMMs, 2400MT/s)
VMware vSphere on internal 8GB Dual SD	**Microsoft Hyper-V** on 2 x 300GB 15K SAS
Embedded 4 x 10Gb NDC	Embedded 4 x 10Gb NDC
2 x QLogic 2562 8Gb DP FC HBA	2 x QLogic 2562 8Gb DP FC HBA
iDRAC8 Enterprise	iDRAC8 Enterprise
2 x 750W PSUs	2 x 750W PSUs

3.2.2 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-CPU 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 CPUs (Broadwell) and 24 DIMMs (768GB RAM) of DDR4 memory in an extremely compact half-height blade form factor.

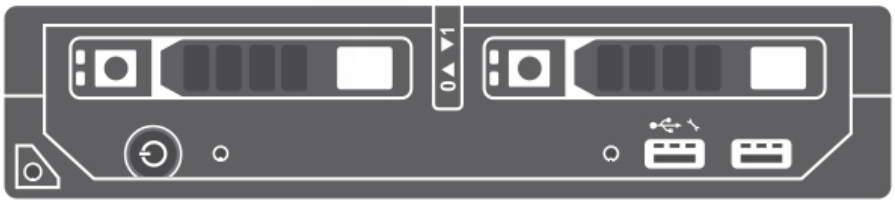

For more information on the PowerEdge M630, please visit: Link

3.2.2.1 Local Tier 1 Blade

In the Local Tier 1 model for blades, VDI desktops or XenApp 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. The recommended provisioning method is MCS with non-persistent desktops. Refer to the Local Tier 1 for Blade Servers section for solution cabling implications.

Local Tier 1 Compute Host – PowerEdge M630
2 x Intel Xeon E5-2698v4 20C CPU (2.2GHz)
512GB Memory (16 x 32GB RDIMMs, 2400MT/s)
VMware vSphere on 2 x 8GB internal SD
2 x SSD
QLogic 57810S-k 10Gb DP KR NDC (LAN)
iDRAC8 Enterprise

Local Tier 1 Management Host – PowerEdge M630
2 x Intel Xeon E5-2660v4 14C CPU (2GHz)
128GB Memory (8 x 16GB RDIMMs, 2400MT/s)
VMware vSphere on 2 x 8GB internal SD
1 x QLogic QME2572 8Gb FC mezz (FC - HA)
iDRAC8 Enterprise

3.2.2.2 Shared Tier 1 Blade (FC)

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.

Shared Tier 1 Compute Host – PowerEdge M630		
2 x Intel Xeon E5-2698v4 20C CPU (2.2GHz)	**Or**	2 x Intel Xeon E5-2698v4 20C CPU (2.2GHz)
512GB Memory (16 x 32GB RDIMMs, 2400MT/s)		512GB Memory (16 x 32GB RDIMMs, 2400MT/s)
VMware vSphere on 2 x 2GB internal SD		**Microsoft Hyper-V** on 2 x 300GB 15K SAS
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		iDRAC8 Enterprise

Shared Tier 1 Management Host – PowerEdge M630		
2 x Intel Xeon E5-2660v4 14C CPU (2GHz)	**Or**	2 x Intel Xeon E5-2660v4 14C CPU (2GHz)
128GB RAM (8 x 16GB RDIMMs @ 2400MT/s)		128GB RAM (8 x 16GB RDIMMs @ 2400MT/s)
VMware vSphere on 2 x 8GB internal SD		**Microsoft Hyper-V** on 2 x 300GB 15K SAS
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		iDRAC8 Enterprise

3.3 Storage

3.3.1 XtremIO X-Brick – Tier 1

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, XtremIO provides customers with a great return on their VDI investment.

A Citrix XenDesktop environment backed by the XtremIO X-Brick for desktop storage and the EMC 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.

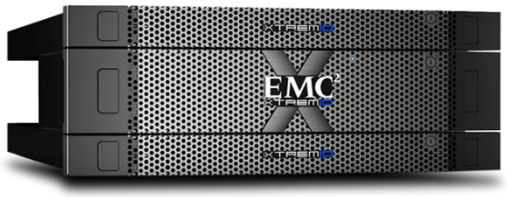

The XtremIO all-flash storage array has a revolutionary architecture with the following elements to enable the agile data center: linear scale-out, inline all-the-time data services, and ample data center services for the workloads.

The basic hardware building block for the XtremIO array is the EMC 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 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. An 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.

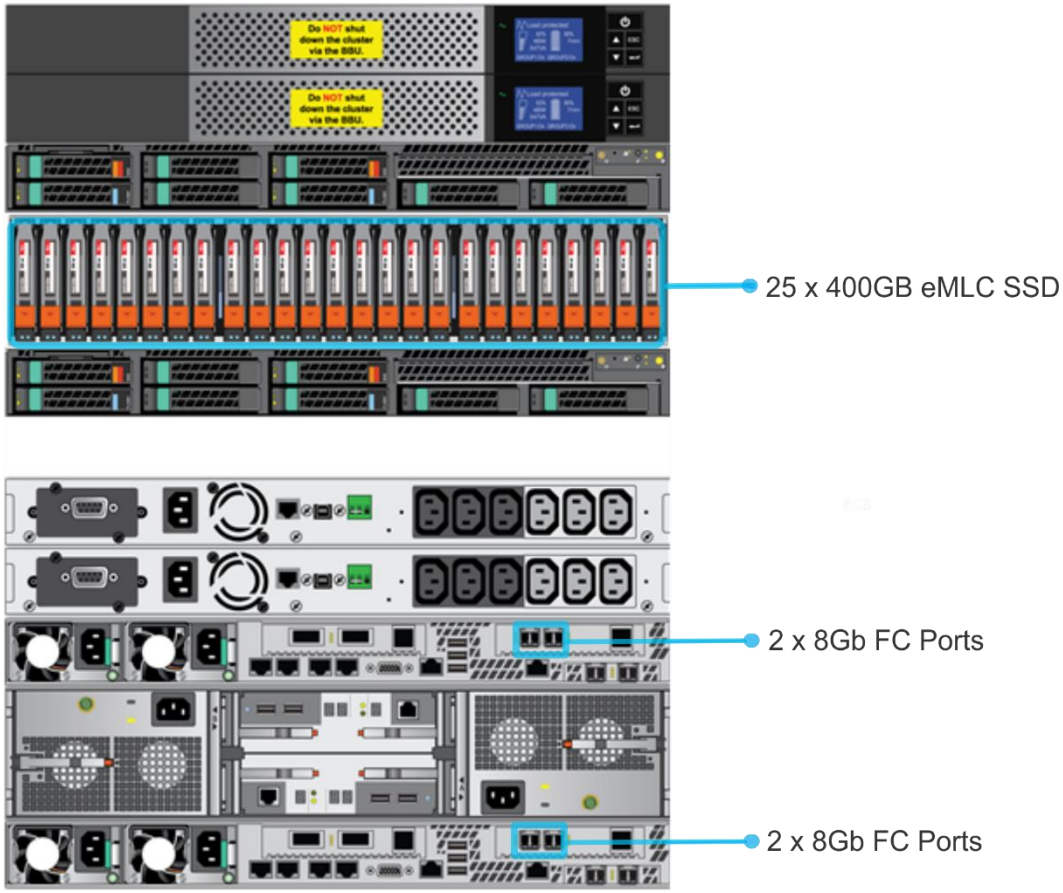

3.3.2 Unity 300 – Tier 2

Dell EMC Unity™ is the only storage system that successfully meets all 4 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, EMC's 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).

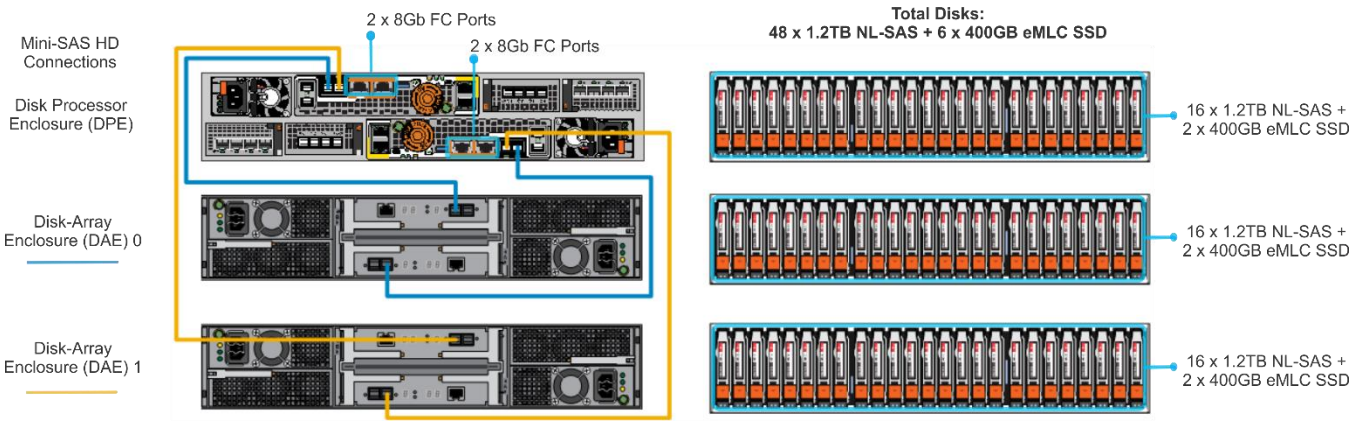

3.3.3 Storage Configuration

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.

3.3.4 XtremIO Starter X-Brick - Combined Tier 1 and Tier 2

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

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)

3.3.5 Network Attached Storage (NAS)

3.3.5.1 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. UnityVSA allows the advanced unified storage and data management features of the 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 Worker 4.1.5+ workload

Note: The Unity VSA VM must be deployed on vSphere.

3.3.5.2 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.

3.4 GPUs

3.4.1 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. Plus, Tesla accelerators deliver the horsepower needed to run bigger simulations faster than ever before. For enterprises deploying VDI, Tesla accelerators are perfect for accelerating virtual desktops.

3.4.1.1 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

3.4.1.2 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.

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

Specs	Tesla M60
Number of GPUs	2 x NVIDIA Maxwell™ GPUs
Total CUDA cores	4096 (2048 per GPU)

Base Clock	899 MHz (Max: 1178 MHz)
Total memory size	16GB GDDR5 (8GB per GPU)
Max power	300W
Form Factors	Dual slot (4.4” x 10.5”)
Aux power	8-pin connector
PCIe	x16 (Gen3)
Cooling solution	Passive/ Active

3.5 Dell Wyse Endpoints

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

3.5.1 Wyse 3020 Zero Client for Citrix

The Wyse 3020 zero client for Citrix provides an exceptional user experience at a highly affordable price for Citrix XenDesktop and XenApp published desktops. With no attack surface, these zero clients offer protection against network-borne viruses and malware zero attacks. The Wyse 3020 draws less than 7 watts of power in full operation and powers on in just six seconds. The Wyse 3020’s extremely compact and efficient design sports gigabit LAN and dual band wireless a/b/g/n, so it’s ready for serious tasks whether you’re wired in or not. These can be connected to LAN, WAN, or Wi-Fi networks, as these zero clients easily adapt to the latest network security and protocols. For more information, please visit: Link

3.5.2 Wyse 3030 LT Thin Client with ThinOS

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

3.5.3 Wyse 5010 Thin Client with ThinOS

Designed for knowledge workers and power users, the Wyse 5010 is a high performance thin client based on Wyse ThinOS, the virus-resistant firmware base designed for optimal thin client security, performance, and ease-of-use. Highly secure, compact and powerful, it combines a dual-core AMD 1.4 GHz CPU with a revolutionary unified graphics engine for an outstanding user experience. It addresses the performance challenges of processing-intensive applications like computer-aided design, multimedia, HD video and 3D modelling. Scalable on premise or cloud-based management provides simple deployment, patching, and updates. Take a unit from box to productivity in minutes with auto configuration. Delivering outstanding processing speed and power, security and display performance, the Wyse 5010 offers a unique combination of performance, efficiency, and affordability. For more information, please visit: Link

3.5.4 Wyse 5060 Thin Client with ThinOS and Xenith mode

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. ThinOS-based Wyse 5060 thin clients can be deployed as zero clients within a Citrix installation giving customers the ability to leverage the features or capabilities that define the Citrix zero client experience. 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) or cloud-based Wyse Cloud Client Manager (CCM). For more information, please visit: Link

3.5.5 Wyse 7020 Thin Client with ThinLinux

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. Designed to provide a superior user experience, ThinLinux features broad broker support including Citrix Receiver, VMware Horizon and Amazon Workspace, and support for unified communication platforms including Skype for Business, Lync 2013 and Lync 2010. For additional security, ThinLinux also supports single sign-on and VPN. 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

3.5.6 Wyse 7020 Thin Client with 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. This 64-bit thin client with Windows 10 IoT delivers a great user experience and support for local applications while ensuring security with features such as Trusted Platform Module (TPM), BitLocker Drive Encryption, Secure Boot, and Windows Defender. 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. 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

3.5.7 Wyse 7040 Thin Client with Windows Embedded Standard 7P

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

3.5.8 Latitude E7270 mobile with client

The high-powered Latitude E7270 mobile thin client is designed to securely deliver virtual desktops to mobile users. Equipped with a powerful quad core processor and integrated graphics engine, the mobile thin client provides an exceptional user experience with a high level of compute power for the most demanding applications. This thin client for on-the-go performance, offers a 12'' display with a Full HD anti-glare display and longer battery life. The Latitude E7270 mobile thin client includes a rich array of thin client features on a mobile form factor. It supports a wide variety of virtual desktop brokers including Citrix XenDesktop and VMware Horizon and can connect to a broad range of peripherals and desktop / application virtualization environments using a fast, rich and familiar Windows 7 user experience. The Latitude E7270 mobile thin client is easily manageable through Wyse Device Manager (WDM), Wyse Cloud Client Manager and Microsoft's System Center Configuration Manager (SCCM). For enhanced security, optional Advanced Threat Protection in the form of Dell Endpoint Security Suite Enterprise and Dell Threat Defense offer proactive malware protection on both virtual desktops and the endpoints themselves, respectively. For more information, please visit: Link

4 Software Components

4.1 Citrix

4.1.1 Citrix XenDesktop

The solution is based on Citrix XenDesktop which provides a complete end-to-end solution delivering Microsoft Windows virtual desktops or server-based hosted shared sessions 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.

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

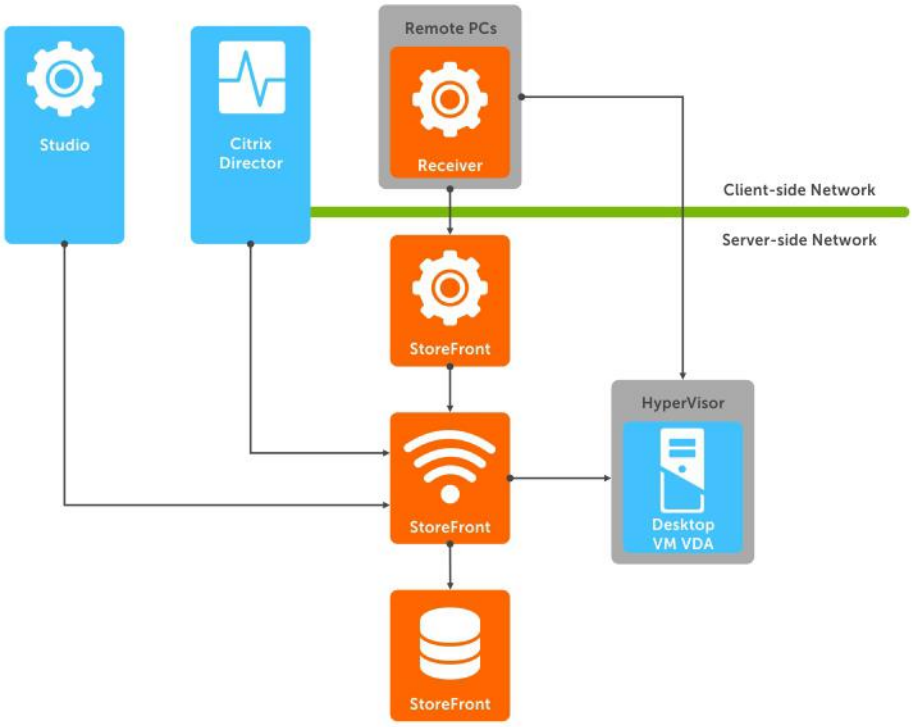

The core XenDesktop components include:

- **Studio**

Studio is the management console that enables you to configure and manage your deployment, eliminating the need for separate management consoles for managing delivery of applications and desktops. Studio provides various wizards to guide you through the process of setting up your environment, creating your workloads to host applications and desktops, and assigning applications and desktops to users.

- **Delivery Controller** (DC)

Installed on servers in the data center, the controller authenticates users, manages the assembly of users' virtual desktop environments, and brokers connections between users and their virtual desktops. The Controller also manages the state of desktops, starting and stopping them based on demand and administrative configuration.

- **Database**

At least one Microsoft SQL Server database is required for every XenApp or XenDesktop Site to store configuration and session information. The Delivery Controller must have a persistent connection to the database as it stores data collected and managed by the Controller services.

- **Director**

Director is a web-based tool that enables IT support teams to monitor an environment, troubleshoot issues before they become system-critical, and perform support tasks for end users. You can also view and interact with a user's sessions using Microsoft Remote Assistance. Starting in version 7.12, Director now includes detailed descriptions for connection and machine failures, one month historical data (Enterprise edition), custom reporting, and notifications via SNMP traps.

- **Receiver**

Installed on user devices, Citrix Receiver provides users with quick, secure, self-service access to documents, applications, and desktops from any of the user's devices including smartphones, tablets, and PCs. Receiver provides on-demand access to Windows, Web, and Software as a Service (SaaS) applications. For devices that cannot install the Receiver software, Citrix Receiver for HTML5 provides connectivity through a HTML5-compatible web browser.

- **StoreFront**

StoreFront authenticates users to sites hosting resources and manages stores of desktops and applications that user's access. StoreFront version 3.8 (released with XenDesktop 7.12) and above includes ability to create and use multiple IIS websites each having its own domain name.

- **License Server**

The Citrix License Server is an essential component at any Citrix-based solution. Every Citrix product environment must have at least one shared or dedicated license server. License servers are computers that are either partly or completely dedicated to storing and managing licenses. Citrix products request licenses from a license server when users attempt to connect.

- **Machine Creation Services** (MCS)

A collection of services that work together to create virtual servers and desktops from a master image on demand; optimizing storage utilization and providing a pristine virtual machine to users every time they log on. Machine Creation Services is fully integrated and administrated in Citrix Studio.

- **Provisioning Services** (PVS)

The Provisioning Services infrastructure is based on software-streaming technology. This technology allows computers to be provisioned and re-provisioned in real-time from a single shared-disk image.

