Dell Wyse Datacenter for Microsoft Remote Desktop Services Reference Architecture

A Reference Architecture for the design, configuration and implementation of a Microsoft RDS environment.

Dell Engineering
May 2017

A Dell Reference Architecture

Revisions

Date	Description
May 2017	Initial release

Table of 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 Microsoft Remote Desktop Services (RDS) on Windows Server 2016 Hyper-V hypervisor. 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. Content topics additional to what is listed in this template should be included in a heading at level two or lower.

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

This is the initial release of the Dell Wyse Datacenter reference architecture for Microsoft RDS with Windows Server 2016 and Hyper-V hypervisor.

2 Soluiton 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 for Microsoft RDS include:

- Desktop-based virtualization using Remote Desktop Virtualization Hosts (RDVH)
  - Pooled Virtual Desktops (Non-persistent)
  - Personal Virtual Desktops (Persistent)
- Session-based virtualization using Remote Desktop Sessions Hosts (RDSH)
  - RemoteApp programs
  - Personal or Shared session desktops

Please refer to the Remote Desktop Services section for an explanation of acronyms and components.

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.

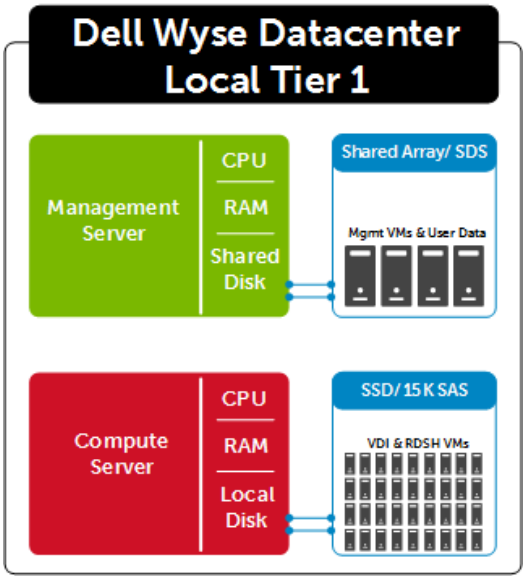

In the Shared Tier 1 solution model, an additional high-performance shared storage array is added to handle the execution of the VDI sessions. All compute and management layer hosts in this model are diskless where possible or utilize two mirrored disks if required for the operating system.

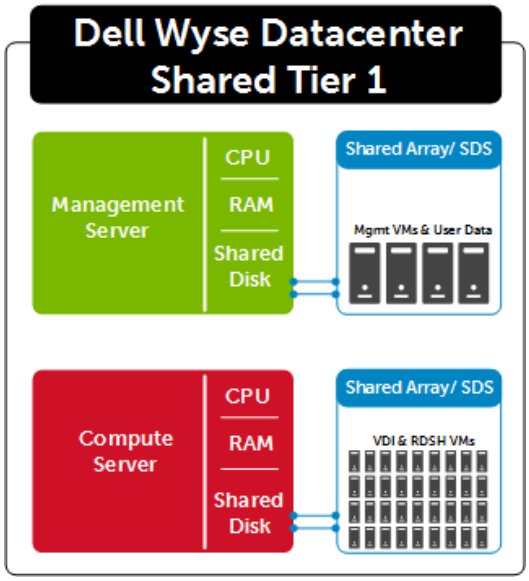

2.3 Solution Layers

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.

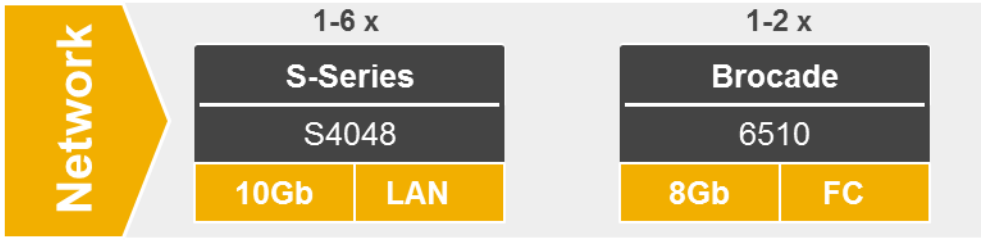

2.3.2 Compute

The compute layer consists of the server resources responsible for hosting the desktop or RDSH VMs on Hyper-V hypervisor with local or shared Tier 1 solution models (local Tier 1, all-flash, pictured below).

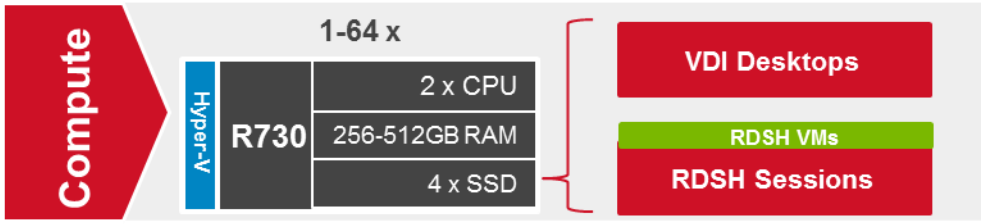

2.3.3 Management

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. In the image below, RDCB stands for Remote Desktop Connection Broker, RDWA stands for Remote Desktop Web Access, and RDG stands for Remote Desktop Gateway. Refer to the Remote Desktop Services section for information on these components.

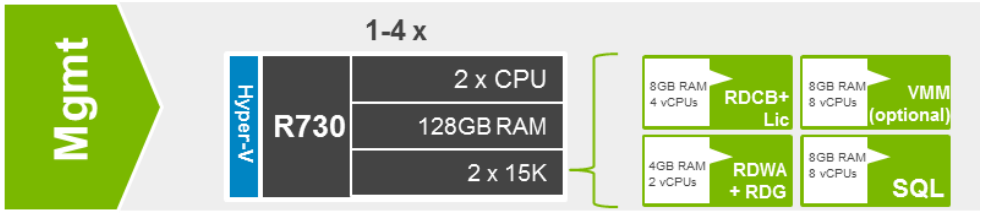

2.3.4 Storage

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 along with Dell EMC Unity arrays for discrete Tier 2 management VM storage and user data file storage. The discrete T1 and T2 array configurations support 3,000 knowledge worker users with 5GB of user data (disk configurations can be adjusted to support large amounts of user data).

Combined Tier 1 and Tier 2 (including user data) for up to 500 users:

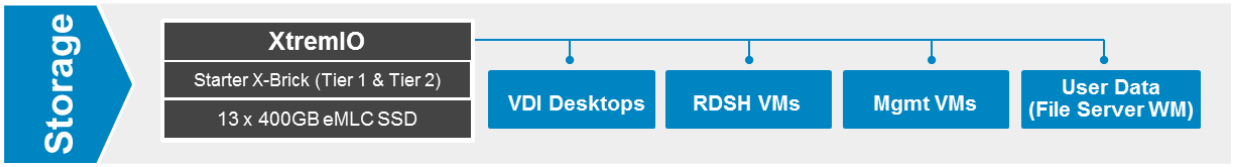

Distinct arrays for Tier 1 and Tier 2 for up to 3,000 knowledge users:

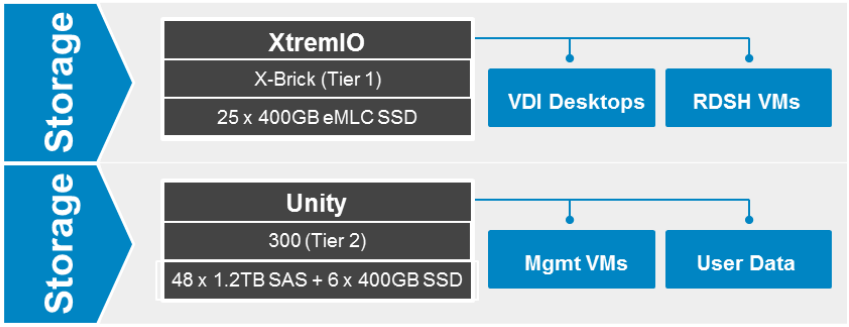

2.4 Local Tier 1

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. Disk size depends on total capacity requirements of all VMs but a minimum of 4 x 800GB SSDs is recommended. SQL server is optional for a pilot/POC deployment, but required if intending to scale the setup for production use and ensuring the broker is highly available. SQL is also required if installing the optional System Center Virtual Machine Manager (SCVMM) for overall management of the environment.

Note: 150 user density based on Task Worker workload.

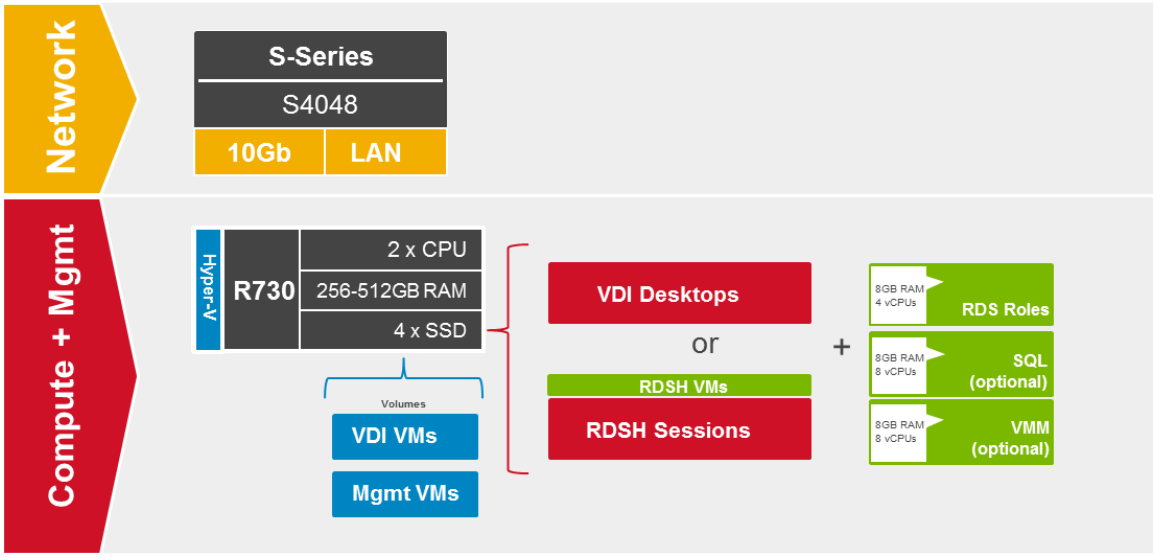

2.4.2 Local Tier 1 for Rack Servers

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. All-flash pictured below, if spinning disk is desired substitute the SSDs with 12 x 600GB 15K SAS HDDs. A Unity 300 array is added to provide shared storage (T2) for the management VMs as well as for user data.

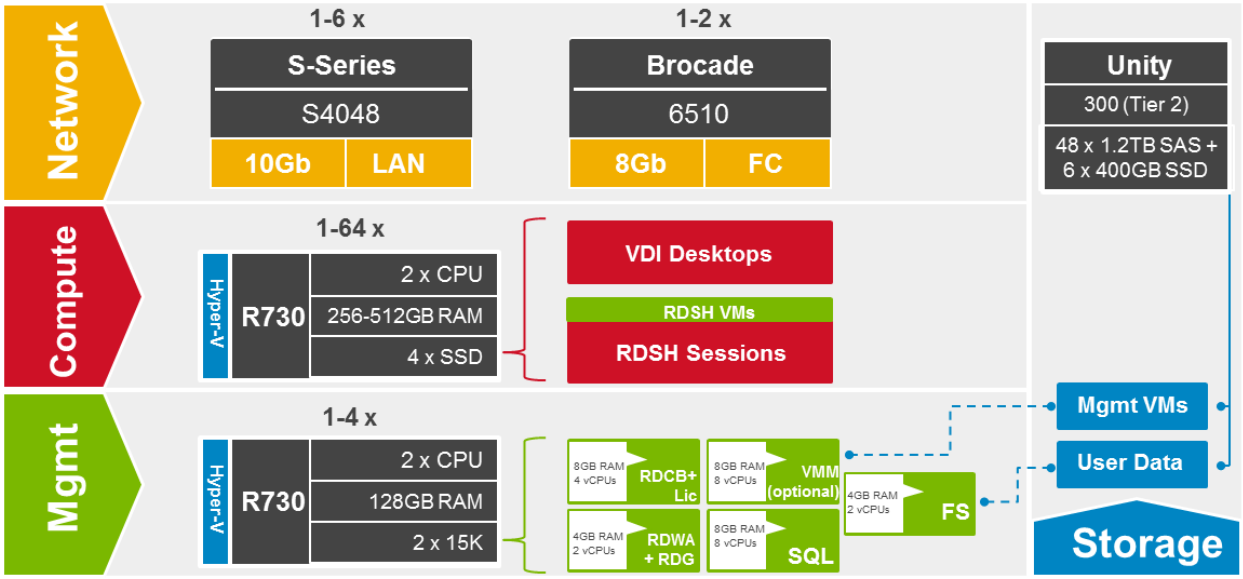

Note: Compute disk size depends on total capacity requirements. Recommended minimum of 4 x 800GB SSDs.

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.

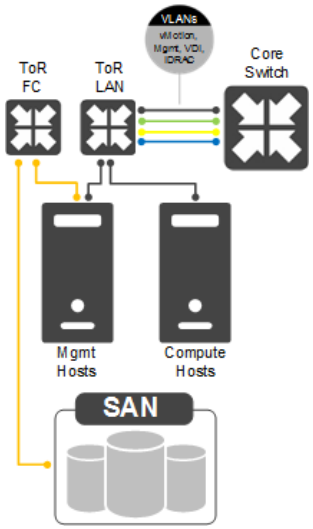

2.4.2.2 Local Tier 1 – Rack Cabling (HA)

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

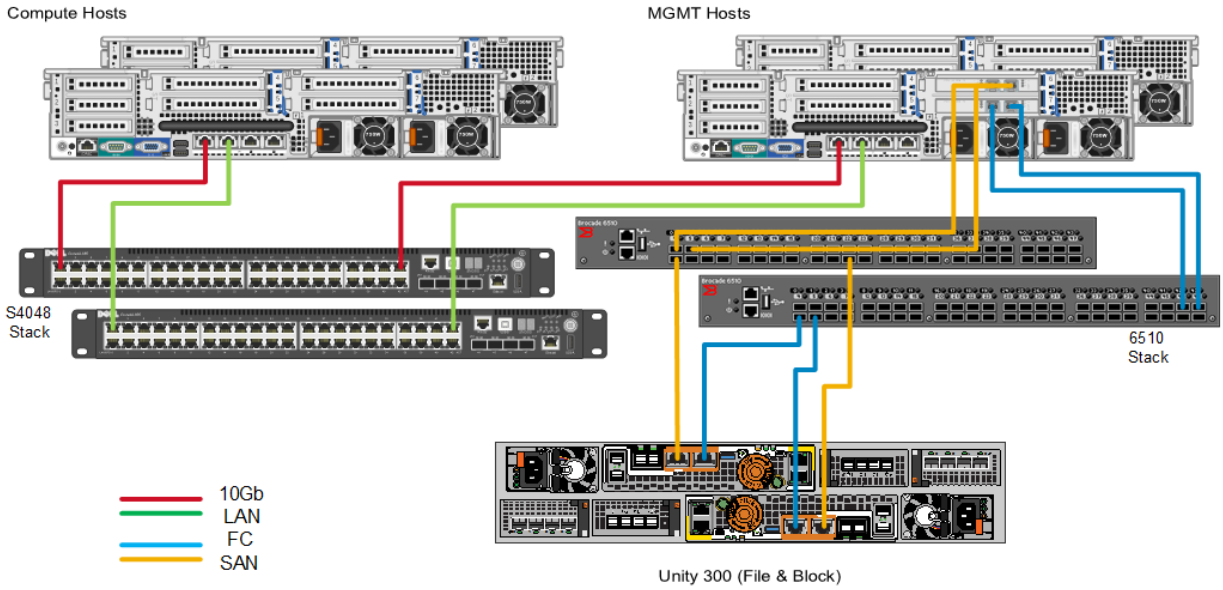

Refer to the Unity 300 Tier 2 section for a diagram of disk enclosure connections via mini-SAS HD.

2.4.2.3 Local Tier 1 – Rack Storage Scaling Guidance

Local Tier 1 Storage 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: Density is based on empirical data and has not been validated with Windows Server 2016 with Hyper-V.

2.4.3 Local Tier 1 for Blade Servers

The Local Tier 1 solution model for blade servers provides a high-performance 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. 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.