- **Virtual Delivery Agent** (VDA)

The Virtual Desktop Agent is a transparent plugin that is installed on every virtual desktop or XenApp host (RDSH) and enables the direct connection between the virtual desktop and users' endpoint devices. Windows and Linux VDAs are available.

4.1.2 Machine Creation Services (MCS)

Citrix Machine Creation Services is the native provisioning mechanism within Citrix XenDesktop for virtual desktop image creation and management. Machine Creation Services uses the hypervisor APIs to create, start, stop, and delete virtual desktop images. Desktop images are organized in a Machine Catalog and within that catalog there are a number of options available to create and deploy virtual desktops:

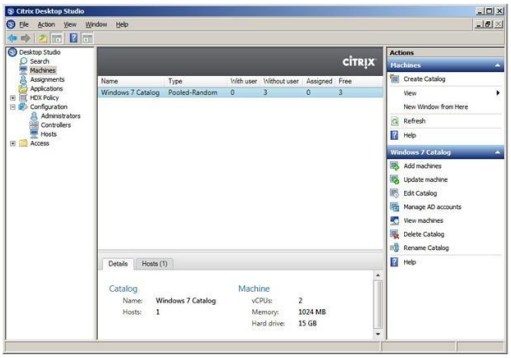

- **Random**: Virtual desktops are assigned randomly as users connect. When they logoff, the desktop is reset to its original state and made free for another user to login and use. Any changes made by the user are discarded at log off.
- **Static**: Virtual desktops are assigned to the same user every time with three options for how to handle changes made to the desktop: Store on local vDisk, Personal vDisk, or discarded on user log off.

All the desktops in a random or static catalog are based off a master desktop template which is selected during the catalog creation process. MCS then takes snapshots of the master template and layers two additional virtual disks on top: an Identity vDisk and a Difference vDisk. The Identity vDisk includes all the specific desktop identity information such as host names and passwords. The Difference vDisk is where all the writes and changes to the desktop are stored. These Identity and Difference vDisks for each desktop are stored on the same data store as their related clone.

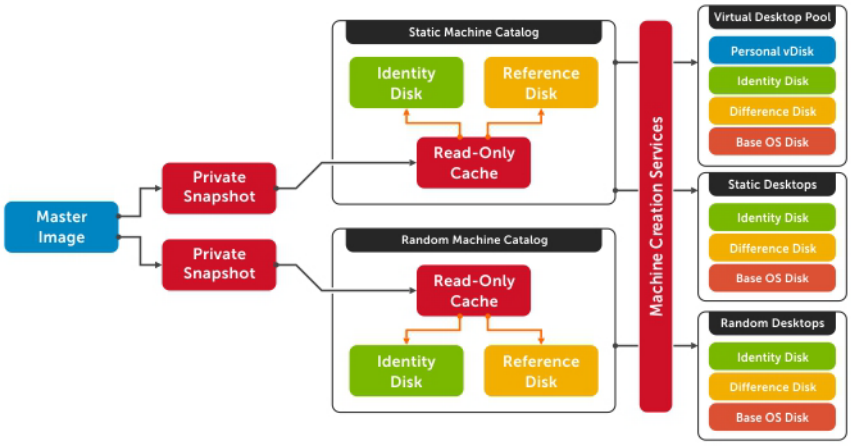

While traditionally used for small to medium sized XenDesktop deployments, MCS can bring along with it some substantial Tier 1 storage cost savings because of the snapshot/identity/difference disk methodology. The Tier 1 disk space requirements of the identity and difference disks when layered on top of a master image snapshot, is far less than that of a dedicated desktop architecture.

4.1.3 Provisioning Services (PVS)

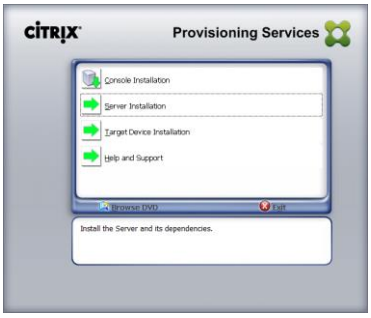

PVS is an alternative method of image provisioning which uses streaming to share a single base vDisk image instead of copying images to VMs. PVS are used to deliver shared vDisk images to physical or virtual machines. Another potential use is the serial provisioning of XenApp to enable scale-out hosted shared desktop infrastructure. Provisioning Services enables real-time streamed provisioning and re-provisioning which enable administrators to completely eliminate the need to manage and patch individual systems.

Desktop images are organized in a Machine Catalog and within that catalog there are a number of options available to create and deploy virtual or physical desktops:

- **Random**: Virtual or physical desktops are assigned randomly as users connect. When they logoff, the desktop is reset to its original state and made free for another user to login and use. Any changes made by the user are discarded at log off.
- **Static**: Virtual desktops are assigned to the same user every time with user changes stored on a separate Personal vDisk.

Using Provisioning Services, vDisk images are configured in Standard Image mode, read-only, or Private Image mode, read/write. A vDisk in Standard Image mode allows multiple desktops to boot from it simultaneously greatly reducing the number of images that must be maintained and the amount of storage that is otherwise required (non-persistent). Private Image mode vDisks are equivalent to dedicated hard disks and can only be used by one target device at a time (persistent). The Provisioning Server runs on a virtual instance of Windows Server 2012 R2 or Windows 2016 on the Management Server(s).

4.1.3.1 PVS Write Cache

Citrix Provisioning Services delivery of standard images relies on write-caches to store any writes made by the target OS. The most common write-cache implementation places write-cache on the target machine's storage. Independent of the physical or virtual nature of the target machine, this storage has to be allocated and formatted to be usable.

While there are 4 possible locations for storage of the write cache in PVS, the Dell Wyse Datacenter solution recommends placement of the PVS write cache in the target compute host's RAM with overflow enabled. We recommend using a cache size of 512MB for virtual desktops and 21GB for XenApp VMs delivered via PVS.

4.1.4 Personal vDisk

Citrix Personal vDisk is an enterprise workspace virtualization solution that is built into Citrix XenDesktop. Personal vDisk provides the user customization and personalization benefits of a persistent desktop image with the storage savings and performance of a single/shared image.

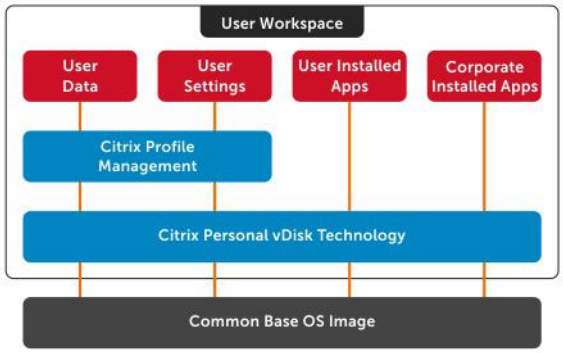

Used in conjunction with a static desktop experience, Citrix Personal vDisk allows each user to receive personal storage in the form of a layered vDisk (3GB minimum). This personal vDisk enables users to personalize and persist their desktop environment while providing storage for any user or departmental apps.

Personal vDisk provides the following benefits to XenDesktop:

- Persistent personalization of user profiles, settings and data
- Enables deployment and management of user installed and entitlement based applications
- Fully compatible with Microsoft SCCM and App-V
- 100% persistence with VDI pooled Storage management
- Near Zero management overhead

4.1.5 AppDisks

Citrix AppDisk, included in XenDesktop 7.8 and higher, provides layering technology to manage departmental applications as an independent storage layer. AppDisk eases the management burden of maintaining multiple departmental images by instantly layering applications onto a single golden image that remains separate and pristine. AppDisks can be associated with either published desktops or published applications via XenApp. AppDisk does not replace the functionality provided by Personal vDisk but currently cannot be used within the same golden image either. AppDisks, when integrated with AppDNA, provide the ability to analyze OS and application performance, compatibility, and remediation capabilities.

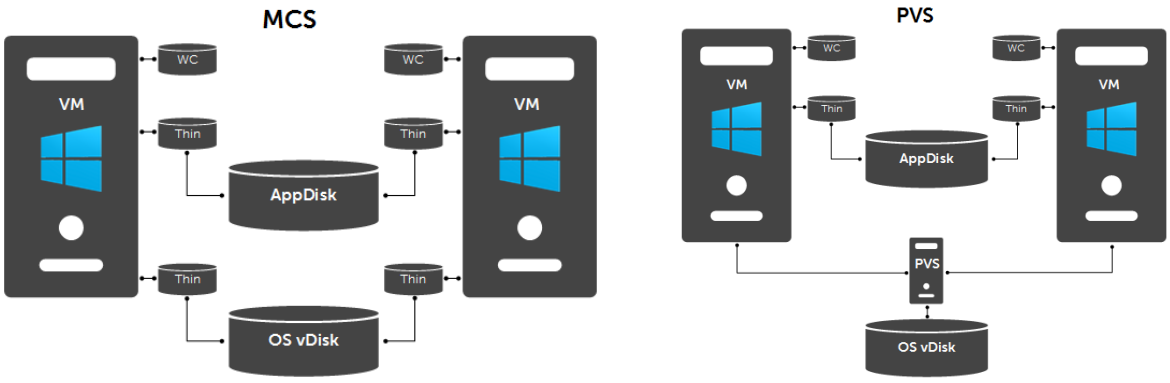

4.1.6 HDX 3D Pro

XenDesktop with HDX 3D Pro is a desktop and app virtualization solution that supports high-end designers and engineers of 3D professional graphics applications as well as provides cost-effective support to viewers and editors of 3D data. With XenDesktop, you can deliver a persistent user experience and leverage other virtualization benefits such as single-image management and improved data security.

Use HDX 3D Pro technologies with:

- Computer-aided design, manufacturing, and engineering (CAD/CAM/CAE) applications
- Geographical information system (GIS) software
- Picture Archiving Communication System (PACS) workstations for medical imaging
- Latest OpenGL, DirectX, CUDA and CL versions supported
- Latest NVIDIA Grid cards
- Shared or dedicated GPUs or a mix of both on desktop or server OS VMs

HDX 3D Pro provides the best user experience over any bandwidth using Framehawk integration:

- On wide area network (WAN) connections: Deliver an interactive user experience over WAN connections with bandwidths as low as 1.5 Mbps.
- On local area network (LAN) connections: Deliver a user experience equivalent to that of a local desktop on LAN connections.

Framehawk is a display remoting technology implemented as an ICA virtual channel that optimizes delivery of virtual desktops and applications to users on broadband wireless connections where high packet loss or congestion occurs.

4.1.7 Citrix Profile Manager

Citrix Profile Management is a component of the XenDesktop suite which is used to manage user profiles and minimize many of the issues associated with traditional Windows roaming profiles in an environment where users may have their user profile open on multiple devices at the same time. The profile management toolset has two components: the profile management agent, installed on any device where the user profiles is managed, and a Group Policy Administrative Template, which is imported to a group policy.

In order to further optimize, the profile management folders within the user profile is redirected the users' home drive. The folder redirection is managed via group policy objects within Active Directory. The following folders are redirected:

- Contacts
- Downloads
- Favorites
- Links
- My Documents
- Searches
- Start Menu
- Windows
- My Music
- My Pictures
- My Videos
- Desktop

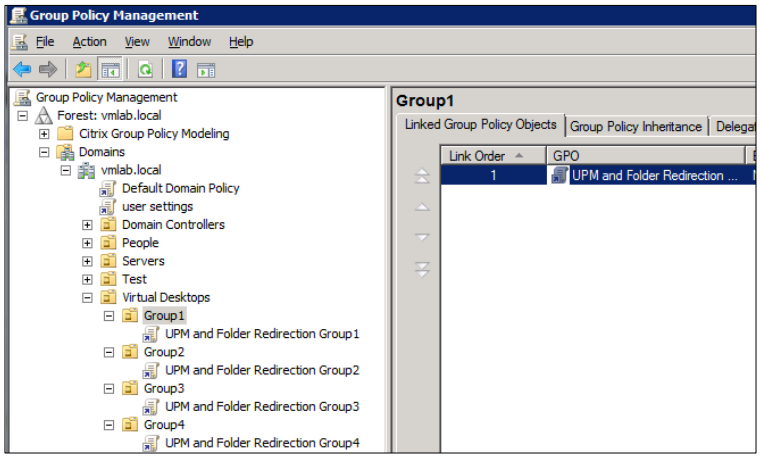

4.1.8 Citrix XenApp

Citrix XenApp version 7.6 and above includes enhancements in the areas of faster access to virtual apps with higher connection resiliency, improved graphics rendering, and new app-usage reporting and monitoring tools.

Citrix XenApp delivers Windows apps as secure mobile services. With XenApp, IT can mobilize the business - increasing user productivity, while reducing costs by centralizing control and security of intellectual property. XenApp delivers high-performance apps to any PC, Mac, laptop, tablet or smartphone that enable the delivery of a native experience that is optimized for the type of device, as well as the network. XenApp is built on a 3rd generation FlexCast Management Architecture (FMA) and is the only hybrid cloud-ready platform that separates the management plane from the workload to enable IT to securely deliver published apps on-premises, and manage workers and mobile workspaces either on-premises or in the cloud.

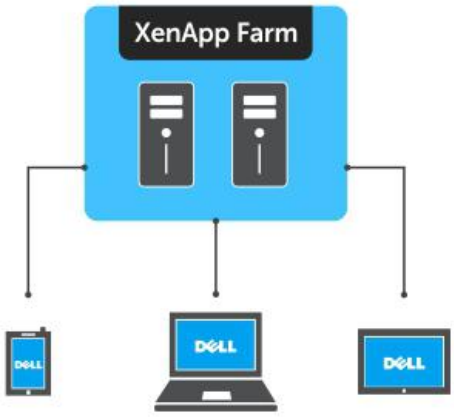

Benefits of hosted desktop sessions and applications:

- Management of applications (single instance)
- Management of simple desktop images (no applications installed)
- PVS to stream XenApp servers as well as user desktops
- Scalability of XenDesktop compute hosts: CPU and IOPS reduction via application offload
- Shared storage scalability: less IOPS = more room to grow

Citrix XenDesktop with XenApp integration can effectively deliver a desktop/application hybrid solution as well. Specifically where a single or small number of shared VDI desktop images are deployed via XenDesktop, each with common shared applications installed within the golden image. A user-specific application set is then deployed and made accessible via the hosted application compute infrastructure, accessible from within the virtual desktop.

User Environment	XenDesktop	XenApp
User-Specific Applications		✓
Profile and User Data	✓	✓
Shared Applications	✓	
Shared Virtual Desktop Image	✓	

Alternatively, XenApp provides a platform for delivering Windows server-based sessions to users who may not need a full desktop VM. Hosted desktops increase infrastructure resource utilization while reducing complexity as all applications and sessions are centrally managed.

User Environment	XenDesktop	XenApp
User-Specific Applications		✓
Profile and User Data		✓
Dedicated Virtual Desktop Image		✓

4.1.8.1 XenApp Integration into Dell Wyse Datacenter Architecture

The XenApp servers can exist as physical or virtualized instances of Windows Server 2012 R2 or Windows 2016. A minimum of one (1), up to a maximum of ten (10) virtual servers are installed per physical compute host. Since XenApp instances are easily added to an existing XenDesktop stack, the only additional components required are:

- One or more Windows Server OS instances running the Citrix VDA added to the XenDesktop site

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

4.1.8.2 XenDesktop with XenApp and Personal vDisk Integration

In a XenDesktop implementation that leverages hosted applications, these execute from a centralized Windows Server and are then accessed via the Citrix Receiver. There are some instances, however, where certain departmental or custom applications cannot run using XenApp. At the same time, for organizational policy or certain storage considerations, delivering these applications as a part of a base image is not possible either. In this case, Citrix Personal vDisk technology is the appropriate solution.

With Citrix Personal vDisk, each user of that single shared virtual desktop image also receives a personal layered vDisk, which enables the user to personalize their desktop and receive native application execution within a Windows client OS and not from a server. When leveraging the integration of XenApp within XenDesktop, all profile and user data is seamlessly accessed within both environments.

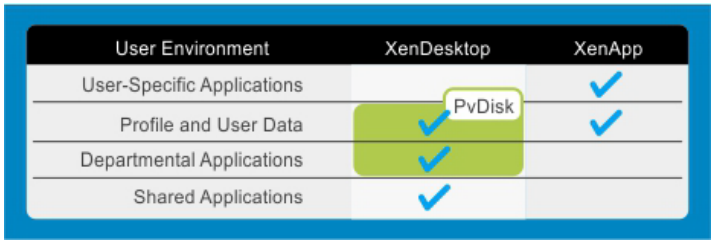

4.1.8.3 PVS Integration with XenApp

One of the many benefits of PVS is the ability to quickly scale the XenApp instances within a farm. Bandwidth is a key consideration and PVS bandwidth utilization is mostly a function of the number of target devices and the portion of the image(s) they utilize. Network impact considerations include:

- PVS streaming is delivered via UDP, yet the application has built-in mechanisms to provide flow control, and retransmission as necessary.
- Data is streamed to each target device only as requested by the OS and applications running on the target device. In most cases, less than 20% of any application is ever transferred.
- PVS relies on a cast of supporting infrastructure services. DNS and DHCP need to be provided on dedicated service infrastructure servers, while TFTP and PXE Boot are functions that may be hosted on PVS servers or elsewhere.

4.1.8.4 NUMA Architecture Considerations

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

The general guidance for XenApp NUMA-alignment on the Dell Wyse Datacenter solution is as follows:

4.1.8.5 NUMA Alignment

20 physical cores per CPU given the 2698v4 part, 40 logical cores with Hyper-Threading active, gives us a total of 80 consumable cores per compute node and falls in line with a 2x oversubscription rate. Configuring the XenApp VMs as shown below (as XA VMs) will ensure that no physical NUMA node spanning occurs which could lower performance for an effected VM.