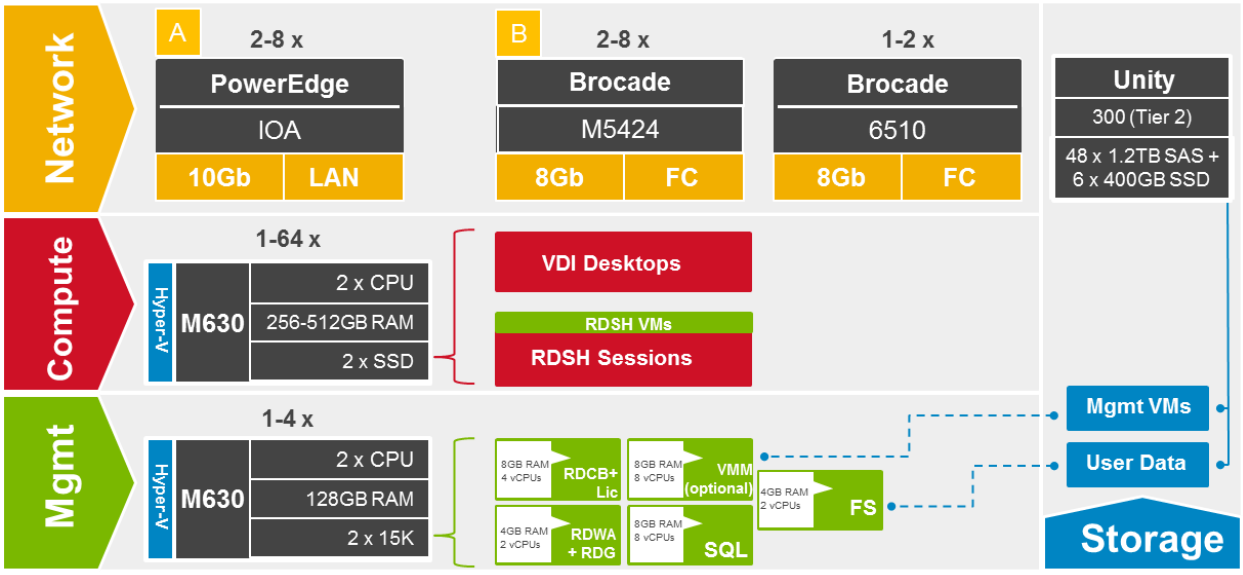

Note: This configuration only possible by installing OS and virtual machines on same set of disks and RAID controller which may impact performance. Compute disk size depends on total capacity requirements.

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.

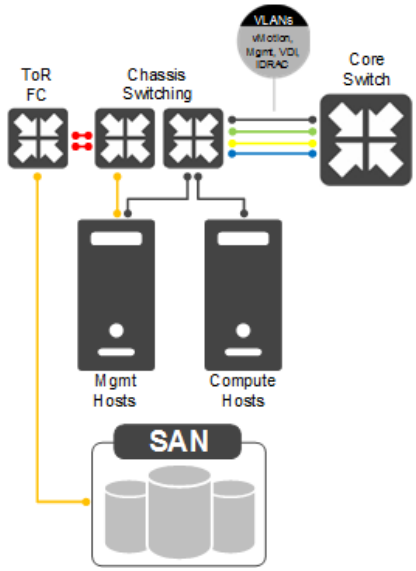

2.4.3.2 Local Tier 1 – Blade Chassis Cabling (HA)

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

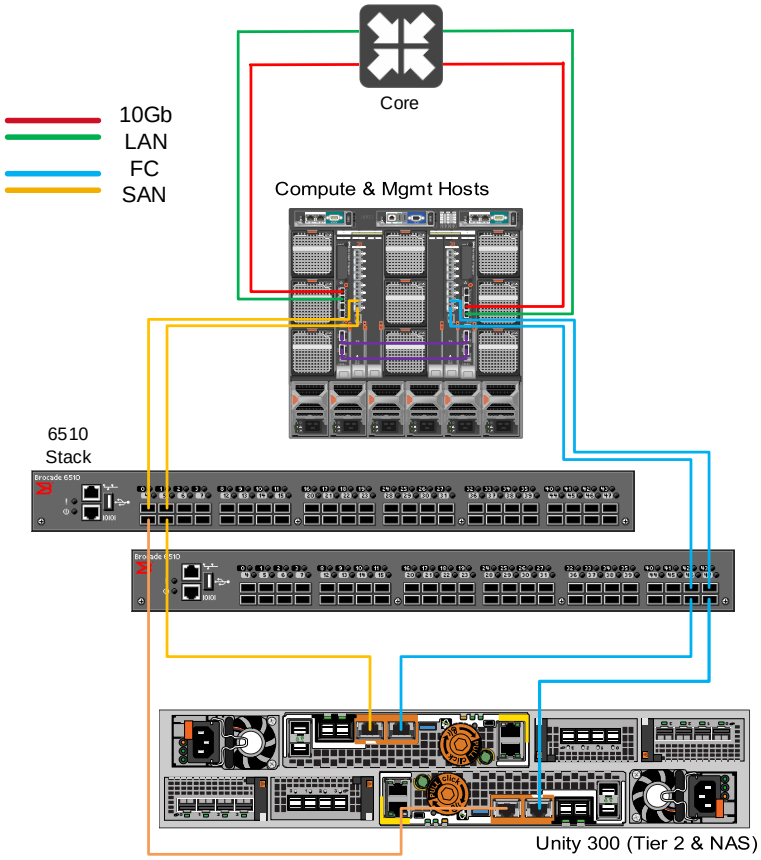

Refer to the Unity 300 Tier 2 section for a diagram of disk enclosure connections via mini-SAS HD.

2.4.3.3 Local Tier 1 – Blade Storage Scaling Guidance

Local Tier 1 Storage 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: Density is based on empirical data and has not been validated with Windows Server 2016 with Hyper-V.

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, storage for user data is provided via file server VMs with disks residing on the X-Brick array.

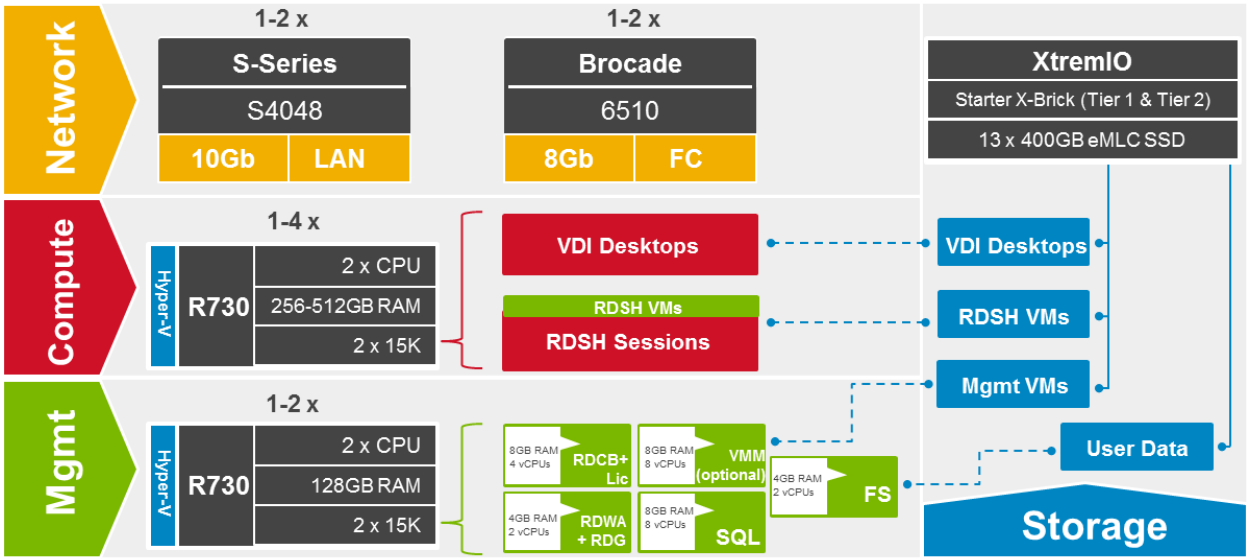

More Than 500 Users

For over 500 users, the storage layers are separated into discrete arrays, as depicted in the figures below. Additional arrays are added for Tier 1 and Tier 2 as the user count grows. The first diagram depicts disk configurations for each array that can support up to 1,500 knowledge users while the second diagram shows the configuration for up to 3,000 users. Additional arrays and/or larger disk sizes can also be added to Tier 2 if necessary due to larger user data capacity needs.