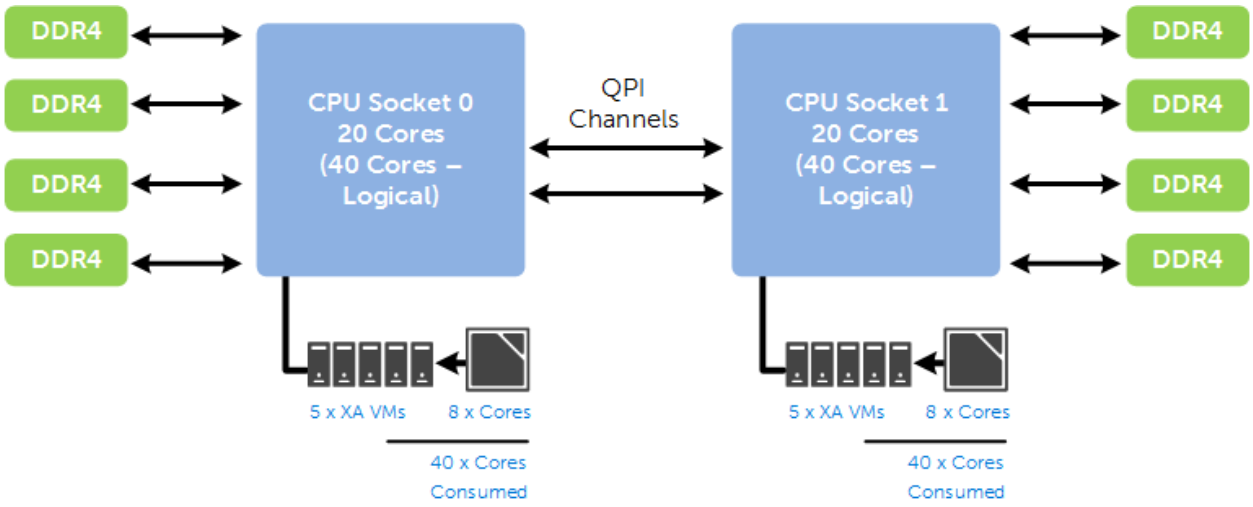

4.1.9 Local Host Cache

In XenApp and XenDesktop version 7.12 and above, the Local Host Cache (LHC) feature allows connection brokering operations to continue when connectivity to the Site database has been interrupted. This includes both failures between the Delivery Controller and Site database in on-premises deployments and when the WAN link between the Site and Citrix control plane fails in a Citrix Cloud environment. LHC replaces the connection leasing feature as the recommended XenApp and XenDesktop high availability solution.

During an outage, LHC will support new users and existing users launching new resources, as well as users accessing pooled resources (shared desktops). Earlier versions of XenApp had a feature named Local Host Cache but this is an entirely different implementation that is more robust and immune to corruption.

The following diagram shows the communication paths during normal operations. The principal broker on a delivery controller accepts requests and communicates with the Site database to connect users. A check is made every two minutes to determine if changes have been made to the principal broker's configuration and if so, the information is synchronized with the secondary broker. All configuration data is copied to ensure the LocalDB database matches the site database.

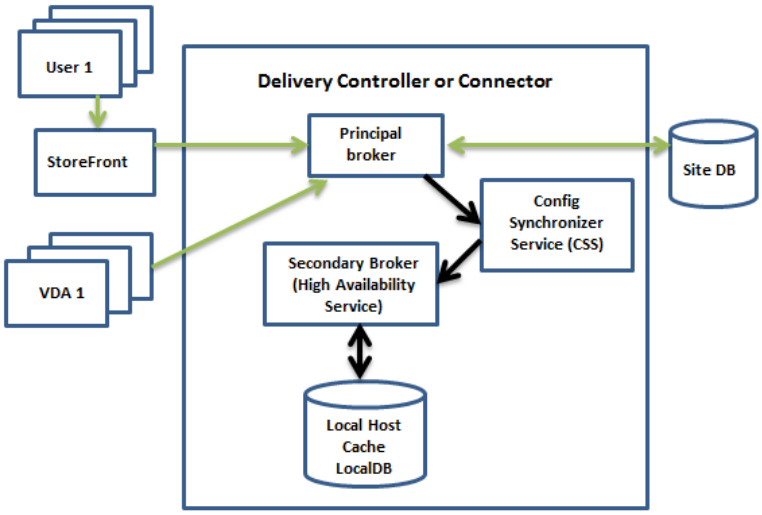

Image provided courtesy of Citrix Systems, Inc.

The following diagram illustrates changes in communication when the principal broker is unable to connect to the Site database.

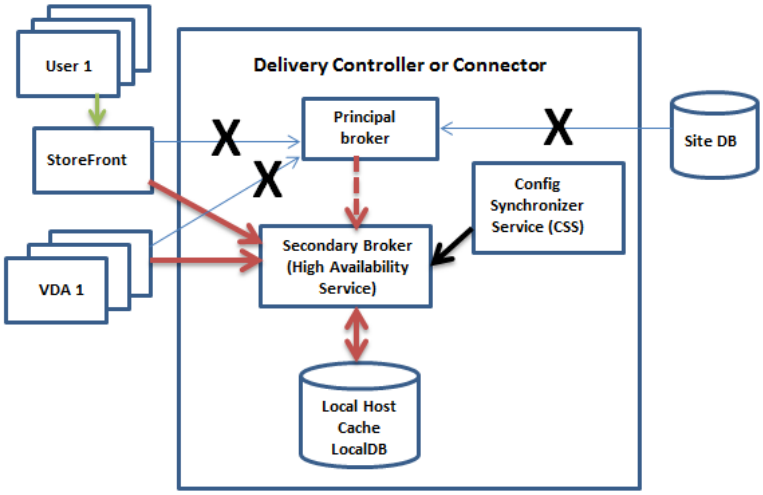

Image provided courtesy of Citrix Systems, Inc.

The principal broker stops listening for requests and instructs the secondary broker to begin listening and processing requests. When a VDA communicates with the secondary broker, a re-registration process is triggered during which current session information is delivered. During this time, the principal broker

continually monitors the connection to the Site database. Once restored, the principal broker resumes brokering operations and instructs the secondary broker to stop listening for connection information.

4.1.10 Citrix NetScaler

Citrix NetScaler is an all-in-one web application delivery controller that makes applications run better, reduces web application ownership costs, optimizes the user experience, and makes sure that applications are always available by using:

- Proven application acceleration such as compression and caching
- High application availability through advanced L4-7 load balancer
- Application security with an integrated application firewall
- Server offloading to significantly reduce costs and consolidate servers

A NetScaler appliance resides between the clients and the servers, so that client requests and server responses pass through it. In a typical installation, virtual servers (vservers) configured on the NetScaler provide connection points that clients use to access the applications behind the NetScaler. In this case, the NetScaler owns public IP addresses that are associated with its vservers, while the real servers are isolated in a private network. It is also possible to operate the NetScaler in a transparent mode as an L2 bridge or L3 router, or even to combine aspects of these and other modes. NetScaler can also be used to host the StoreFront function eliminating complexity from the environment.

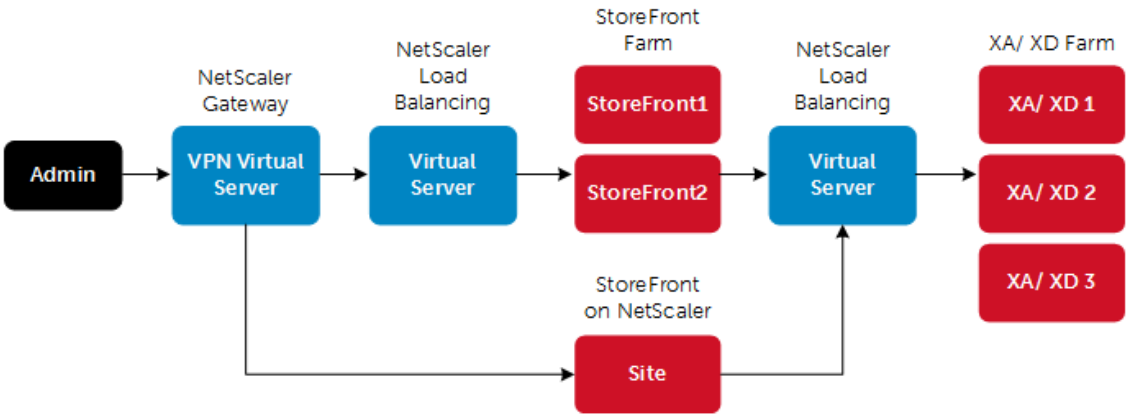

Global Server Load Balancing

GSLB is an industry standard function. It is in widespread use to provide automatic distribution of user requests to an instance of an application hosted in the appropriate data center where multiple processing facilities exist. The intent is to seamlessly redistribute load on an as required basis, transparent to the user community. These distributions are used on a localized or worldwide basis. Many companies use GSLB in its simplest form. They use the technology to automatically redirect traffic to Disaster Recovery (DR) sites on an exception basis. That is, GSLB is configured to simply route user load to the DR site on a temporary basis

only in the event of a catastrophic failure or only during extended planned data center maintenance. GSLB is also used to distribute load across data centers on a continuous load balancing basis as part of normal processing.

NetScaler and XenDesktop Deployment Guide: Link

Several of the management components of the XenDesktop stack are made highly-available using NetScaler to load balance traffic. The following management components require the use of a load balancer to function in a high availability mode:

- StoreFront Servers
- Licensing Server
- XenDesktop XML Service
- XenDesktop Desktop Director
- Provisioning Services TFTP Service
- Framehawk UDP virtual channel (supported on NetScaler Gateway 11.0.62.10 or later and NetScaler Unified Gateway 11.0.64.34 or later)

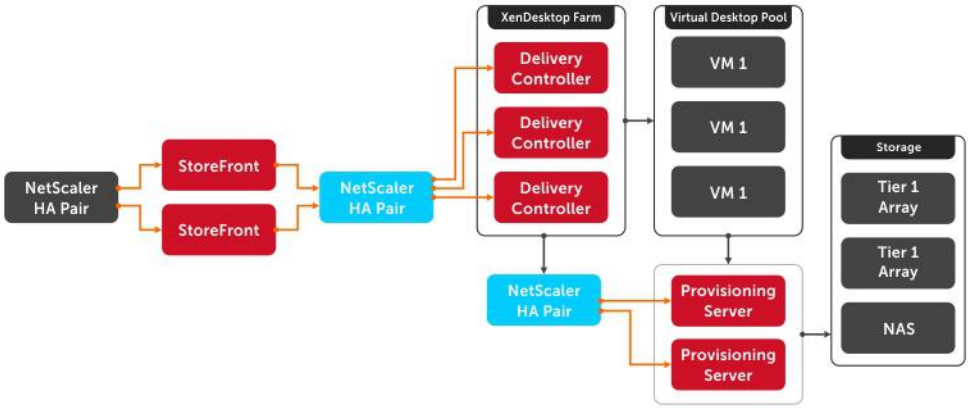

4.2 Hypervisor Platforms

4.2.1 VMware vSphere 6.x

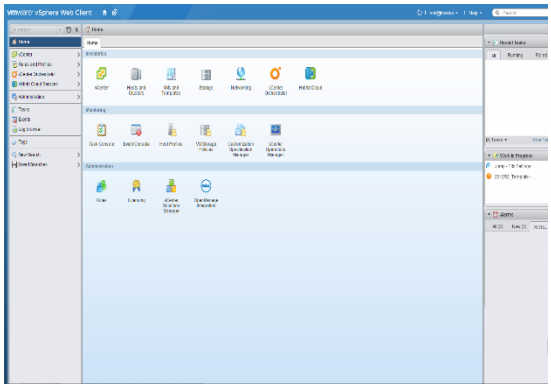

The vSphere hypervisor also known as ESXi is a bare-metal hypervisor that installs directly on top of your physical server and partitions it into multiple virtual machines. Each virtual machine shares the same physical resources as the other virtual machines and they can all run at the same time. Unlike 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 6.x includes three major layers: Virtualization, Management and Interface. The Virtualization layer includes infrastructure and application services. The Management layer is central for configuring, provisioning and managing virtualized environments. The Interface layer includes the vSphere web client.

Throughout the Dell Wyse Datacenter solution, all VMware and Microsoft best practices and prerequisites for core services are adhered to (NTP, DNS, Active Directory, etc.). The vCenter 6.x VM used in the solution is a single VM running Windows Server 2012 R2 or 2016 (version 6.5 and higher - check for current Windows Server OS compatibility at http://www.vmware.com/resources/compatibility) or vCenter 6.x virtual appliance, residing on a host in the management Tier. SQL server is a core component of the Windows version of vCenter and is hosted on another VM also residing in the management Tier. It is recommended that View Composer is installed on a standalone VM running Windows Server 2012 R2 or 2016 when using the vCenter Server Appliance.

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

4.2.2 Microsoft Windows Server Hyper-V

Hyper-V ™ is a powerful virtualization technology that enables businesses to leverage the benefits of virtualization. Hyper-V reduces costs, increases hardware utilization, optimizes business infrastructure, and improves server availability. Hyper-V works with virtualization-aware hardware to tightly control the resources available to each virtual machine. The latest generation of Dell servers includes virtualization-aware CPUs and network adapters. Windows Server 2012 R2 or 2016 can be used with this architecture.

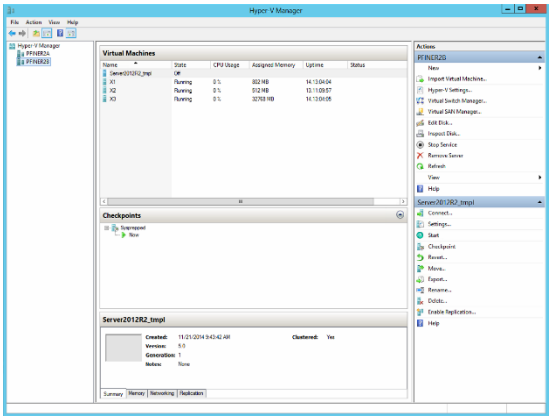

From a network management standpoint, virtual machines are much easier to manage than physical computers. To this end, Hyper-V includes many management features designed to make managing virtual machines simple and familiar, while enabling easy access to powerful VM-specific management functions. The primary management platform

within a Hyper-V based XenDesktop virtualization environment is Microsoft Systems Center Virtual Machine Manager (SCVMM).

SCVMM provides centralized and powerful management, monitoring, and self-service provisioning for virtual machines. SCVMM host groups are a way to apply policies and to check for problems across several VMs at once. Groups are organized by owner, operating system, or by custom names such as “Development” or “Production”. The interface also incorporates Remote Desktop Protocol (RDP); double-click a VM to bring up the console for that VM—live and accessible from the management console.

4.3 NVIDIA GRID vGPU

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

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

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

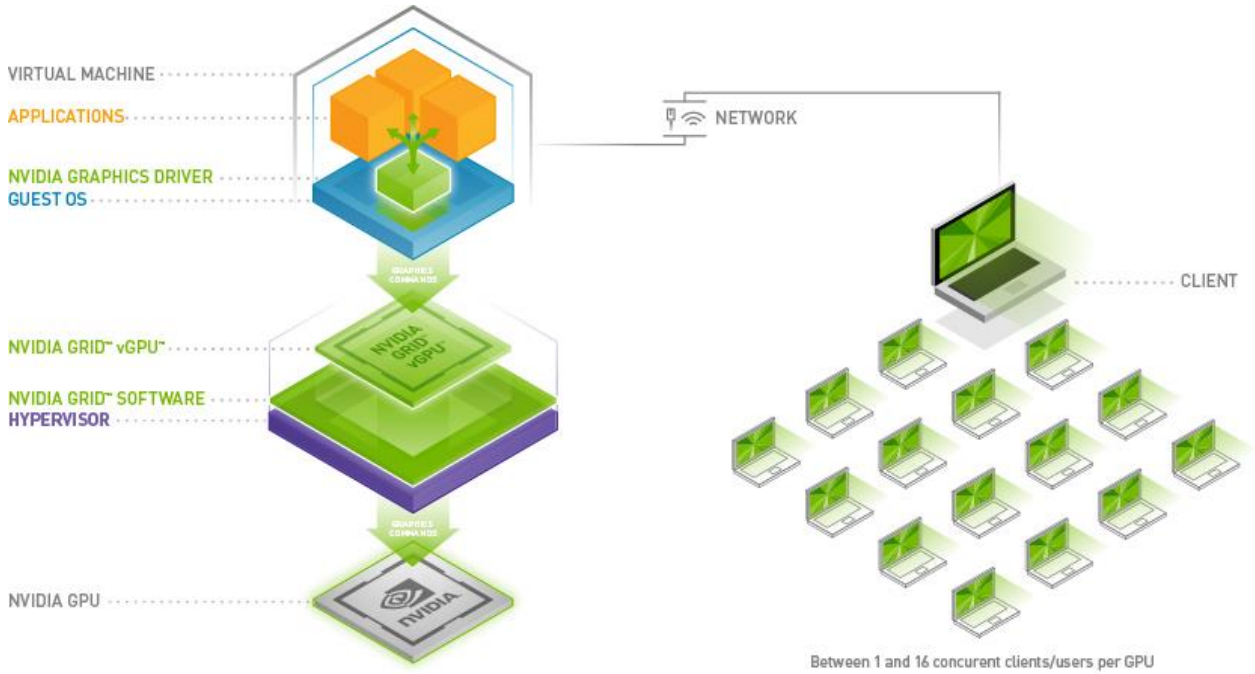

Image provided curtesy of NVIDIA Corporation, Copyright NVIDIA Corporation

4.3.1 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. Every virtual desktop has dedicated graphics memory so they always have the resources they need to launch and run their applications at full performance. 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.

4.3.1.1 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 want 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.

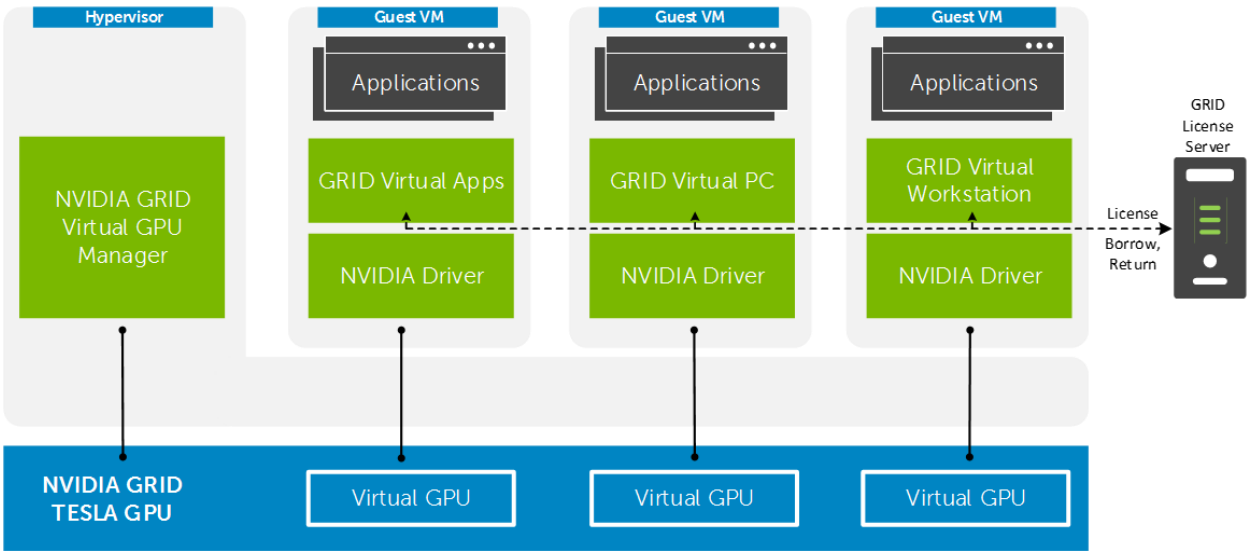

Image provided courtesy of NVIDIA Corporation, Copyright NVIDIA Corporation

5 Solution Architecture for XenDesktop

5.1 Management Role Configuration

5.1.1 vSphere

The Management role requirements for the base solution are summarized below and take into account the use of PVS. If MCS provisioning is desired remove this from your build plan. Depending on the scale of the deployment, the number of these services must be adjusted. Use data disks for role-specific application files such as data, logs and IIS web files in the Management volume.