Up to 1,500 Users

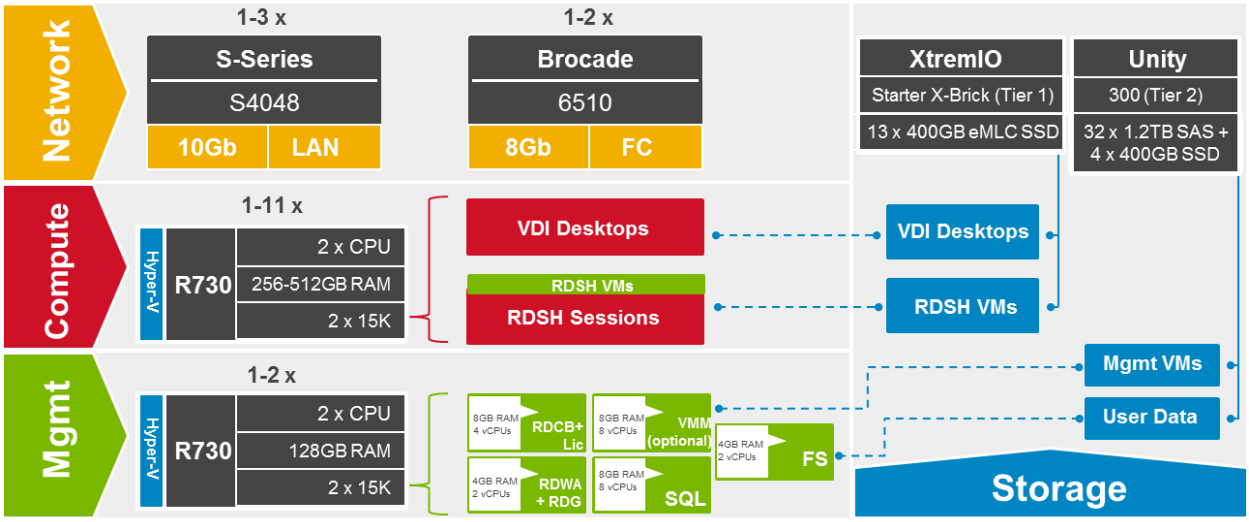

Up to 3,000 Users

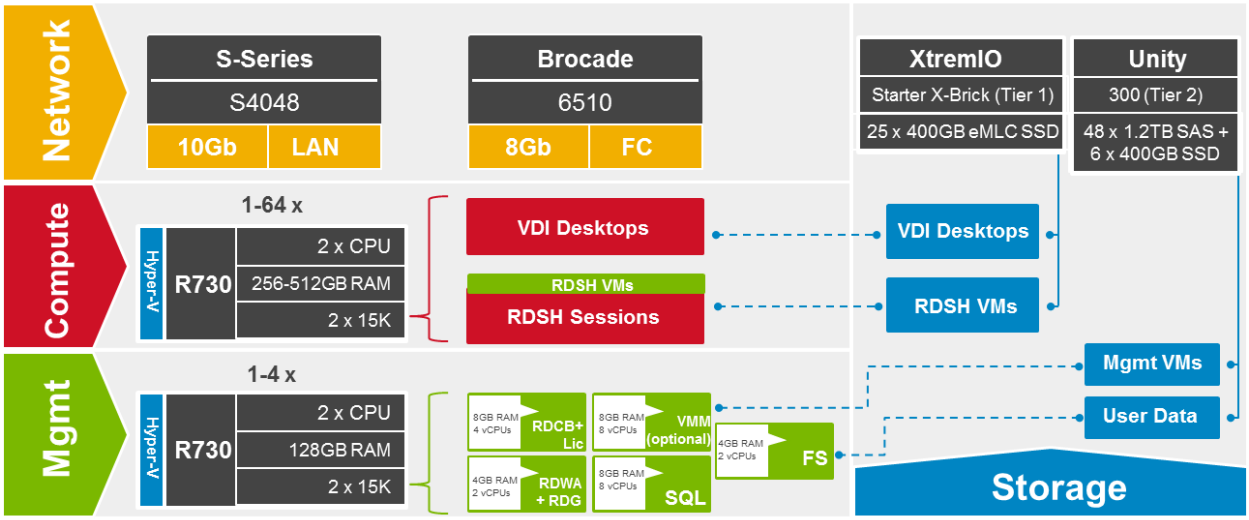

Note: Maximum of 64 nodes supported per cluster, but each cluster limited to a maximum of 8000 running virtual machines.

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.

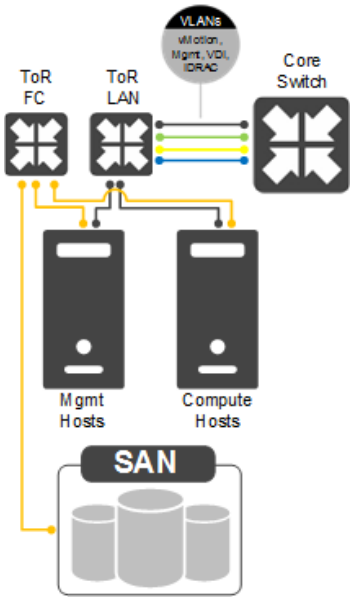

2.5.1.2 Shared Tier 1 – Rack Cabling (HA)

500 Users or Less - Combined T1 and T2 Option

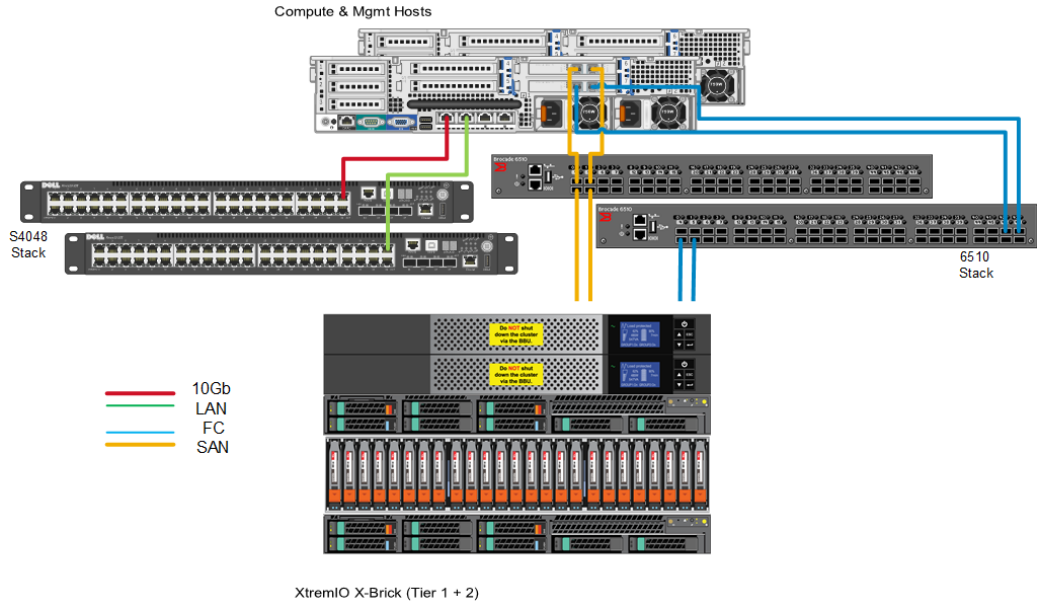

Discrete T1 and T2 Storage Arrays

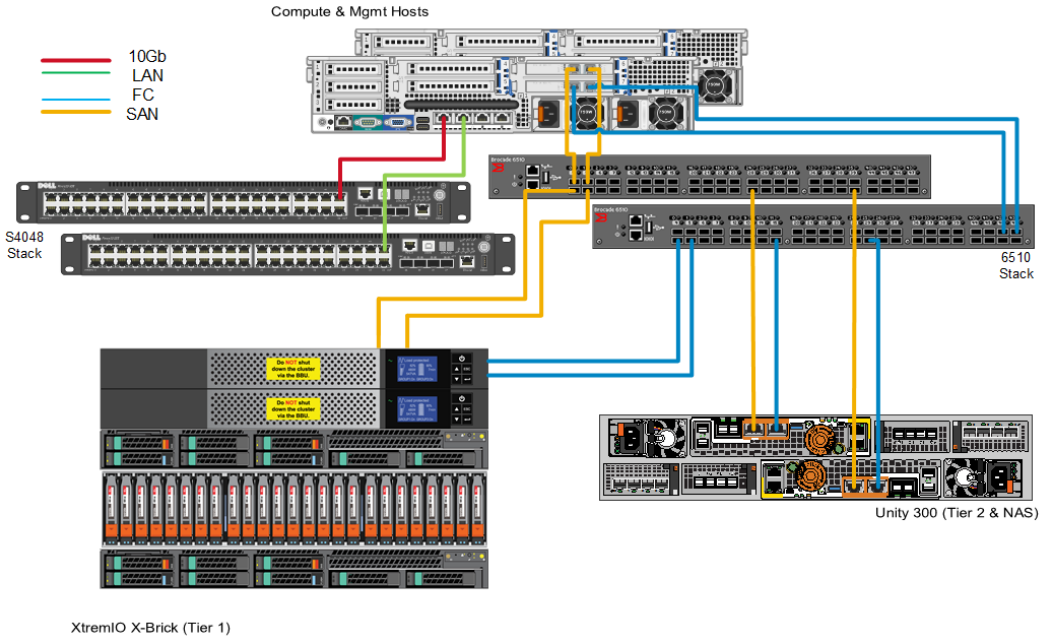

Refer to the Unity 300 Tier 2 section for a diagram of disk enclosure connections via mini-SAS HD.

2.5.1.3 Shared Tier 1 – Rack Storage Scaling Guidance

Shared Tier 1 Storage Scaling (Rack - FC)					
User Scale	**XtremIO T1**	**XtremIO T2**	**Unity T2 & NAS**	**ToR LAN**	**ToR 8Gb FC**
0 - 500	Starter X-Brick SSD	File Server	-	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: Density is based on empirical data and has not been validated with Windows Server 2016 with Hyper-V.

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. 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

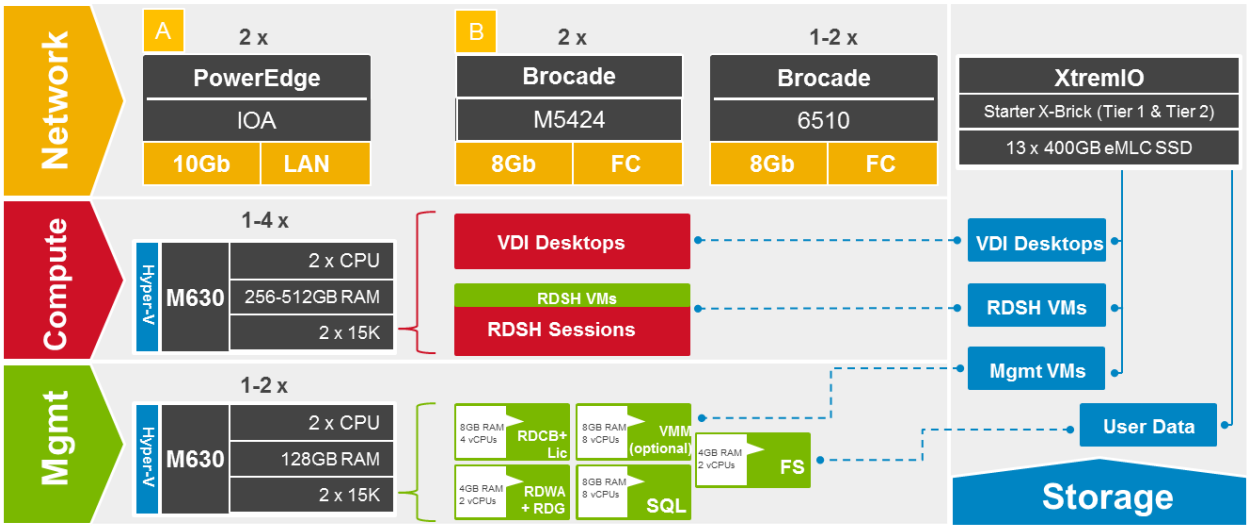

More Than 500 Users

For over 500 users, the storage layers are separated into discrete arrays, as depicted in the figures below. Additional arrays are added for Tier 1 and Tier 2 as the user count grows. The first diagram depicts disk configurations for each array that can support up to 1,500 knowledge users while the second diagram shows the configuration for up to 3,000 users. Additional arrays and/or larger disk sizes can also be added to Tier 2 if necessary due to larger user data capacity needs.

Up to 1,500 Users

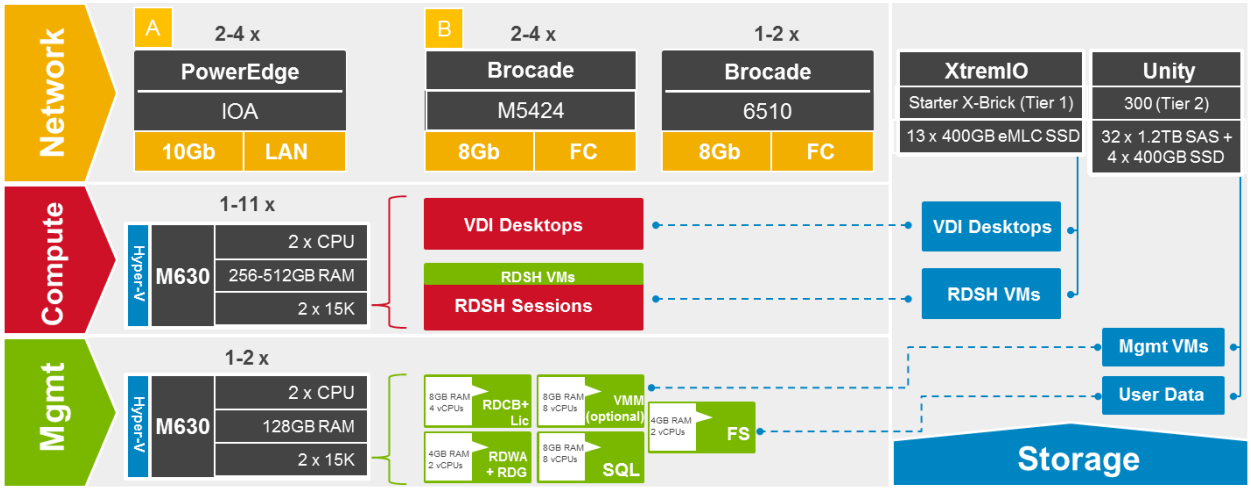

Up to 3,000 Users

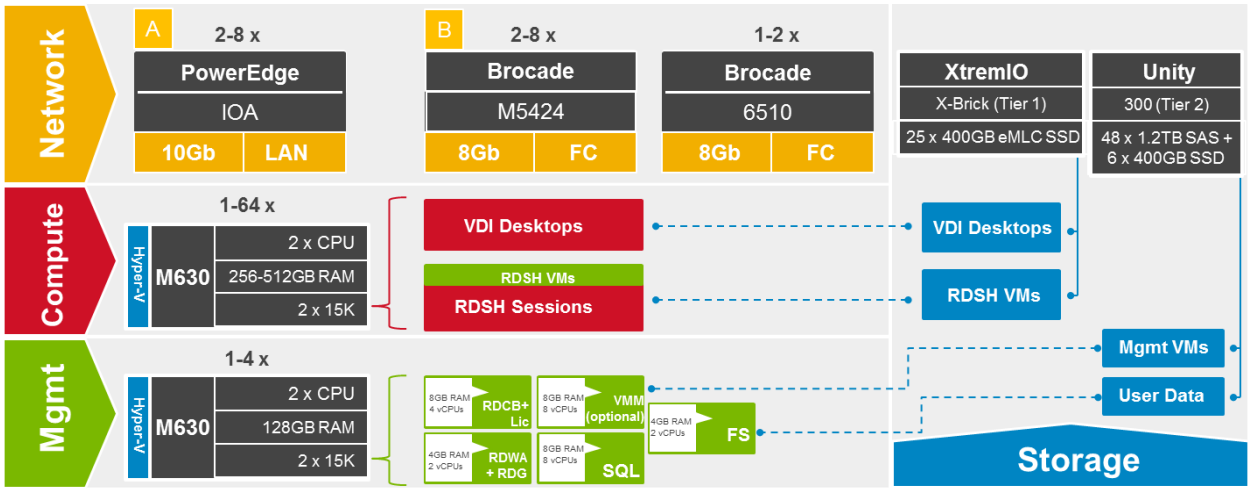

Note: Maximum of 64 nodes supported per cluster, but each cluster limited to a maximum of 8000 running virtual machines.

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.