Role	vCPU	RAM (GB)	NIC	OS + Data vDisk (GB)	Tier 2 Volume (GB)
DDC + Lic	4	8	1	60	-
Storefront	2	4	1	60	-
vCenter	8	24	1	60	50 (VMDK)
PVS	4	24	1	60	100 (VMDK)
SQL Server	4	8	1	60	200 (VMDK)
File Server	2	4	1	60	2048 (RDM)
VSA (optional)	2	12	6	85	100 (VMDK)
TOTALS	**26**	**84**	**12**	**445**	**2,498**

5.1.2 NVIDIA GRID License Server Requirements

When using NVIDIA Tesla M60 or M10 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	-

At the time of this writing, 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 Hyper-V

Role	vCPU	Startup RAM (GB)	Dynamic Memory			NIC	OS + Data vDisk	
			Min|Max	Buffer	Weight		Size (GB)	Tier2 Vol
DDC + Lic	4	8	512MB | 10GB	20%	Med	1	60	-
Storefront	2	4	512MB | 6GB	20%	Med	1	60	-
SCVMM	8	8	512MB | 16GB	20%	Med	1	60	50 (VHDX)
PVS	4	24	512MB | 24GB	20%	Med	1	60	200 (VHDX)
Primary SQL	8	8	512MB | 16GB	20%	Med	1	60	300 (VHDX)
File Server	2	4	512MB | 6GB	20%	Med	1	60	2048 (PTM)
Total	**28**	**56GB**	**3GB | 78GB**	**-**	**-**	**6**	**360GB**	**2598GB**

5.1.4 XenApp VM Configuration

The recommended number of XenApp VMs and their configurations on vSphere or Hyper-V are summarized below and take into account proper NUMA balancing assuming the CPU in use is the E5-2698v4. For more information on NUMA please refer to the NUMA Architecture Considerations section.

XenApp VM configuration on Hyper-V

Role	VMs per host	vCPUs per VM	Startup RAM (GB)	Dynamic Memory			NIC	OS vDisk	
				Min|Max	Buffer	Weight		Size (GB)	Location
XenApp VM	10	8	16	8GB | 48GB	20%	Med	1	80	Tier 1

XenApp VM configuration on vSphere

Role	VMs per host	vCPUs per VM	Memory	NIC	OS vDisk	
					Size (GB)	Location
XenApp VM	10	8	48GB	1	80	Tier 1

5.1.5 SQL Databases

The Citrix, Microsoft and VMware databases are hosted by a single dedicated SQL 2012 R2 (or higher) Server VM in the Management layer. Use caution during database setup to ensure that SQL data, logs, and TempDB are properly separated onto their respective volumes. Create all Databases that are required for:

- Citrix XenDesktop
- vCenter or SCVMM

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 Citrix, Microsoft and VMware are to be adhered to, to ensure optimal database performance.

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

5.1.6 DNS

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

Pay consideration for eventual scaling, access to components that may live on one or more servers (SQL databases, Citrix 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.

5.1.6.1 DNS for SQL

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

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

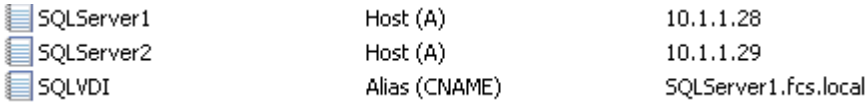

5.2 Storage Architecture Overview

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 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 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/ SCVMM, XD roles, File & SQL	VMFS/ NTFS
PVS	100	Tier 2	PVS vDisks	VMFS/ NTFS
User Data	2048	Tier 2	File Server/ NAS	RDM/PTD
User Profiles	20	Tier 2	User profiles	VMFS/ NTFS
SQL DATA	100	Tier 2	SQL	VMFS/ NTFS
SQL LOGS	100	Tier 2	SQL	VMFS/ NTFS
TempDB Data	5	Tier 2	SQL	VMFS/ NTFS
TempDB Logs	5	Tier 2	SQL	VMFS/ NTFS
Templates/ISO	200	Tier 2	ISO storage (optional)	VMFS/ NTFS

5.2.4 Storage Networking – XtremIO Fiber 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.

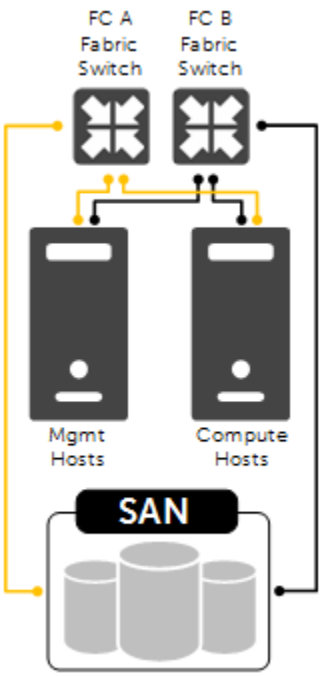

XtremIO actively manages data at a block level using real-time intelligence, providing fully virtualized storage at the disk level. Resources are pooled across the entire storage array. All virtual volumes are thin-provisioned. With inline data compression and dedupe, physical storage requirements can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

5.3.1.1 vSphere

dvSwitches should be used as desired for VM traffic especially in larger deployments to ease the management burden across numerous hosts. In the LT1 rack model, only the mgmt. hosts connect to shared storage so require additional VMK ports. Network share values should be configured equally among the VMKernel port groups that share a physical set of network adapters.

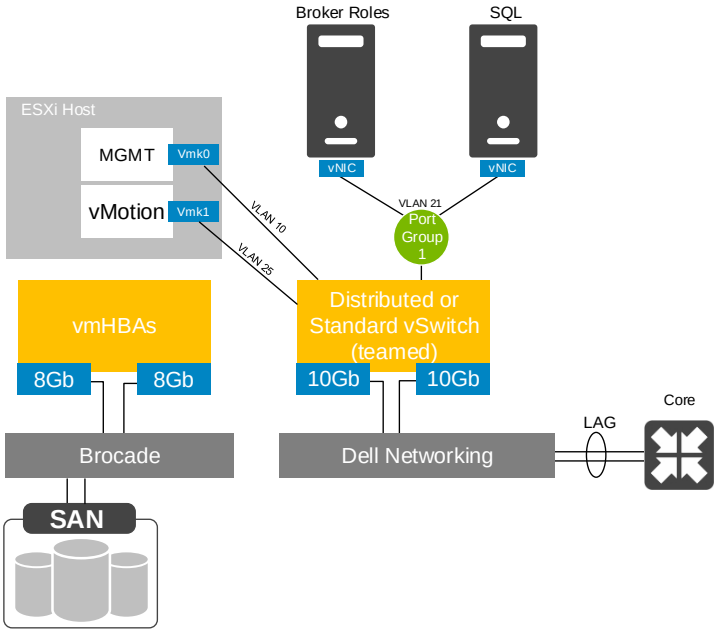

The Compute hosts are configured in the same basic manner, minus the shared storage, with the desktop VMs connecting to the primary port group on the external vSwitch.

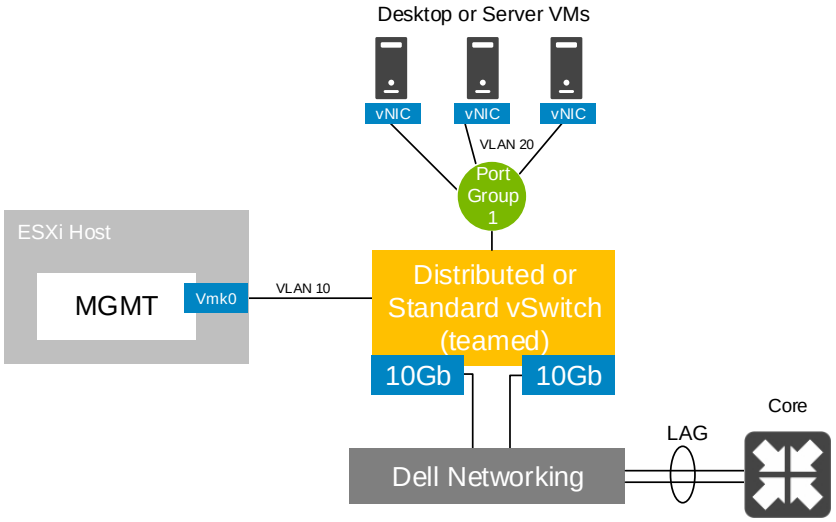

5.3.1.2 Hyper-V

The Hyper-V configuration, while identical in core requirements and hardware, is executed differently due to how Hyper-V and Windows Server implement networking and virtual switches. As shown in the diagram below, native Windows Server NIC Teaming is utilized to load balance and provide resiliency for network connections. For the compute host in this scenario, a single LBFO NIC team is configured to connect to a Hyper-V switch for external traffic. All vNICs associated with the Management OS connect directly to the external Hyper-V switch.

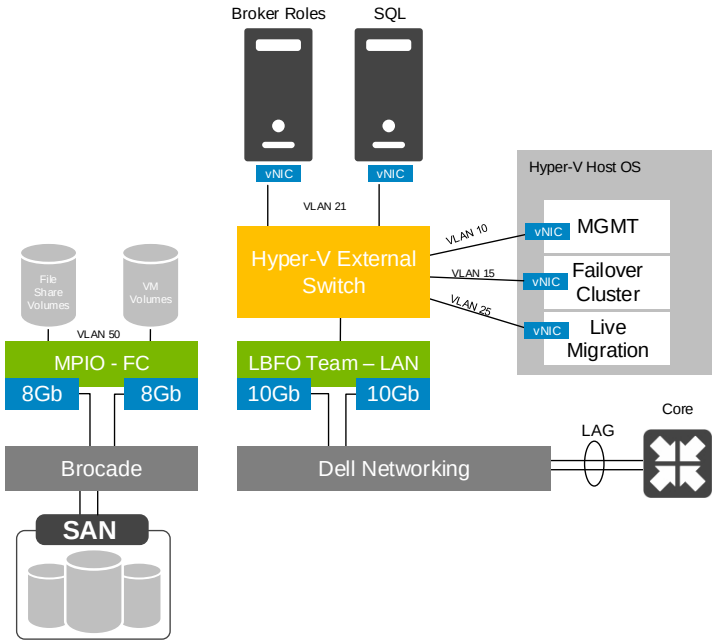

The dedicated compute host configuration is shown in the diagram below and configured similarly to the management host configuration but without the features enabled by shared storage.

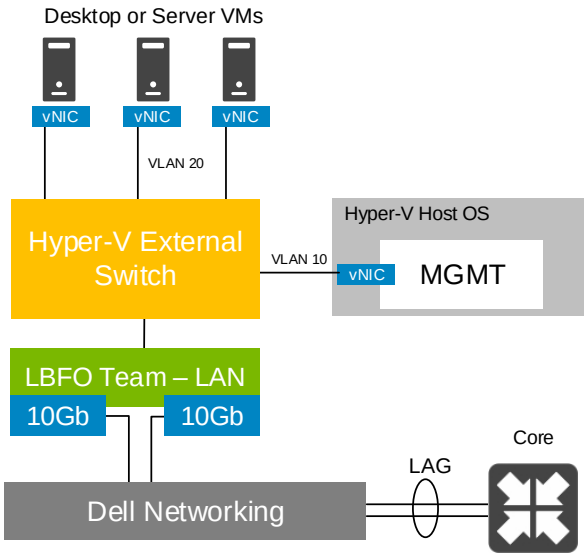

5.3.2 Shared Tier 1 – FC

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.

5.3.2.1 vSphere

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

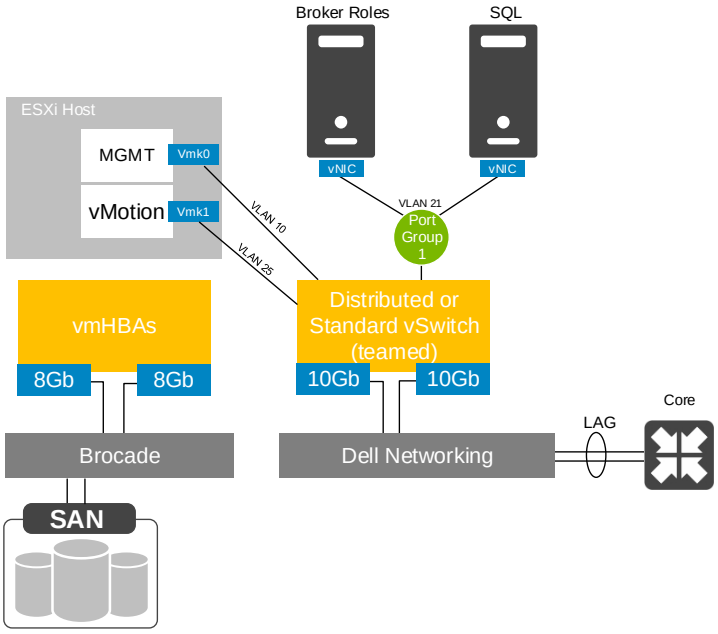

Compute Servers

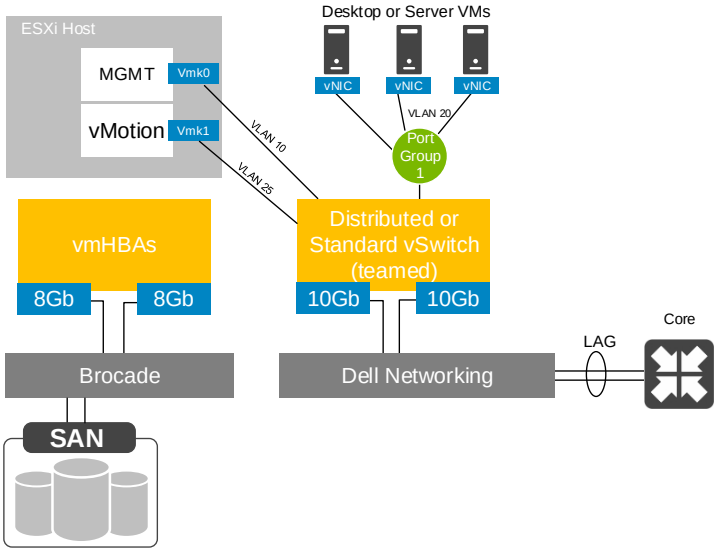

5.3.2.2 Hyper-V

The Hyper-V configuration, while identical in core requirements and hardware, is executed differently due to how Hyper-V and Windows Server implement networking and virtual switches. As shown in the diagram below, native Windows Server NIC Teaming is utilized to load balance and provide resiliency for network connections. For the compute or management hosts in this scenario, a single LBFO NIC team is configured to connect to a Hyper-V switch for external traffic. All vNICs associated with the Management OS connect directly to the external Hyper-V switch with MPIO used to connect to shared storage.

Management Servers

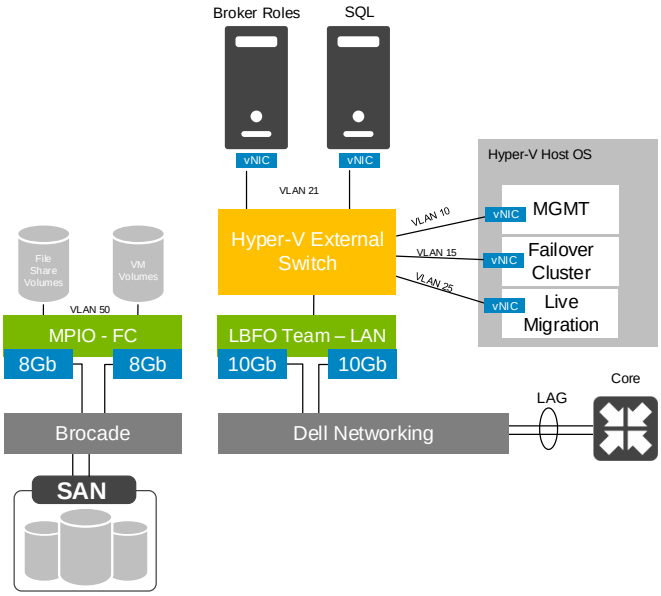

Compute Servers

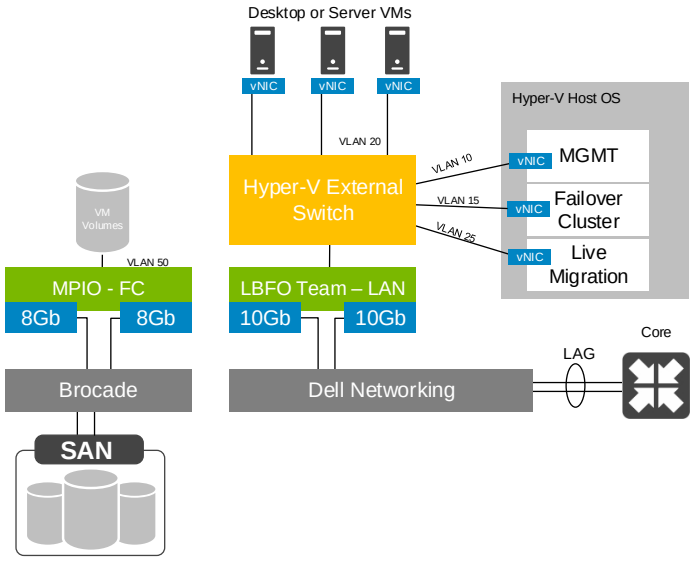

5.4 Scaling Guidance

- 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
Compute Servers	Desktop VMs per physical host based on available CPU	Additional hosts and clusters added as necessary	Additional RAM or CPU compute power
Mgmt Servers	Number of server VMs per host	Add additional hosts	Add RAM or network adapters
Provisioning Servers	Desktops per instance	Additional servers added to the Provisioning Server farm	Additional network and I/O capacity added to the servers
Desktop Delivery Servers	Desktops per instance (dependent on SQL performance as well)	Additional servers added to the XenDesktop Site	Additional virtual machine resources (RAM and CPU)
XenApp Servers	Desktops per instance	Additional virtual servers added to the XenDesktop Site	Additional physical servers to host virtual XenApp servers.
Storefront Servers	Logons/ minute	Additional servers added to the Storefront environment	Additional virtual machine resources (RAM and CPU)
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.