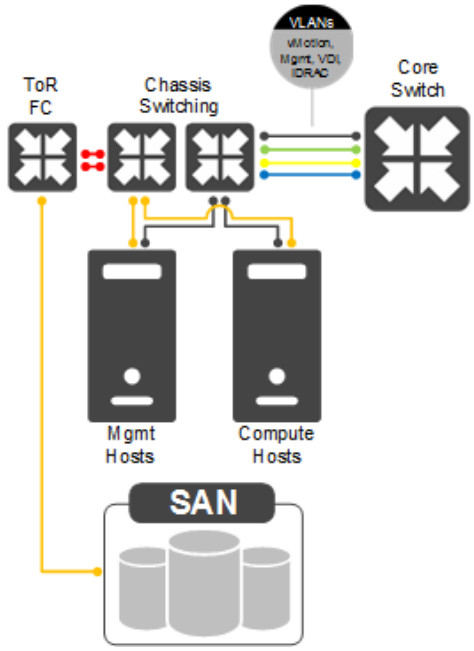

2.6.1.2 Shared Tier 1 – Rack Cabling (HA)

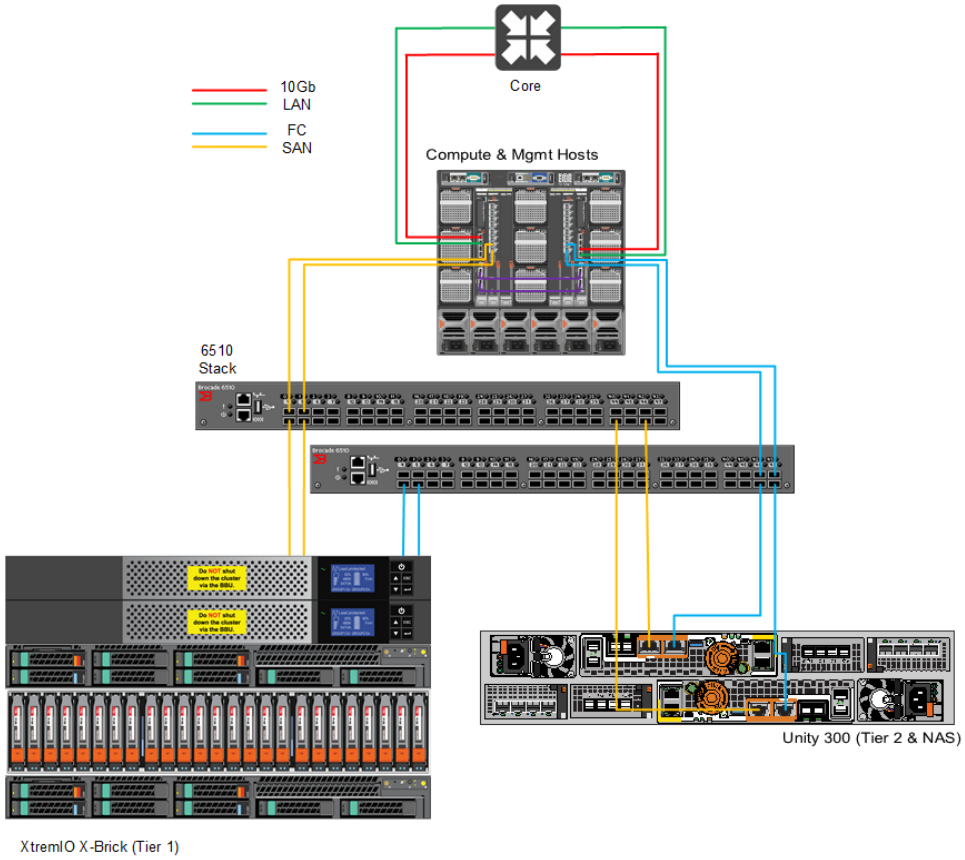

Refer to the Unity 300 Tier 2 section for a diagram of disk enclosure connections via mini-SAS HD.

2.6.1.3 Shared Tier 1 – Blade Storage Scaling Guidance

Shared Tier 1 HW scaling (Blade - FC)						
User Scale	**XtremIO T1**	**XtremIO T2**	**Unity T2 & NAS**	**Blade LAN (A Fabric)**	**Blade FC (B Fabric)**	**ToR 8Gb FC**
0-500	Starter X-Brick SSD	File Server	-	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: Density is based on empirical data and has not been validated with Windows Server 2016 with Hyper-V.

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+	Redundant hot-swap PSUs & fans	1Gb connectivity
	Non-blocking, line-rate performance	VRF-lite, Routed VLT, VLT Proxy Gateway	
	260Gbps full-duplex bandwidth	User port stacking (up to 6 switches)	
	131 Mbps forwarding rate	Open Networking Install Environment (ONIE)	

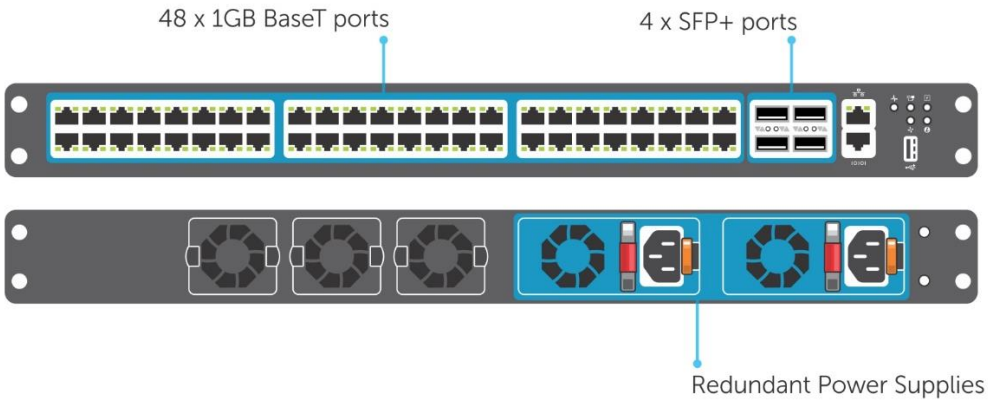

3.1.2 Dell Networking S-Series S4048 (10Gb ToR Switch)

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)	

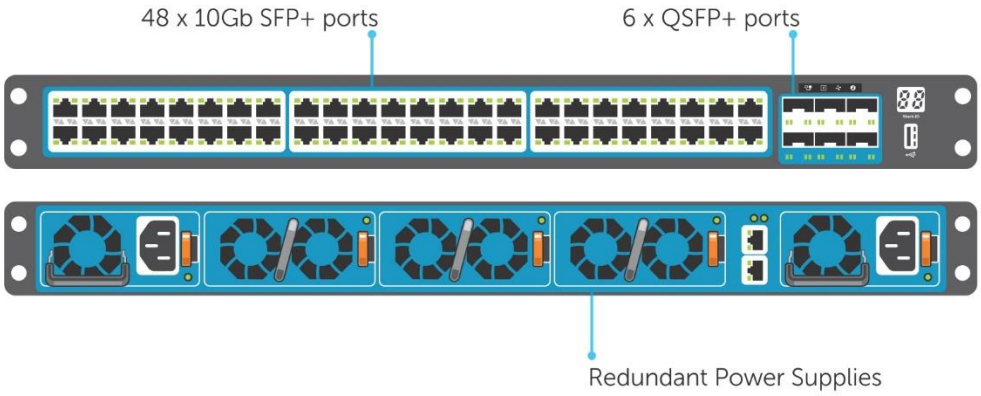

For more information on the S3048, S4048 switches and Dell Networking, please visit: LINK

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

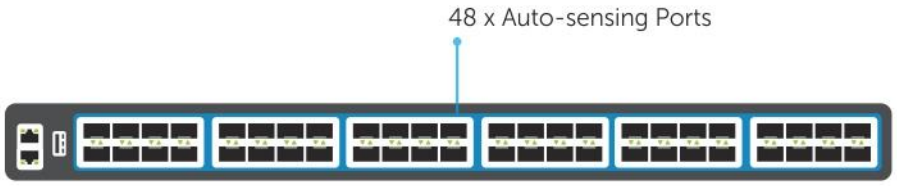

For more information on the Brocade 6510 switch, please visit: LINK

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	

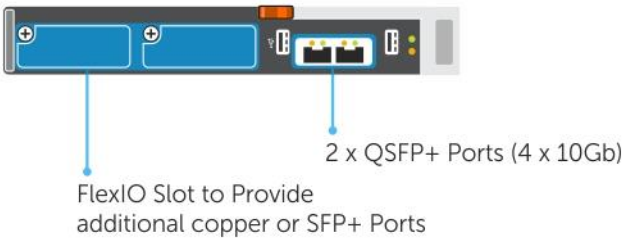

For more information on the PowerEdge IOA switch, please visit: LINK

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, and supports up to 16 x 2.5" SAS disks. The Dell PowerEdge R730 offers uncompromising performance and scalability in a 2U form factor.

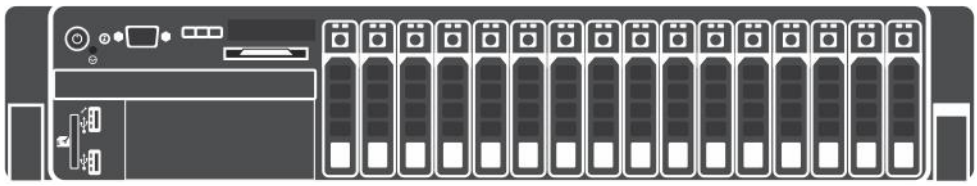

For more information on the R730, please visit: Link

NOTE: Configurations shown below are recommendations for VDI based on stated density determined by our testing. They do not represent absolute platform maximums and can be adjusted as needed.

3.2.1.1 Local Tier 1 Rack

In the Local Tier 1 model, desktop or session-based VMs execute from local storage on each Compute server. Microsoft Hyper-V is the hypervisor used in this solution. 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 RDSH while optionally reducing the amount of RAM.

R730 Local T1	Compute	Management
CPU	2 x E5-2698v4 (20C, 2.2GHz)	2 x E5-2660v4 (14C, 2.0GHz)
Memory	16 x 32GB 2400MT/s RDIMMs **Effective speed: 2400MT/s @ 512GB**	8 x 16GB 2400MT/s RDIMMs **Effective speed: 2400MT/s @ 128GB**
Storage Ctrls	PERC H730P – RAID10	PERC H330 – RAID1
Storage	4 x SSD or 12 x 15K SAS **(Hypervisor + VMs)**	2 x 300GB 15K SAS **(Hypervisor)**
Network	4 x 10Gb SFP+ (BT options available)	4 x 10Gb SFP+ (BT options available) 2 x QLogic 2562 8Gb DP FC HBA
iDRAC	iDRAC8 Enterprise	iDRAC8 Enterprise
Power	2 x 750W PSUs	2 x 750W PSUs

3.2.1.1.1 Shared Tier 1 Rack (FC)

In the Shared Tier 1 model, desktop or session-based VMs 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.

R730 Shared T1	Compute	Management
CPU	2 x E5-2698v4 (20C, 2.2GHz)	2 x E5-2660v4 (14C, 2.0GHz)
Memory	16 x 32GB 2400MT/s RDIMMs **Effective speed: 2400MT/s @ 512GB**	8 x 16GB 2400MT/s RDIMMs **Effective speed: 2400MT/s @ 128GB**
Storage Ctrls	PERC H730P – RAID1	PERC H330 – RAID1
Storage	2 x 300GB 15K SAS **(Hypervisor)**	2 x 300GB 15K SAS **(Hypervisor)**
Network	4 x 10Gb SFP+ (BT options available) 2 x QLogic 2562 8Gb DP FC HBA	4 x 10Gb SFP+ (BT options available) 2 x QLogic 2562 8Gb DP FC HBA
iDRAC	iDRAC8 Enterprise	iDRAC8 Enterprise
Power	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.

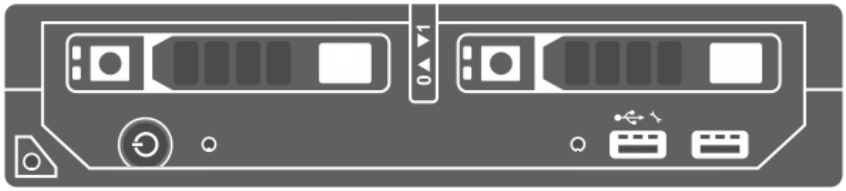

For more information on the PowerEdge M630, please visit: Link

NOTE: Configurations shown below are recommendations for VDI based on stated density determined by our testing. They do not represent absolute platform maximums and can be adjusted as needed.

3.2.2.1 Local Tier 1 Blade

In the Local Tier 1 model for blades, VDI desktops or RDSH sessions execute on local high-performance SSDs on each compute host. 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 pooled, non-persistent desktops. Refer to the Local Tier 1 for Blade Servers section for solution cabling implications.

M630 Local T1	Compute	Management
CPU	2 x E5-2698v4 (20C, 2.2GHz)	2 x E5-2660v4 (14C, 2.0GHz)
Memory	16 x 32GB 2400MT/s RDIMMs **Effective speed: 2400MT/s @ 512GB**	8 x 16GB 2400MT/s RDIMMs **Effective speed: 2400MT/s @ 128GB**
Storage Ctrls	PERC H730P – RAID1	PERC H330 – RAID1
Storage	2 x SSD **(Hypervisor + VMs)**	2 x SSD or 15K SAS **(Hypervisor)**
Network	QLogic 57810S-k 10Gb DP KR NDC	QLogic 57810S-k 10Gb DP KR NDC 1 x QLogic QME2572 8Gb FC mezz
iDRAC	iDRAC8 Enterprise	iDRAC8 Enterprise

Note: This configuration only possible by installing OS and virtual machines on same set of disks and RAID controller which may impact performance. Compute disk size depends on total capacity requirements.

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.

M630 Shared T1	Compute	Management
CPU	2 x E5-2698v4 (20C, 2.2GHz)	2 x E5-2660v4 (14C, 2.0GHz)
Memory	16 x 32GB 2400MT/s RDIMMs **Effective speed: 2400MT/s @ 512GB**	8 x 16GB 2400MT/s RDIMMs **Effective speed: 2400MT/s @ 128GB**
Storage Ctrls	PERC H330 – RAID1	PERC H330 – RAID1
Storage	2 x SSD or 15K SAS **(Hypervisor)**	2 x SSD or 15K SAS **(Hypervisor)**
Network	QLogic 57810S-k 10Gb DP KR NDC 1 x QLogic QME2572 8Gb FC mezz	QLogic 57810S-k 10Gb DP KR NDC 1 x QLogic QME2572 8Gb FC mezz
iDRAC	iDRAC8 Enterprise	iDRAC8 Enterprise

3.3 Storage

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

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 Microsoft RDS 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.

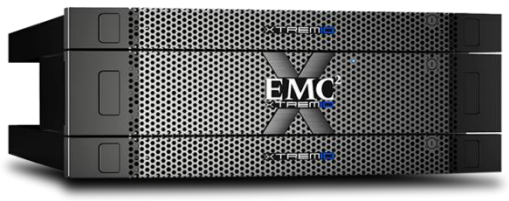

The XtremIO all-flash storage array has a revolutionary architecture with the following elements to enable the agile data center: linear scale-out, inline all-the-time data services, and ample data center services for the workloads.

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

NOTE: Density is based on empirical data and has not been validated with Windows Server 2016 with Hyper-V.

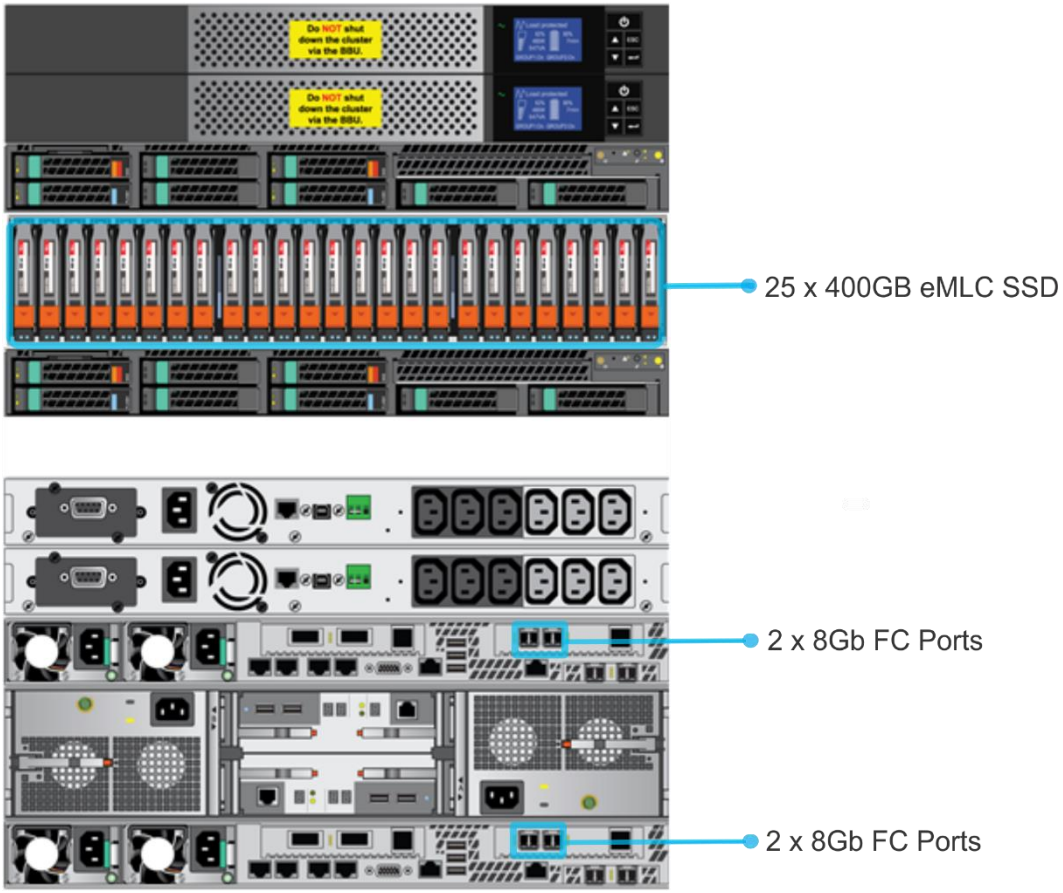

3.3.2 Unity 300 – Tier 2

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 and file with concurrent support for native NAS, iSCSI, and Fibre Channel protocols. Each system leverages dual storage processors, full 12Gb 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).