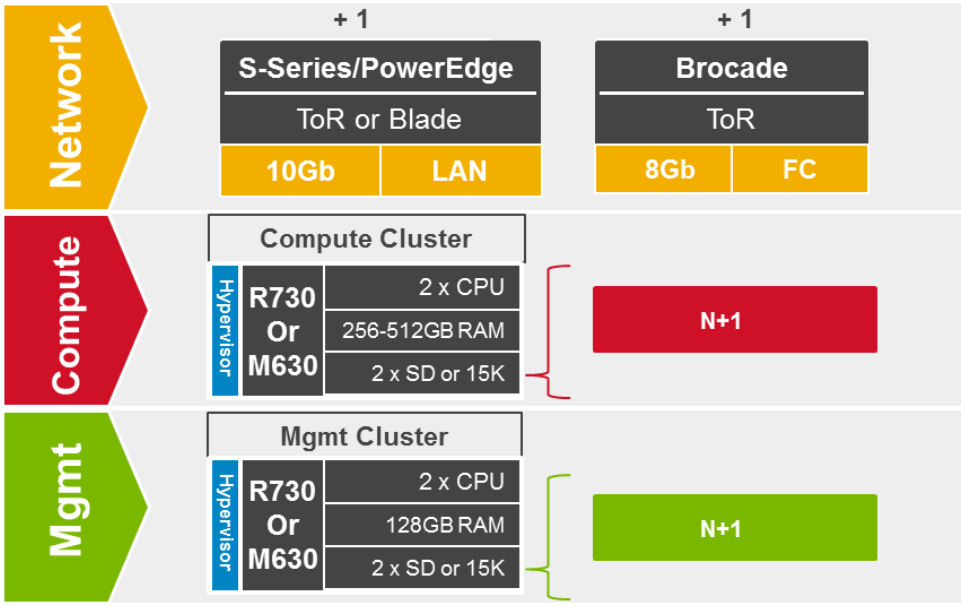

The HA options provide redundancy for all critical components in the stack while improving the performance and efficiency of the solution as a whole.

- Additional switches added to the existing thereby equally spreading each host's network connections across multiple switches.
- Additional ESXi or Hyper-V hosts added in the compute or management layers to provide N+1 protection.
- Applicable Citrix infrastructure server roles are duplicated and spread amongst management host instances where connections to each are load balanced via the addition of virtual NetScaler appliances.
- Each PVS Server instance is responsible for the Write Cache of each desktop that it is hosting. This Write Cache is readable by the resilient Provisioning Server and as it is held on the target device. In the event of a Provisioning Server failure, all desktops that were hosted will transfer to an alternate provisioning server for that site and users will be unaware of the failure.
- SQL Server databases also are protected through the addition and configuration of an "AlwaysOn" Failover Cluster Instance or Availability Group.

5.5.1 SQL Server High Availability

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

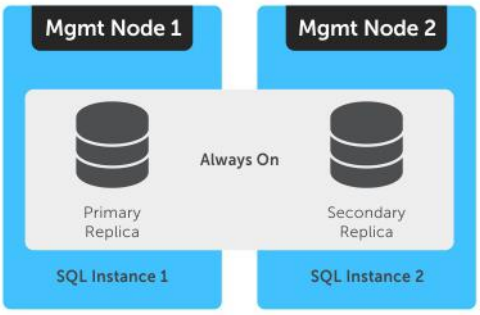

Please refer to these links for more information: LINK1 and LINK2

5.6 Citrix XenDesktop Communication Flow

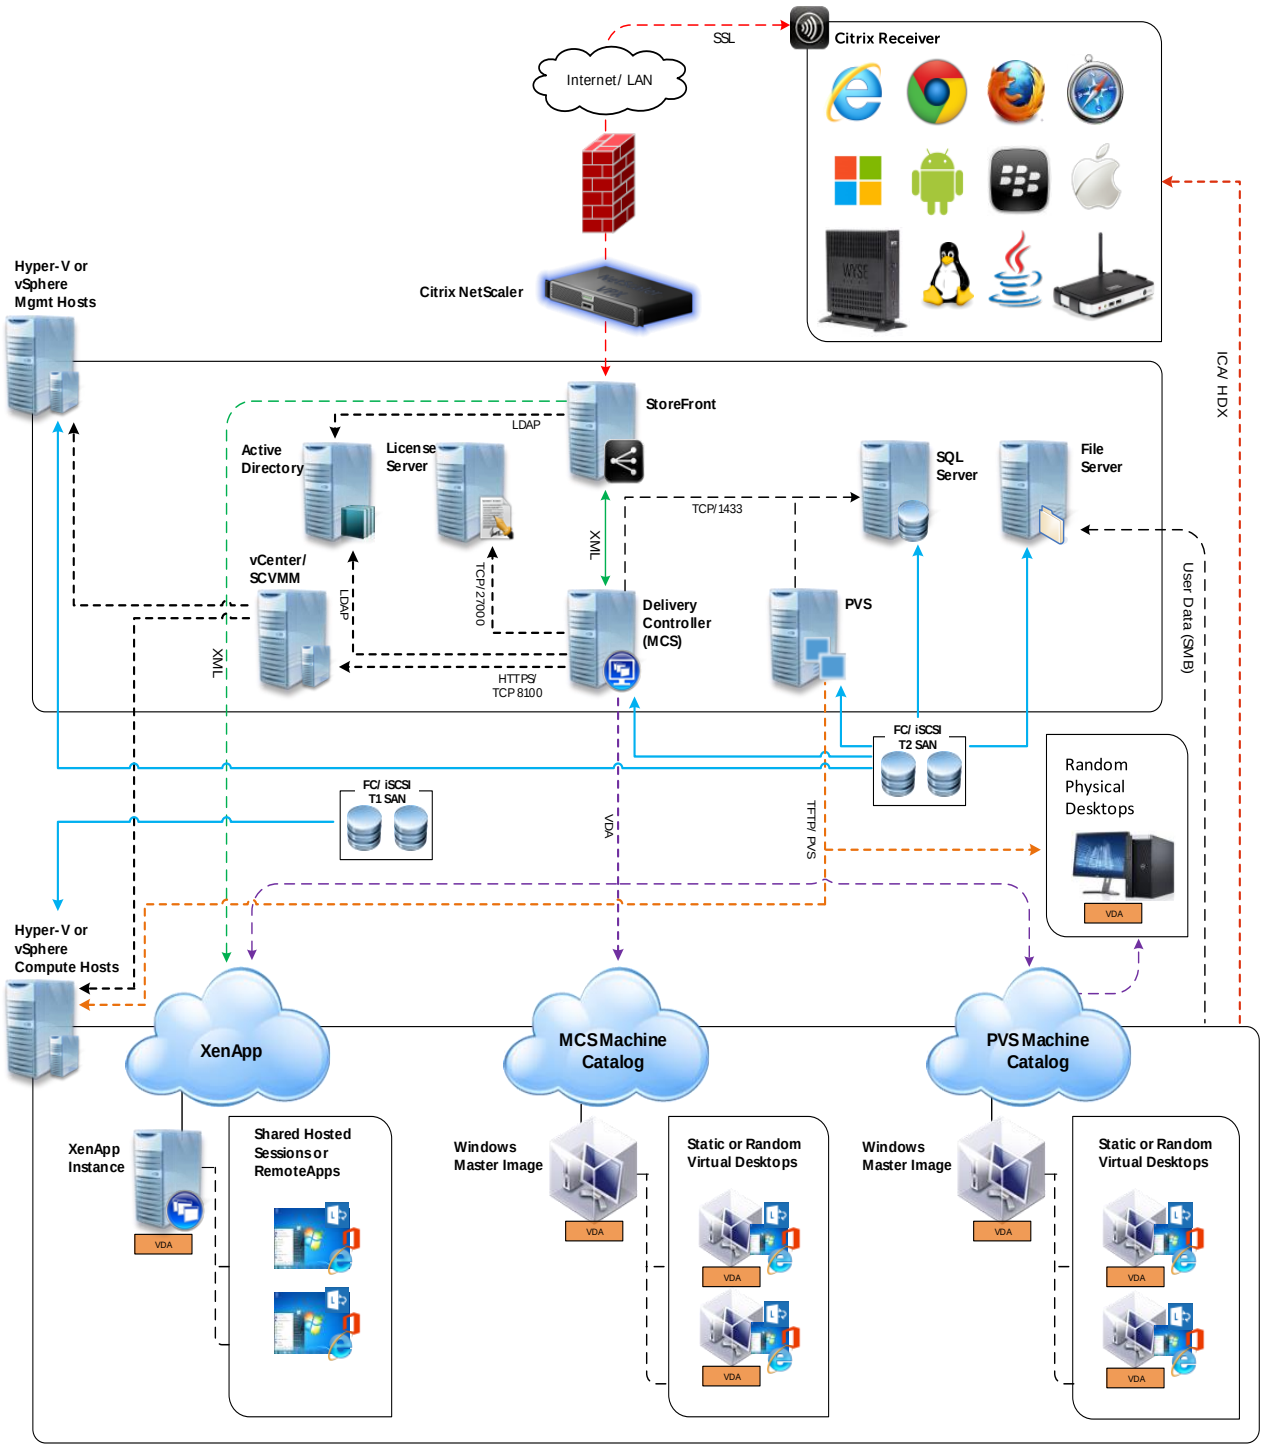

6 Customer-provided solution components

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 user density recommendations per compute server. Please refer to the Platform Configurations section for hardware specifications.

User density summary

Host Config	Hypervisor	Broker & Provisioning	Workload	Template	User Density
R730	ESXi 6.0	XenDesktop 7.x MCS	Task Worker (Light)	Windows 8.1 x64 & Office 2010	330*
R730	ESXi 6.0	XenDesktop 7.x MCS	Knowledge Worker (Medium)	Windows 8.1 x64 & Office 2010	240*
R730	ESXi 6.0	XenDesktop 7.x MCS	Power Worker (Heavy)	Windows 8.1 x64 & Office 2010	180*
R730	ESXi 6.0 U2	XenDesktop 7.x MCS	Task Worker	Windows 10 x64 & Office 2016	270
R730	ESXi 6.0 U2	XenDesktop 7.x MCS	Knowledge Worker	Windows 10 x64 & Office 2016	180
R730	ESXi 6.0 U2	XenDesktop 7.x MCS	Power Worker	Windows 10 x64 & Office 2016	150
R730	Hyper-V 2012 R2	XenDesktop 7.x MCS	Task Worker	Windows 10 x64 & Office 2016	300
R730	Hyper-V 2012 R2	XenDesktop 7.x MCS	Knowledge Worker	Windows 10 x64 & Office 2016	220
R730	Hyper-V 2012 R2	XenDesktop 7.x MCS	Power Worker	Windows 10 x64 & Office 2016	190
R730	Hyper-V 2016	XenDesktop 7.12 MCS	Knowledge Worker	Windows 10 x64 & Office 2016	200
R730	Hyper-V 2016	XenApp 7.12 RDSH	Knowledge Worker	Windows 2016 x64 & Office 2016	300

*Differences in density due to older version Login VSI, Windows OS version, and MS Office version.

Note: Prior Windows 8.1 density included for comparison only - test result data not included. Windows 8.1 testing performed with Office 2010, low screen resolution, and Login VSI 4.0 workloads.

The 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 multiple runs. First, a pilot run to validate that the infrastructure is functioning and valid data can be captured, and then, 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 Microsoft Performance Monitor

Microsoft Performance Monitor is used for Hyper-V 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.

7.1.2.3 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:

Hardware and software test components

Component	Description/Version
Hardware platform(s)	PowerEdge R730
Hypervisor(s)	ESXi 6.0, ESXi 6.0 U2, Hyper-V 2012 R2, and Hyper-V 2016
Broker technology	XenDesktop 7.9 & 7.12
Broker database	Microsoft SQL 2012 and 2016
Management VM OS	Windows Server 2012 R2 and 2016
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	Hyper-V Startup Memory	Hyper-V Min | Max Dynamic	Operating System
Task Worker	1	2GB	1GB	1GB	1GB | 2GB	Windows 10 Enterprise 64-bit
Knowledge Worker	2	3GB	1.5GB	1.5GB	1GB | 3GB	Windows 10 Enterprise 64-bit
Power Worker	2	4GB	2GB	2GB	1GB | 4GB	Windows 10 Enterprise 64-bit
Graphics LVSI Power + ProLibrary	2	4 GB	4GB			Windows 10 Enterprise 64-bit
Graphics LVSI Custom – Density	2	4 GB	4GB			Windows 10 Enterprise 64-bit

Graphics LVSI Custom - Performance	4	8GB	8GB			Windows 10 Enterprise 64-bit

Screen resolutions

User Profile	Screen Resolution
Task Worker	1280 X 720
Knowledge Worker	1920 X 1080
Power Worker	1920 X 1080
Graphics LVSI Power + ProLibrary	1920 X 1080
Graphics LVSI Custom – Density	1920 X 1080
Graphics LVSI Custom - Performance	1920 X 1080

RDSH VM specifications

Platform Config	vCPUs	Hyper-V Startup Memory	Hyper-V Min | Max Dynamic	Operating System
XenApp 7.12	8	16GB	8GB | 48GB	Windows Server 2016

7.2.2 Platform Configurations

The hardware configurations that were tested are summarized in the table(s) below.

R730 Compute hardware configuration

Platform Config	CPU	Memory	RAID Ctlr	HD Config	Network
R730	E5-2698v4 (20 Core, 2.2GHz)	512GB @2400 MT/s	PERC H730, 2GB Cache	4 X 800GB, Intel S3710, 6GB/s SATA SSD‘s 2.5” (RAID 10)	2 x 1Gb Broadcom 2P 5720 adapter 2 X 10GbE Broadcom 2P 57810 adapter

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.

Test result summary

Platform Config	Hypervisor	Broker & Provisioning	Login VSI Workload	Density per Host	Avg CPU	Avg Mem Consumed	Avg Mem Active	Avg IOPS / User	Avg Net Mbps / User
R730	ESXi 6.0 U2	XenDesktop 7.9 MCS	Task Worker	270	99%	499 GB	254 GB	11.07	1.064
R730	ESXi 6.0 U2	XenDesktop 7.9 MCS	Knowledge Worker	180	90%	505 GB	218 GB	15.28	2.341
R730	ESXi 6.0 U2	XenDesktop 7.9 MCS	Power Worker	150	99%	504 GB	283 GB	13.67	4.512
R730	Hyper-V 2012 R2	XenDesktop 7.9 MCS	Task Worker	300	60%	502GB		8.3	1.1
R730	Hyper-V 2012 R2	XenDesktop 7.9 MCS	Knowledge Worker	220	75%	502GB		13.6	2
R730	Hyper-V 2012 R2	XenDesktop 7.9 MCS	Power Worker	190	80%	462GB		15	5
R730	Hyper-V 2016	XenDesktop 7.12 MCS	Knowledge Worker	200	60%	487GB		15	2
R730	Hyper-V 2016	XenApp 7.12 RDSH	Knowledge Worker	300	66%	428GB		<1	1.6

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 AHV/ESXi host CPU metrics will exceed the rated 100% for the host if Turbo Boost is enabled (by default). An additional 35% of CPU is available from the Turbo Boost feature but this additional CPU headroom is not reflected in the metrics where the performance data is gathered. Therefore, CPU usage for AHV/ESXi hosts is adjusted and a line indicating the potential performance headroom provided by Turbo boost is included in each CPU graph.

Avg Consumed Memory: ESXi consumed memory is the amount of host physical memory consumed by a host. For AHV, this is memory consumed expressed as a percentage of total memory. 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. This metric is not available on AHV hosts.

7.3.1 R730 Compute

Refer to the Platform Configurations section for hardware configuration details.

7.3.1.1 Task Worker, 270 Users, ESXi 6.0 U2, XenDesktop 7.9 MCS

The below graph shows the performance data for 270 user sessions on the host The CPU reaches a steady state average of 100% during the test cycle when all 270 users are logged on. However, the core utilization did not exceed 95%.

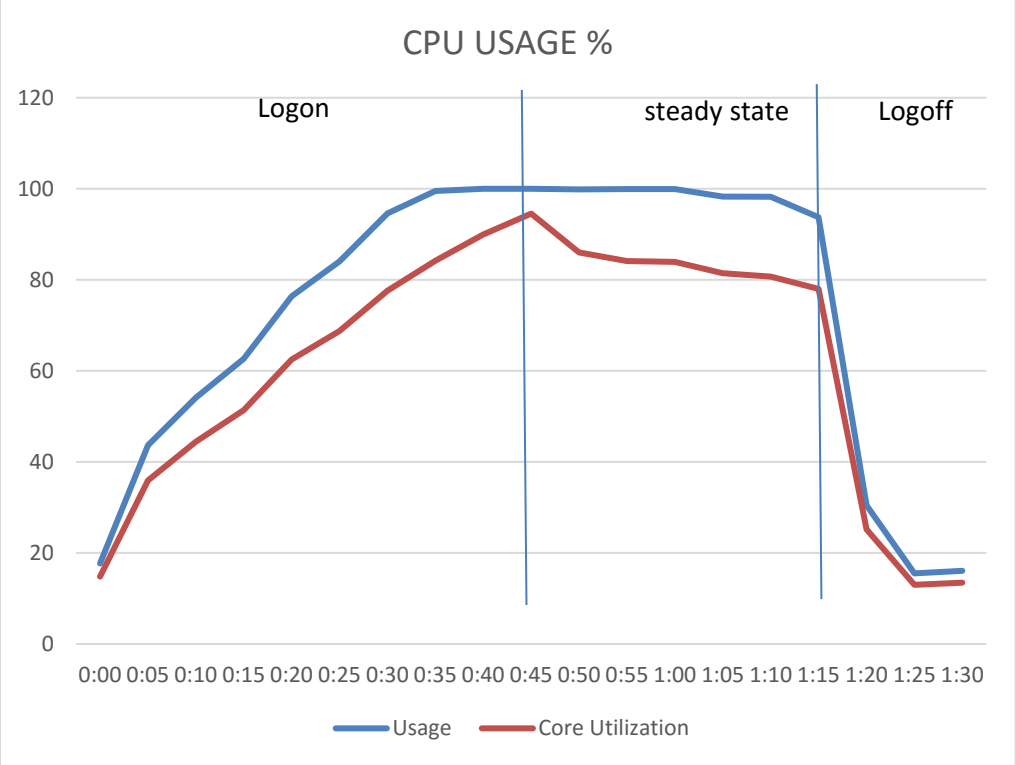

In regards to memory consumption for the cluster, out of a total of 512 GB available memory per node there were no constraints for any of the hosts. The compute hosts consumed almost all the available memory on host with active memory usage reaching a max around 300GB. There was no memory ballooning or swapping on the hosts.

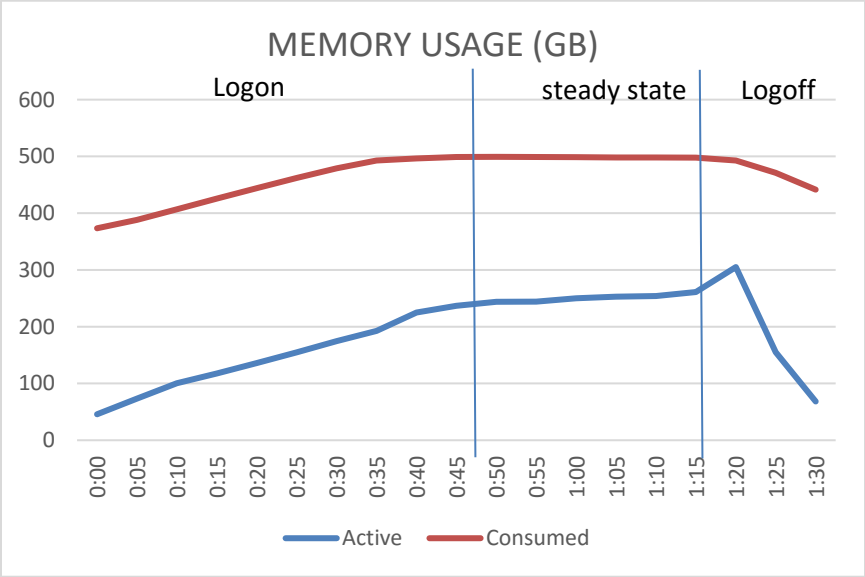

Network bandwidth in this test run with a peak at just over 1Mbps per user, less than 300Mbps per host. The host has pair of teamed 1Gbps NIC, so the bandwidth is not under pressure of saturation.

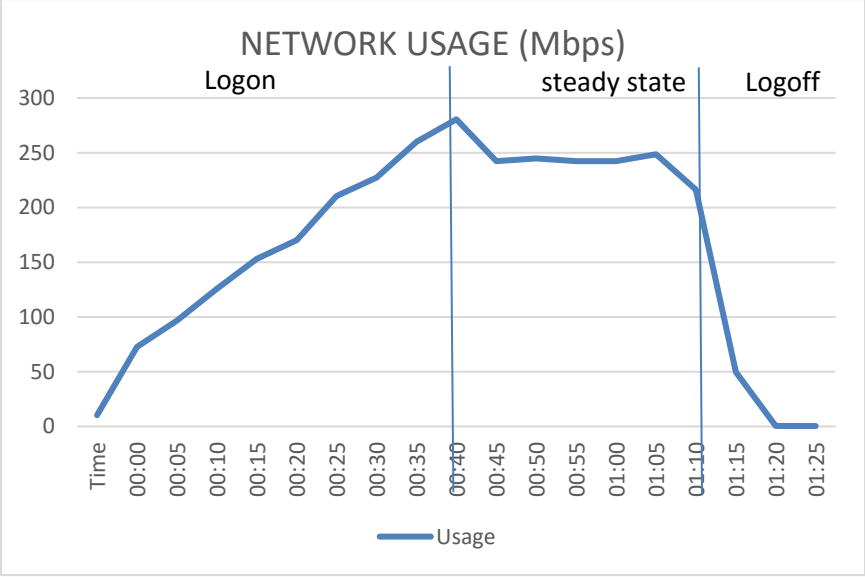

The data store IOPS is shown in below graph, the read IO is higher than Write as expected. And peaked when all users logged on. Average IO Max around 11IOPS, is in line with LoginVSI expectation. The read and write latency remained under 2ms,

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.

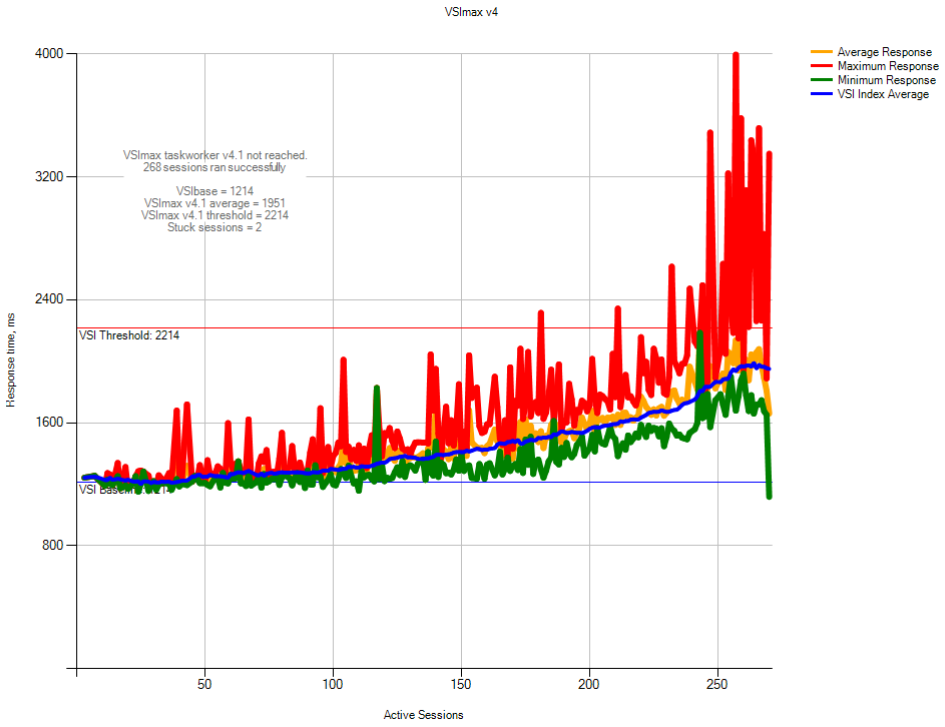

Notes:

- User profile management is configured using Citrix XenDesktop profile management policy in this test. If using other profile management methods, such as Active directory roaming profile may impact logon time and IO, hence impact the performance.

- There were two session failed to complete test cycle. Upon investigation, it's concluded the failure was not related to host performance reason. And since this is the only two sessions out of 270 sessions, it is below the 2% acceptable results threshold. The LoginVSI test "VSImax not reached" result is considered valid.

7.3.1.2 Knowledge Worker, 180 Users, ESXi 6.0 U2, XenDesktop 7.9 MCS

The below graph shows the CPU usage during the test. The CPU usage reaches 100% when all users logged on but the core utilization did not exceed 90%.

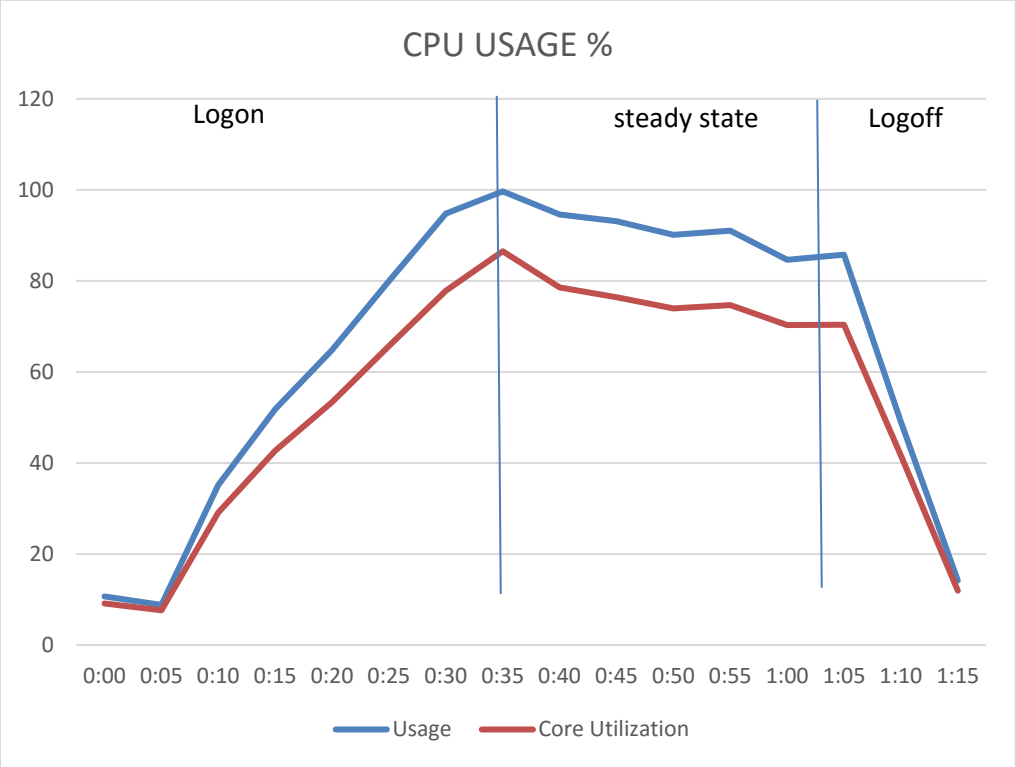

In terms of host memory utilization. Almost all 512GB available host memory were consumed. Though only less than half were active at 218GB. There was no memory swapping or ballooning during the test run.

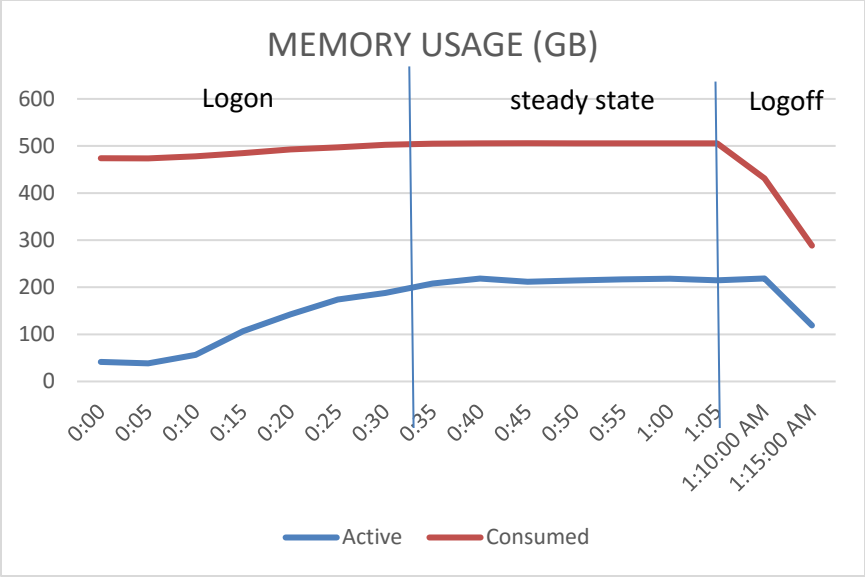

Network bandwidth is not an issue on this test run with a steady state peak of approximately 2Mbps per user. Just over 400Mbps over all on the host.

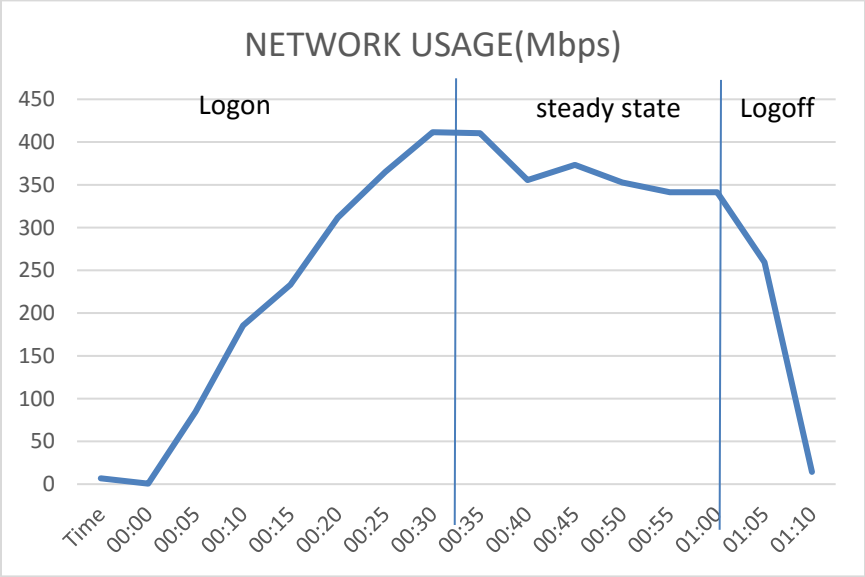

The graph shows the datastore read and write IOPS on the host's local storage. The average IO per user maximum at 15 IOPS.

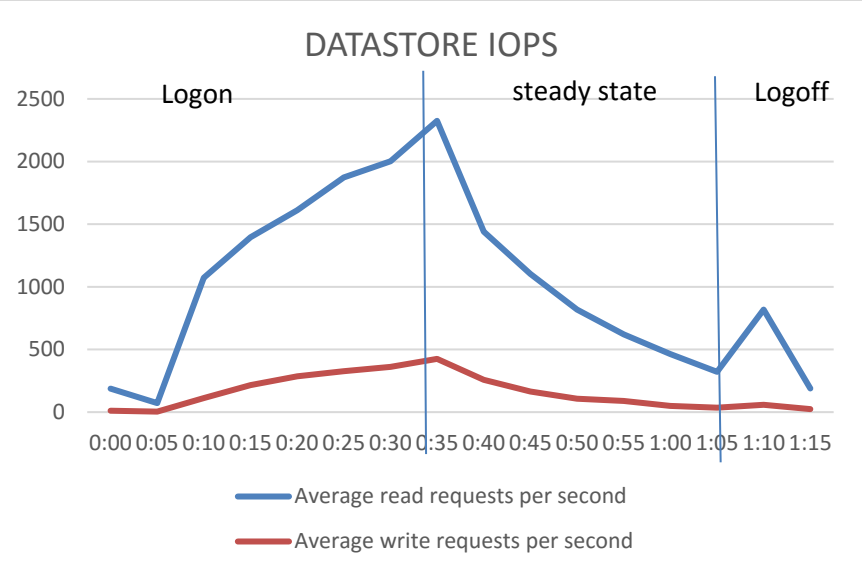

Again storage latency was not an issue. The maximum latency reached was approximately 2ms during a spike in the steady state period of test. This was well below the 20ms threshold that is regarded as becoming potentially troublesome.

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.

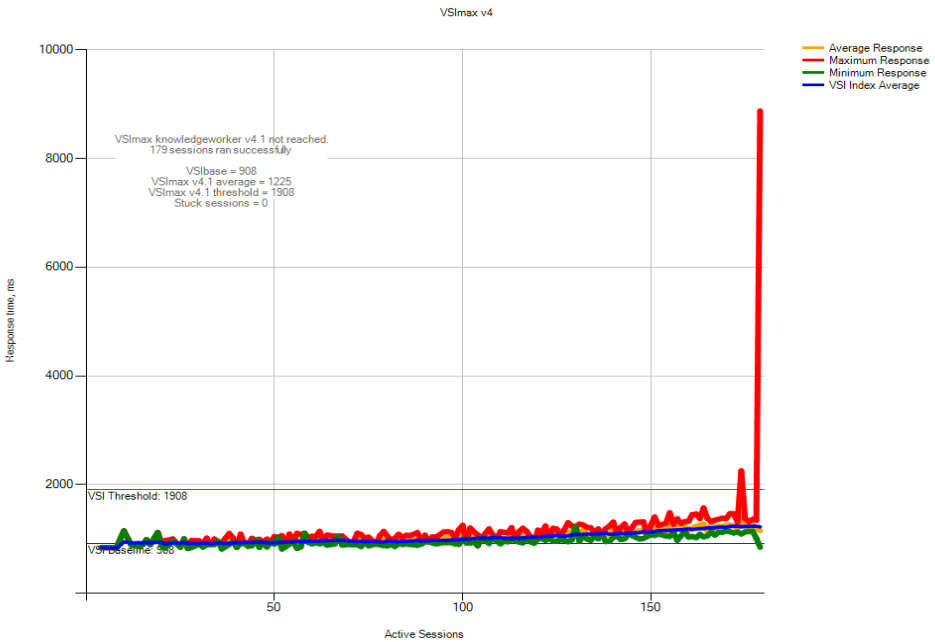

Notes:

There was one session failed to complete test cycle. Upon investigation, it's concluded the failure was not related to host performance reason. And since this is the only session out of 150 sessions, it is below the 2% acceptable results threshold. The LoginVSI test result is considered valid.

7.3.1.3 Knowledge Worker, 200 Users, Hyper-V 2016, XenDesktop 7.12 MCS

The below graph shows the performance data for 200 Knowledge Worker sessions per host. The CPU reaches a steady state average of 60% during the test cycle when 200 users are logged on to the compute host.

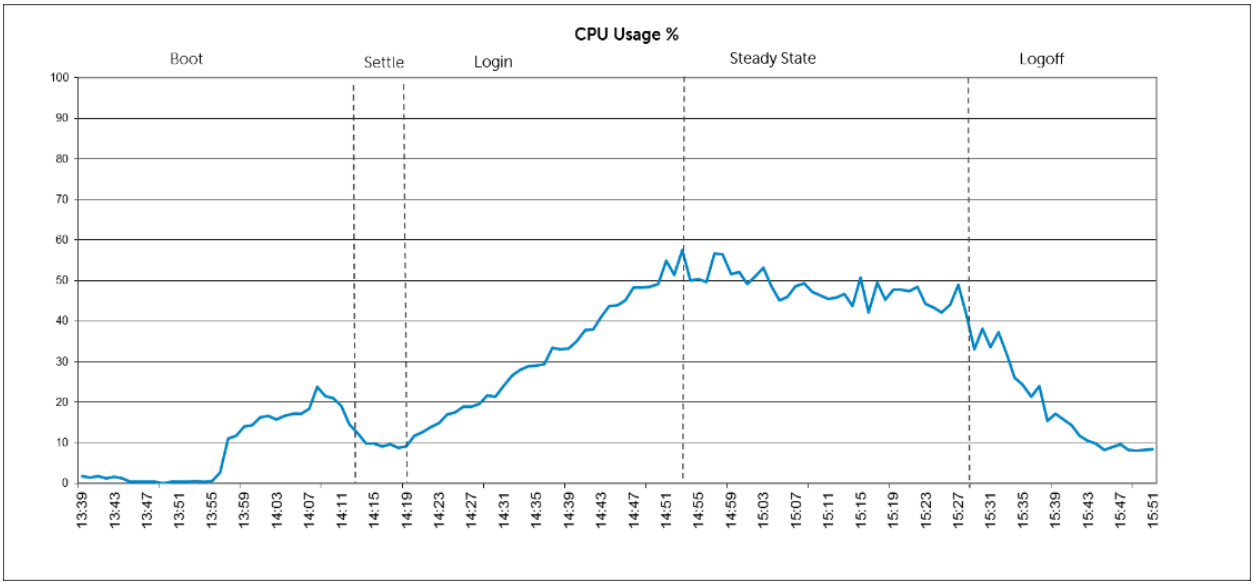

The memory available on the compute host was 512GB. Nearly all the memory was consumed by the VDI VMs. It can be seen that although CPU and IO latency are low the system is limited by the memory available.

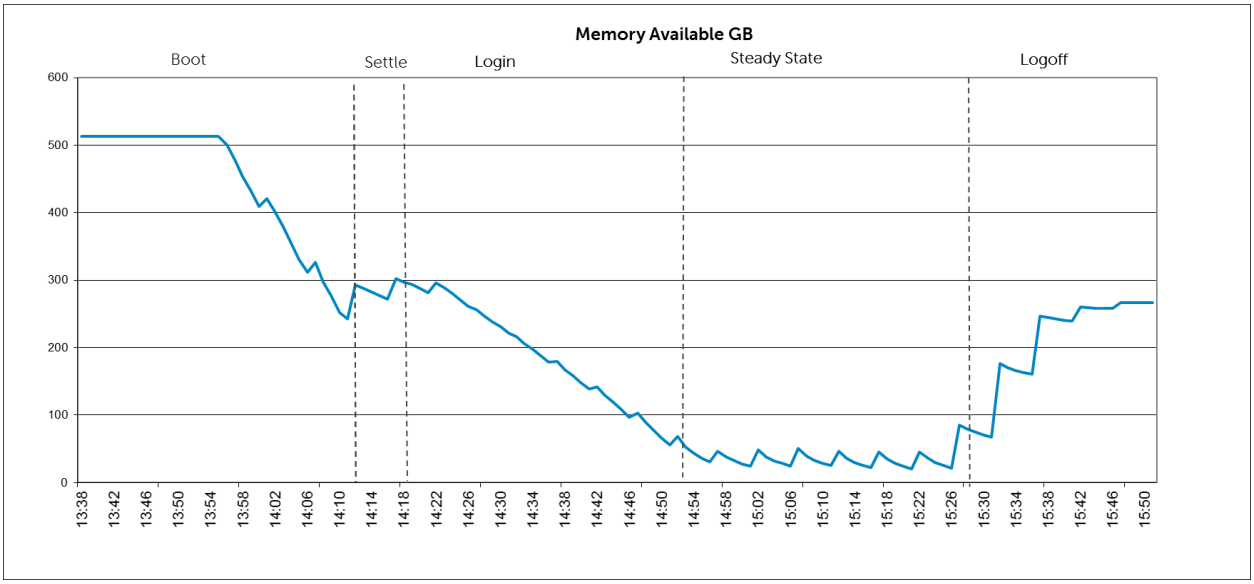

Network bandwidth is not an issue on this test run with a steady state peak of approximately 400 Mbps. Network usage is affected by the Login VSI content and profile and home folder redirection. The output queues on the network adapters remained at zero (0) throughout the testing.

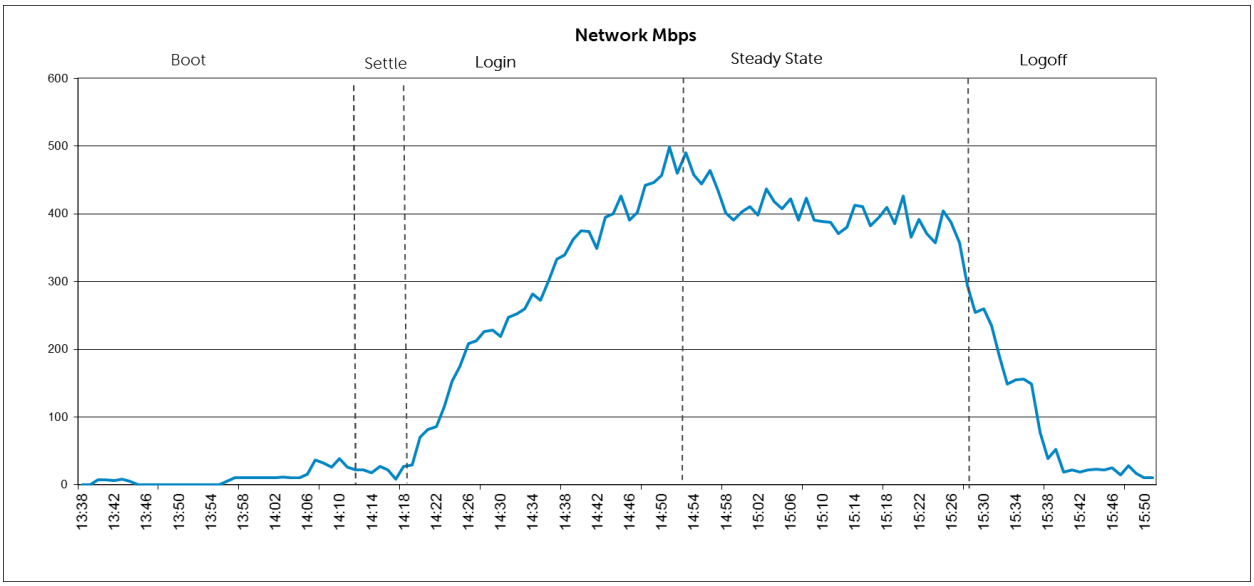

Citrix temporary memory caching for XenDesktop MCS desktops was not used because of the limited memory available to the compute host. All temporary data was stored on local disk of the compute host. It can also be seen that the latency was low throughout the logon steady state and logoff periods

The host reached a maximum of 10000 IOPs during the boot phase, 8000 during the login phase and dropped to 2000 during the steady state resulting in about 10 IOPS per user in the later state of the test. In all case the latency remained below 20 ms during the test.

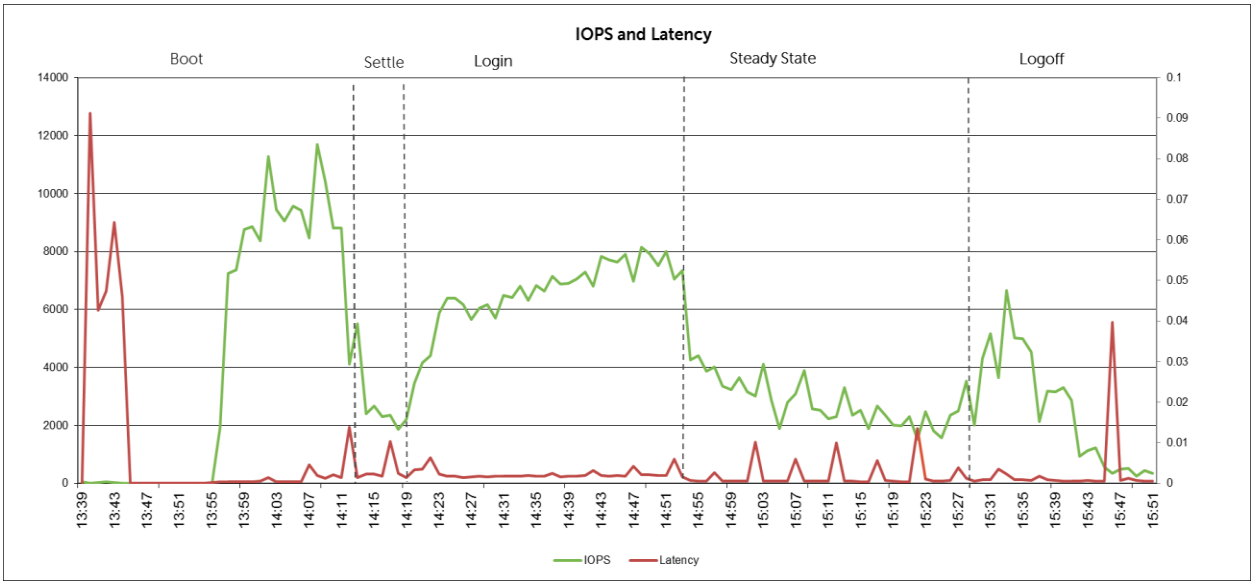

The Login VSI Max user experience score shown below for this test was not reached indicating there was little deterioration of user experience during testing and manually interacting with the test sessions confirmed this. Other indicators form the Login VSI Analyzer show that Logon times remained low throughout the test as did disk IO latency and application load times.

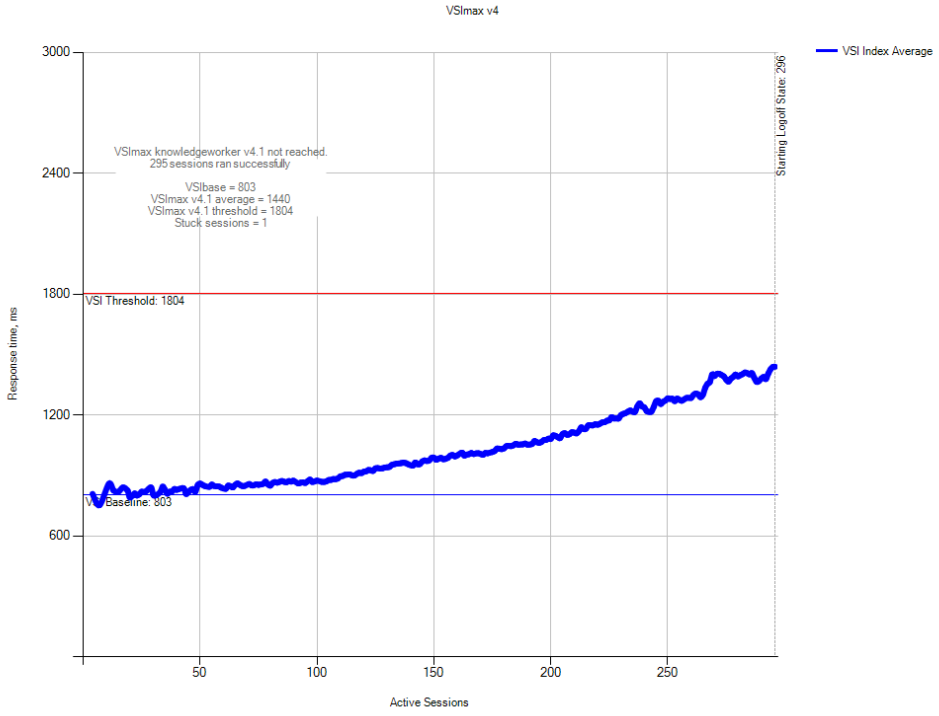

Notes:

- The results obtained on Server 2016 with Hyper-V role are in line with previous tests on Server 2012 R2 Hyper-V
- 512 GB of memory installed on each node is just about sufficient for the number of desktops. With both memory and CPU going close to maximum usage no extra desktops could have been accommodated for this configuration.

7.3.1.4 Knowledge Worker, 300 Users, Hyper-V 2016, XenApp 7.12 RDSH

The below graph shows the performance data for 300 Knowledge Worker sessions per host. The CPU reaches a steady state average of 60% during the test cycle when 300 users are logged on to the compute host.

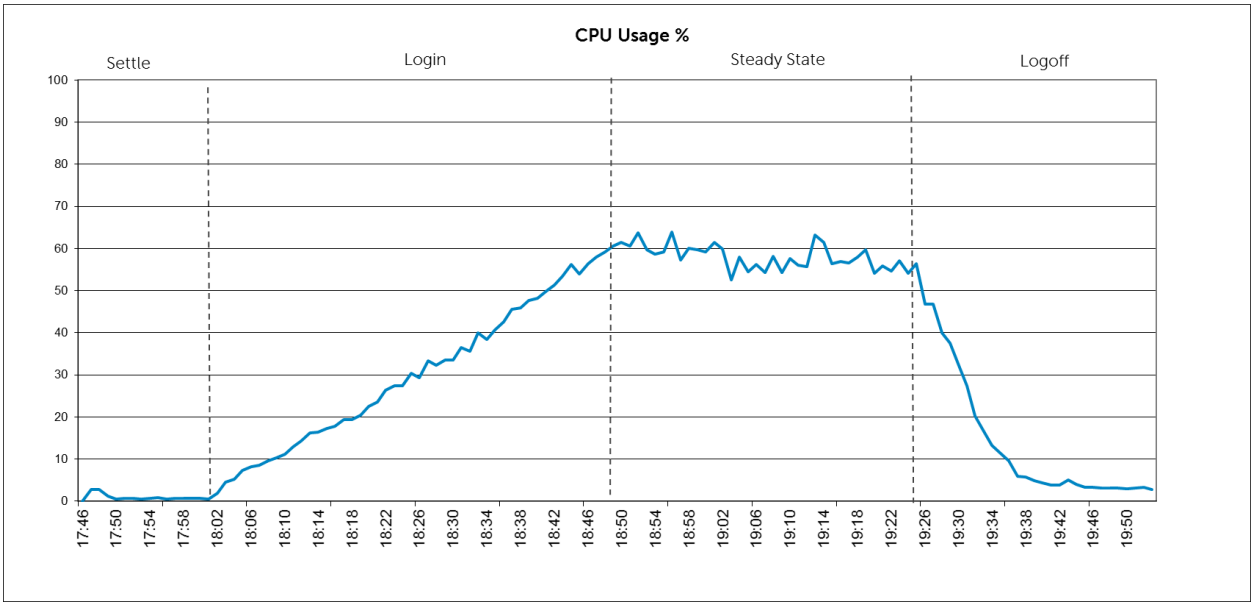

The memory available on the compute host was 512GB. A static assignment of 48GB was given to each Xenapp VM using 384GB for the VM memory. Each VM was assigned a temporary memory cache of 4GB. The available memory on the Hyper-V server was 84GB when all VMs were running and the Login VSI tests were in progress.

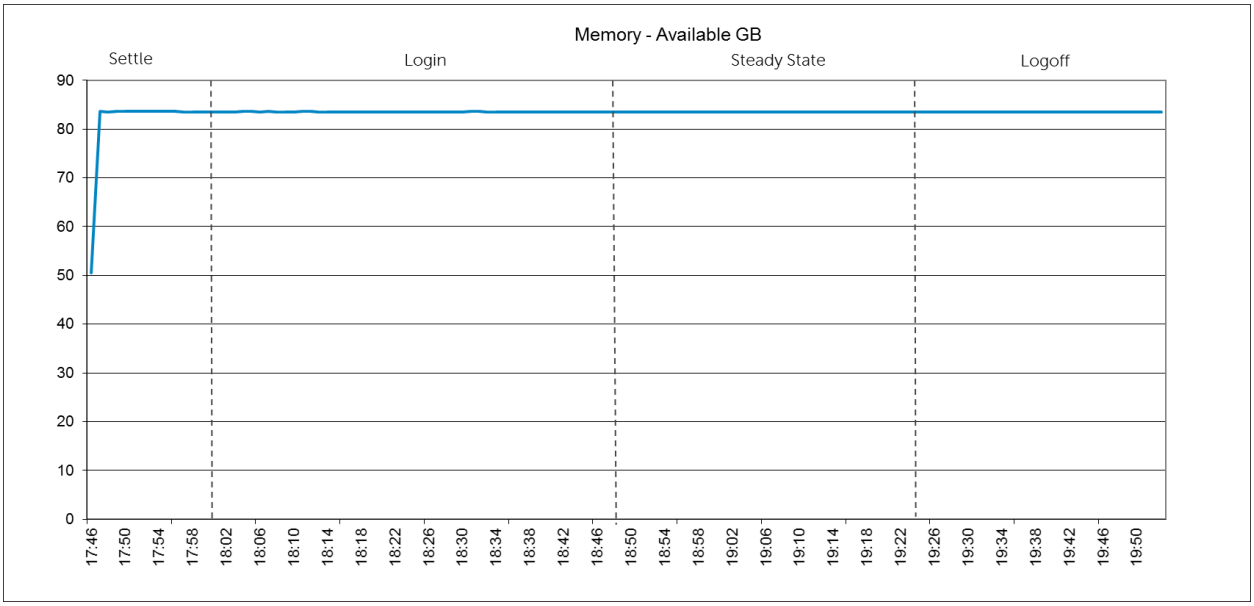

Network bandwidth is not an issue on this test run with a steady state peak of approximately 500 Mbps. Network usage is affected by the Login VSI content and profile and home folder redirection. The output queues on the network adapters remained at zero (0) throughout the testing.

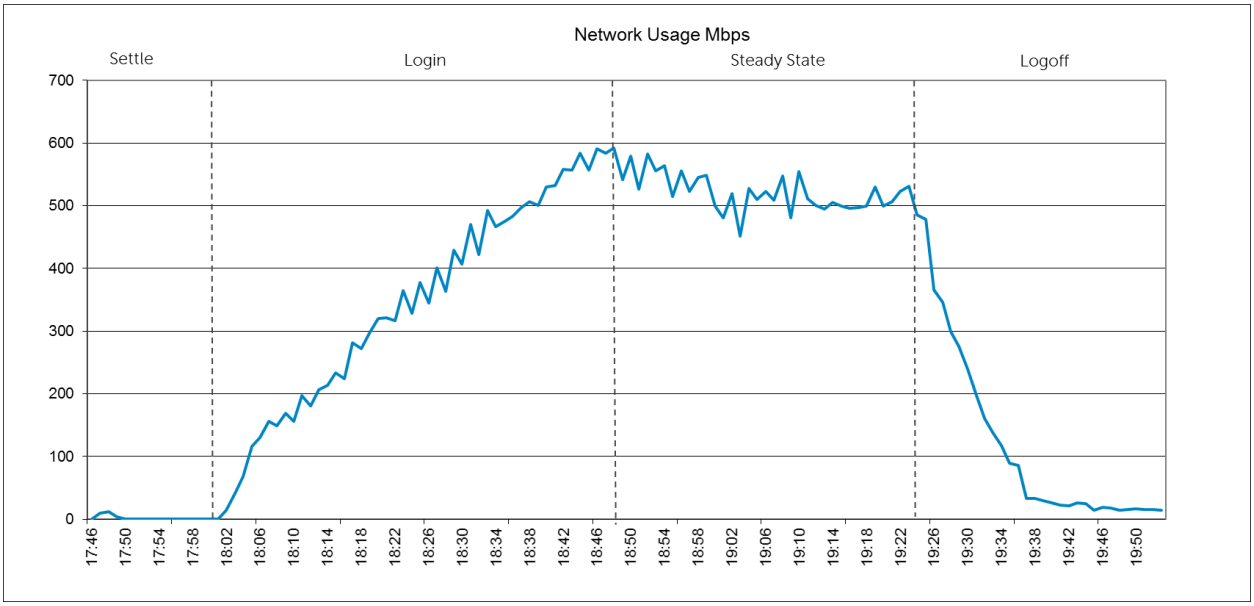

The IOPS are significantly lower than previous testing due to the effect of the temporary memory caching of Citrix Xenapp and XenDesktop. It can also be seen that the latency was at zero (0) throughout the logon steady state and logoff periods

The host reached a maximum of 600 IOPs during the login phase. This reduced to nearly 0 during the steady state showing the significant impact of the Xenapp MCS caching. Caching in RAM provides a small improvement over local SSD for temporary data, but also saves on write wear on the SSDs.

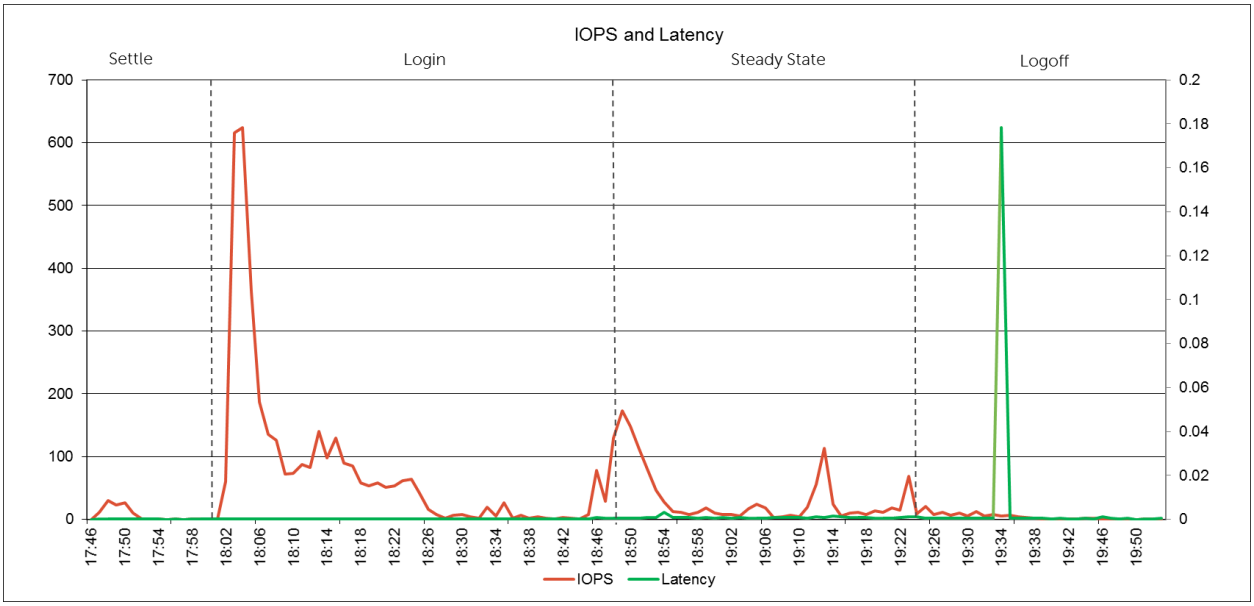

The Login VSI Max user experience score shown below for this test was not reached indicating there was little deterioration of user experience during testing and manually interacting with the test sessions confirmed this. Other indicators form the Login VSI Analyzer show that Logon times remained low throughout the test as did disk IO latency and application load times.

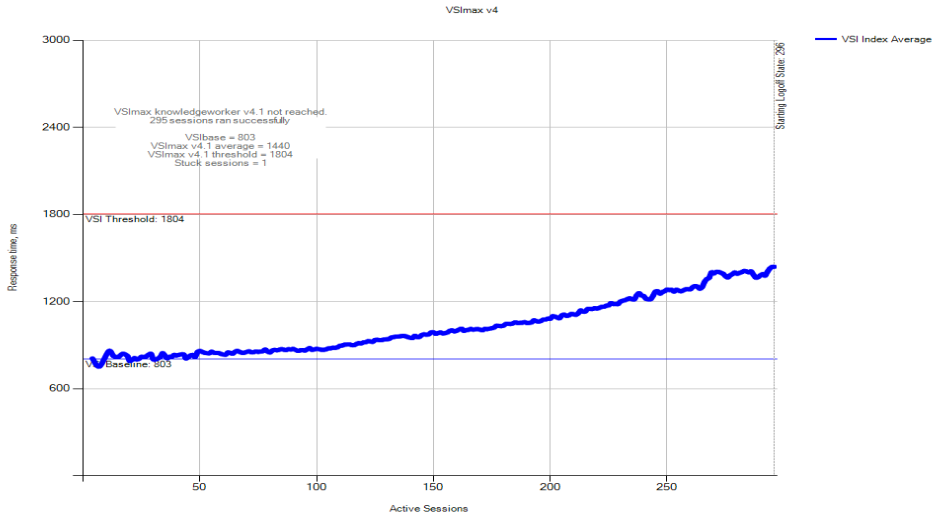

Notes:

- Although all the performance indicators are very good and VSIMAX is not reached, efforts to increase the density on the compute host resulted in stuck Login VSI sessions. This is currently being investigated and it is expected that further improvements in density can be achieved
- 512 GB of memory installed on each node allows significant memory to be assigned to each VM
- NUMA operation was enabled in the server BIOS and Cluster on Die snooping was in effect as recommended for highly NUMA aware workloads. CPU oversubscription of 2.0x was employed as recommended in the Citrix article "New" Citrix Best Practices 2.0.

7.3.1.5 Power Worker, 150 Users, ESXi 6.0 U2, XenDesktop 7.9 MCS

The CPU usage has long sustained period of 100% utilization, however, the core utilization is no more than 90%. From the LoginVSI analysis below, we can see that the resource utilization did not reach saturation point. This is another indication that core utilization is a better indicator for CPU resource utilization.

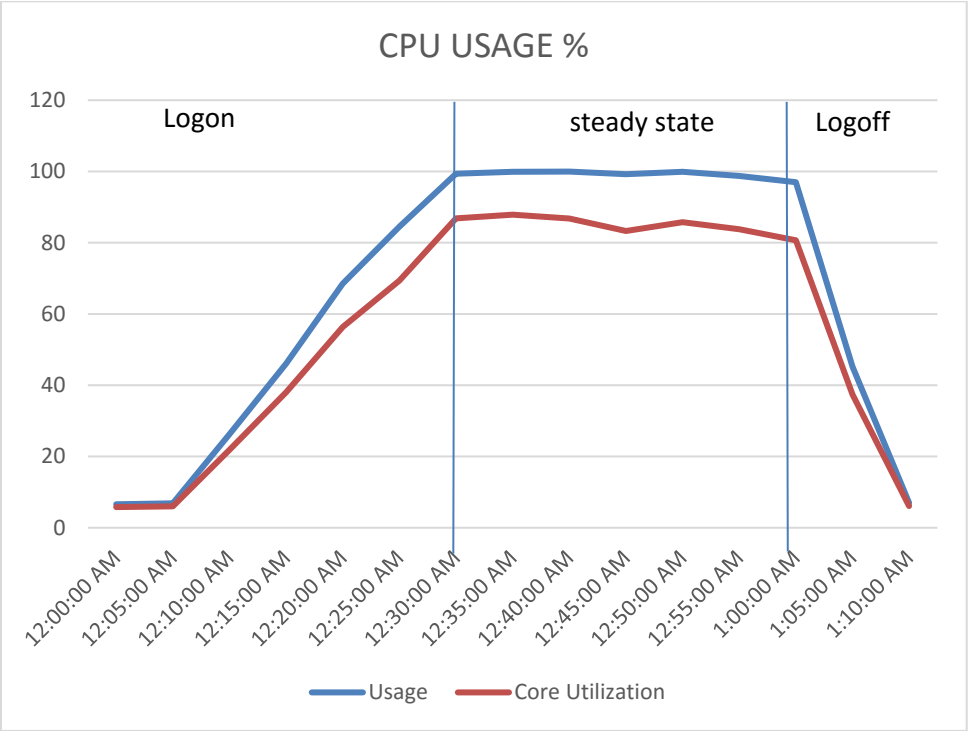

In regards to memory consumption for the host, out of a total of 512GB available memory, the VMs consumed almost all. However active memory usage is only slightly over half, peaked at 284GB.

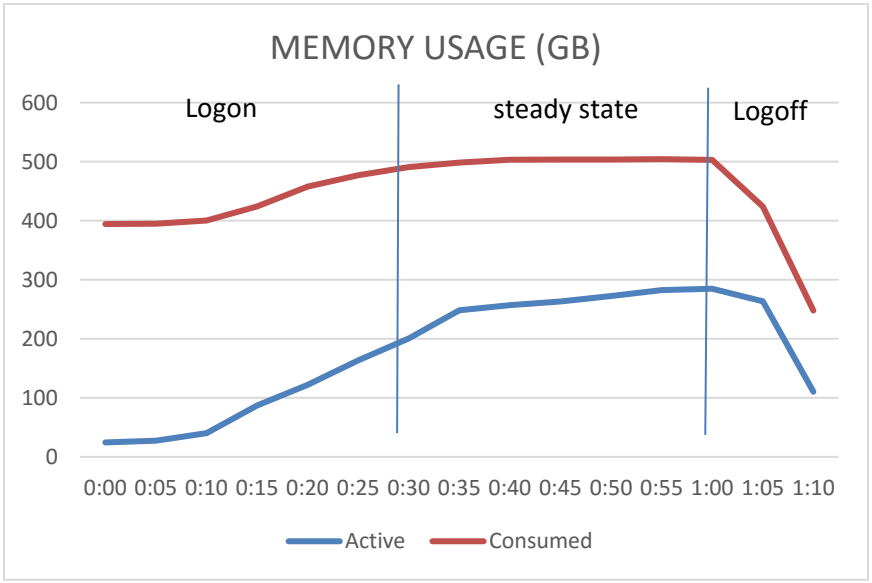

Network usage of the Power work load showing heavier usage than other two workloads. Peaked at around 4.5 Mbps per user, and under 700Mbps on the host. It is still well under the 1Gbps NIC capacity.

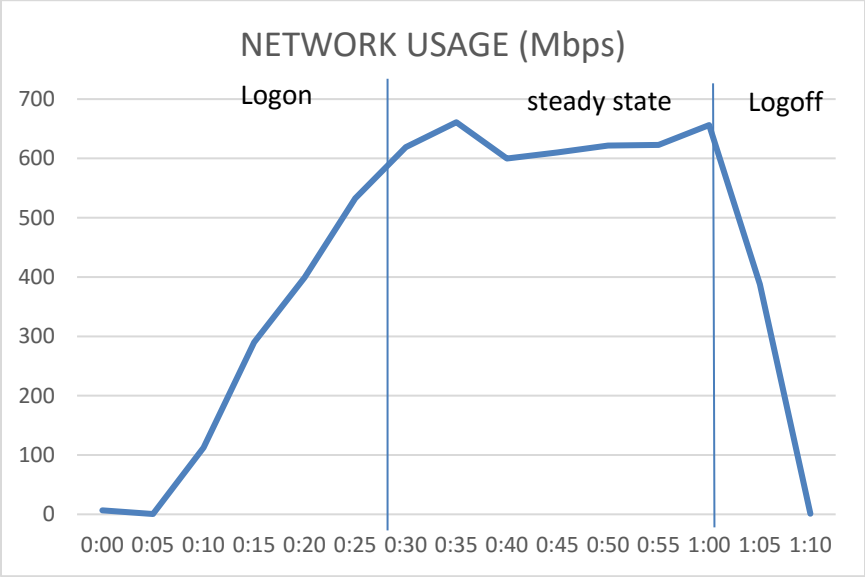

The datastore IOPS reached maximum 1670 of read IO and 380 write IO, gives an average 13.67 IOPS per user. All the time throughout the testing period, the data store latency remains under 2ms. Well below the 20ms latency threshold.

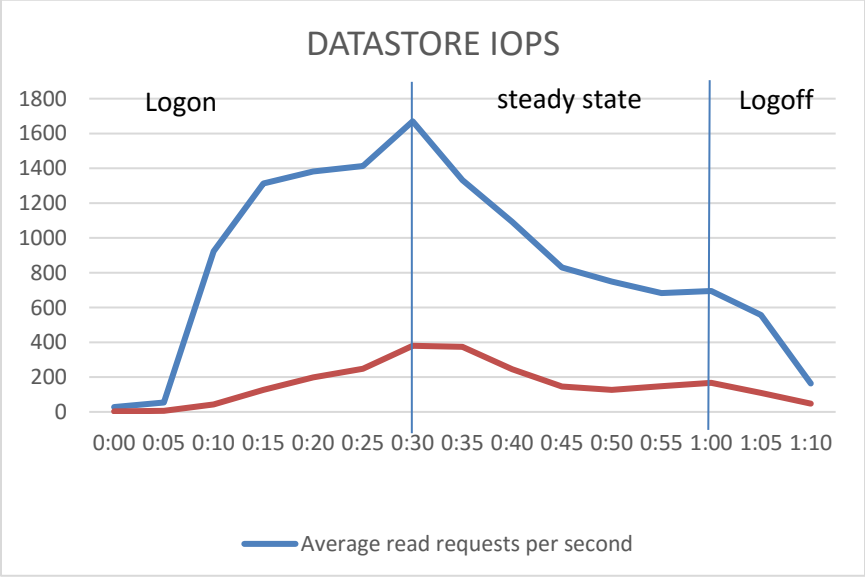

The Login VSI Max user experience score for this test was not reached indicating there was little deterioration of user experience during testing and manually interacting with the test sessions backed this up, mouse and window responses were fast and video play back was of good quality.

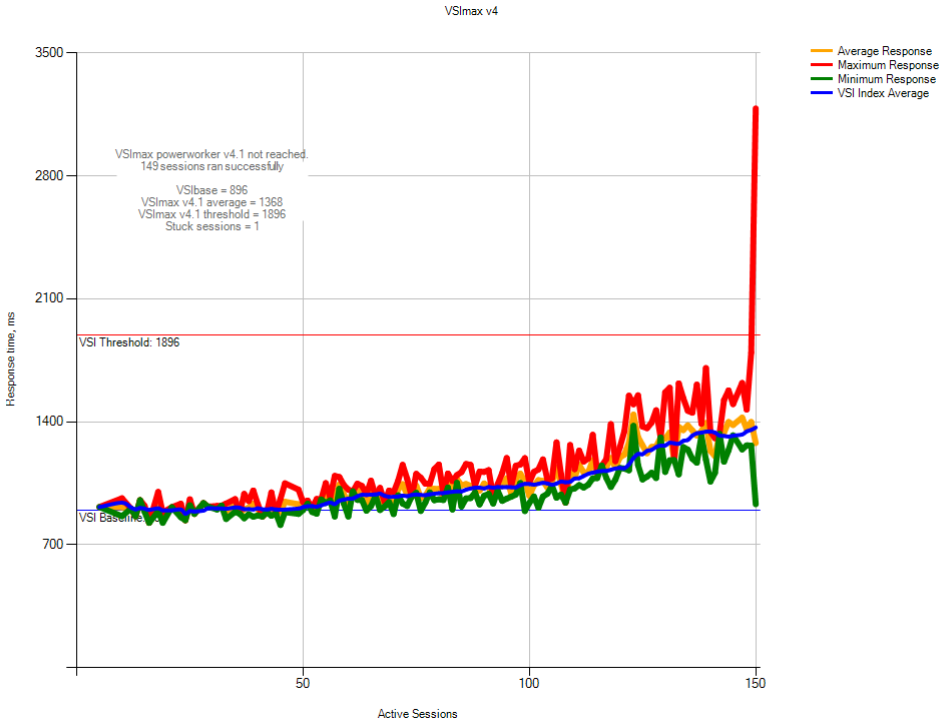

Notes:

There was one session failed to complete test cycle. Upon investigation, it's concluded the failure was not related to host performance reason. And since this is the only session out of 150 sessions, it is below the 2% acceptable results threshold. The LoginVSI test "VSImax not reached" result is considered valid.

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:

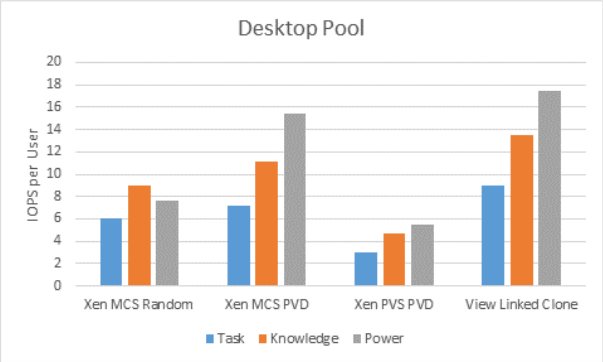

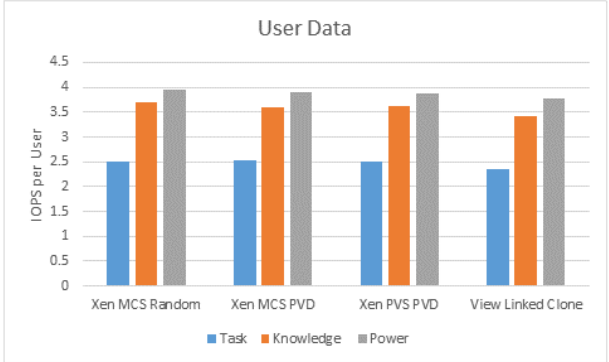

If your workload requirements vary significantly either in IOPS or capacity, we recommend placing T2 data on a discrete array. Our synthetic workload resulted in an overall data reduction rate of 3.7:1, as observed below:

UNITY VSA WITH XTREMIO FOR COMBINED T1/T2

DATA REDUCTION RATIO

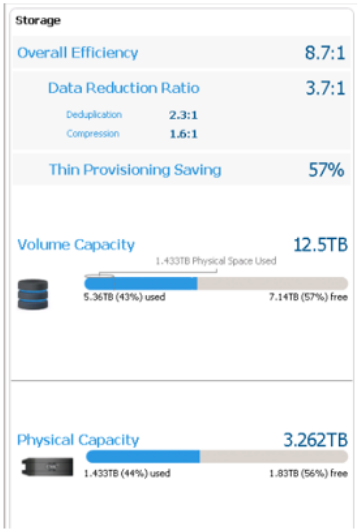

- Observed data reduction ratio across Login VSI workloads was 3.7:1
- 3.262TB physical capacity can support 3.262 * 3.7 = 12.0694TB data
- About 24GB capacity consumed (including desktop and user data) for each user in this 500-user test

Acknowledgements

Thanks to David Hulama of the Dell Wyse Technical Marketing team for his support and assistance with datacenter 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.

About the Authors

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. Follow Peter @ExitTheFastLane or www.ExitTheFastLane.com.

Jerry Van Blaricom is a Lead Architect in the Cloud Client Solutions Engineering Group at Dell. Jerry has extensive experience with the design and implementation of a broad range of enterprise systems and is focused on making Dell's virtualization offerings consistently best in class.