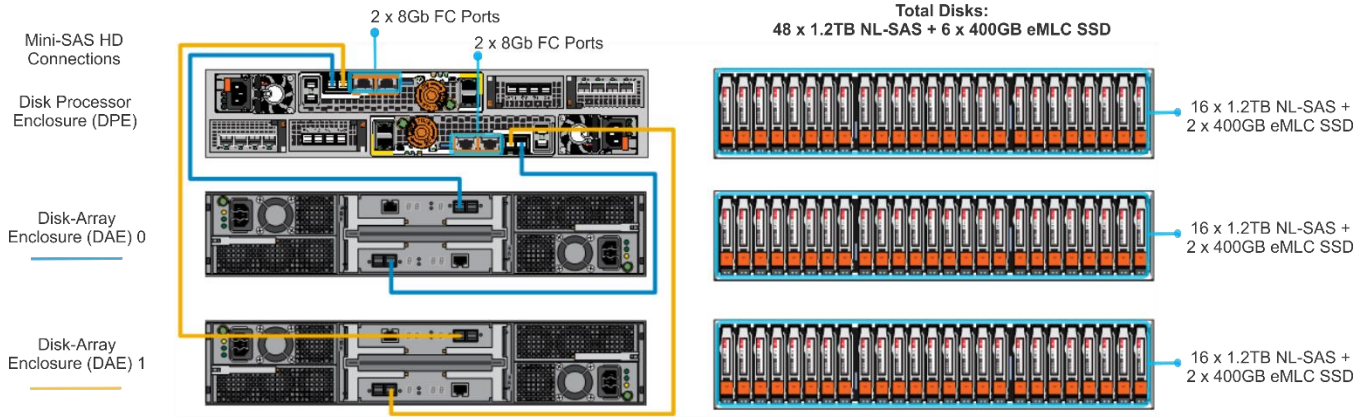

NOTE: Density is based on empirical data and has not been validated with Windows Server 2016 with Hyper-V.

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

For 500 users or less, customers can deploy both 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, user data can be stored via file server VMs with disks stored on the array. For deployments with more than 500 users, we recommend utilizing discrete Unity arrays to provide Tier 2 services as described in the prior section.

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)

NOTE: Density is based on empirical data and has not been validated with Windows Server 2016 with Hyper-V.

3.4 Dell Wyse Endpoints

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

3.4.1 Wyse 3030 LT Thin Client (ThinOS, ThinLinux)

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 (with or without PCoIP), for environments in which security is critical—there’s no attack surface to put your data at risk. The 3030 LT delivers outstanding performance based on its dual core Intel Celeron 1.58GHz processor, 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

3.4.2 Wyse 3040 Thin Client (ThinOS, ThinLinux)

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

3.4.3 Wyse 5040 AIO Thin Client (ThinOS)

The Dell Wyse 5040 AIO all-in-one (AIO) thin client runs ThinOS (with or without PCoIP), has a 21.5" Full HD display and offers versatile connectivity options for use in a wide range of industries. With four USB 2.0 ports, Gigabit Ethernet and integrated dual band Wi-Fi options, users can link to their peripherals and quickly connect to the network while working with processing-intensive, graphics-rich applications. Built-in speakers, a camera and a microphone make video conferencing and desktop communication simple and easy. It even supports a second attached display for those who need a dual monitor configuration. A simple one-cord design and out-of-box automatic setup makes deployment effortless while remote management from a simple file server, Wyse Device Manager (WDM), or Wyse Management Suite can help lower your total cost of ownership as you grow from just a few thin clients to tens of thousands. For more information, please visit: Link

3.4.4 Wyse 5060 Thin Client (ThinOS, ThinLinux, WES7P, WIE10)

The Wyse 5060 offers high performance and reliability, featuring all the security and management benefits of Dell thin clients. It come with flexible OS options: ThinOS (with or without PCoIP), ThinLinux, Windows Embedded Standard 7P (WES7P) or Windows 10 IoT Enterprise (WIE10). Designed for knowledge workers demanding powerful virtual desktop performance, and support for unified communications solutions like Skype for Business, the Wyse 5060 thin client delivers the flexibility, efficiency and security organizations require for their cloud environments. It is powered by a quad-core AMD 2.4GHz processor, supports dual 4K (3840x2160) monitors and provides multiple connectivity options with six USB ports, two of which are USB 3.0 for high-speed peripherals, as well as two DisplayPort connectors, wired networking or wireless 802.11 a/b/g/n/ac. The Wyse 5060 can be monitored, maintained, and serviced remotely via Wyse Device Manager (WDM), cloud-based Wyse Management Suite or Microsoft SCCM (5060 with Windows versions). Customers choosing WIE10 licenses can save about $50/device/year as WIE10 qualifies under Microsoft Software Insurance, without the need to have more expensive VDA licenses to connect to a Windows virtual desktop. For more information, please visit: Link

3.4.5 Wyse 7020 Thin Client (WES 7/7P/8, WIE10, ThinLinux)

The versatile Dell Wyse 7020 thin client is a powerful endpoint platform for virtual desktop environments. It is available with Windows Embedded Standard 7/7P/8 (WES), Windows 10 IoT Enterprise (WIE10), Wyse ThinLinux operating systems and it 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 delivers a great user experience and support for local applications while ensuring security. 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 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 (it only consumes about 15 watts) 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. Customers choosing WIE10 licenses can save about $50/device/year as WIE10 qualifies under Microsoft Software Insurance, without the need to have more expensive VDA licenses to connect to a Windows virtual desktop. For more information, please visit Link

3.4.6 Wyse 7040 Thin Client (WES7P, WIE10)

The Wyse 7040 is a high-powered, ultra-secure thin client running Windows Embedded Standard 7P (WES7P) or Windows 10 IoT Enterprise (WIE10) operating systems. Equipped with an 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 WES7P 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). Customers choosing WIE10 licenses can save about $50/device/year as WIE10 qualifies under Microsoft Software Insurance, without the need to have more expensive VDA licenses to connect to a Windows virtual desktop. For more information, please visit: Link

3.4.7 Latitude E7270 mobile with client (WES7P)

The high-powered Latitude E7270 mobile thin client runs Windows Embedded Standard 7P (WES7P) and is designed to securely deliver virtual desktops to mobile users. Equipped with a powerful Intel i5 quad core processor with an integrated HD 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 12.5'' 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 ensuring a fast, rich end user experience. The Latitude E7270 mobile thin client is easily manageable through Wyse Device Manager (WDM), Wyse Management Suite and Microsoft's System Center Configuration Manager (SCCM). For more information, please visit: Link

Enhanced Security

Note that all the above thin clients running Windows Embedded Standard 7 or Windows 10 IoT can be protected against viruses, ransomeware and zero-day threats by installing **Dell Threat Defense**, a revolutionary anti-malware software solution using artificial intelligence and mathematical modeling and is not signature-based. Threat Defense prevents 99% of executable malware, far above the average 50% of threats identified by the top anti-virus solutions. It doesn't need a constant internet connection nor frequent updates (only about twice a year), it only uses 1-3% CPU and has only a ~40MB memory footprint, making it an ideal choice to protect thin clients without impacting the end user productivity.

If you also want to protect **virtual desktops** against such malware and threats with a similar success, Dell recommends using **Dell Endpoint Security Suite Enterprise**, a full suite featuring advanced threat prevention and data-centric encryption using an on-premise management console. This suite can also be used to protect physical PCs, MAC OS X systems and Windows Server.

4 Software Components

4.1 Microsoft

4.1.1 Remote Desktop Services

Microsoft Remote Desktop Services (RDS) accelerates and extends desktop and application deployments to any device, improves remote worker efficiency, and helps secure critical intellectual property while simplifying regulatory compliance. Remote Desktop Services enables virtual desktop infrastructure (VDI), session-based desktops, and applications, allowing users to work anywhere.

The core RDS components include:

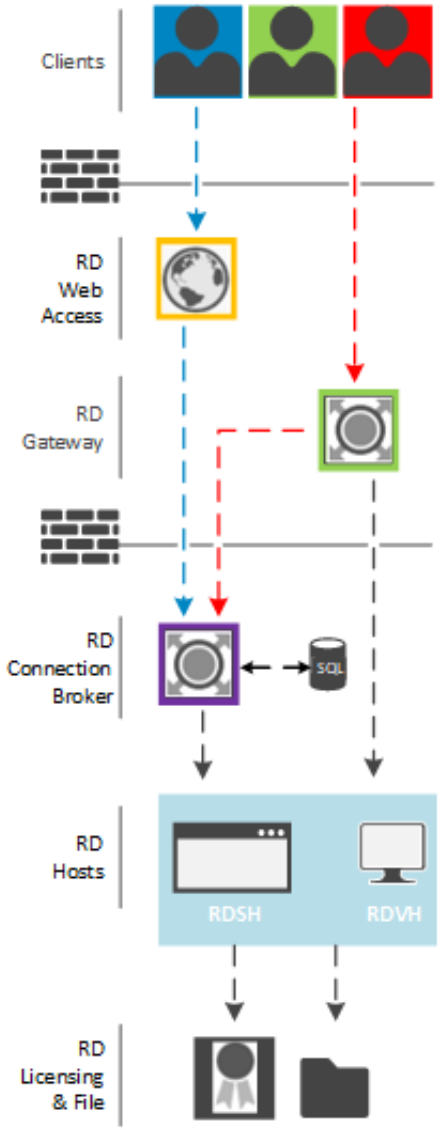

- **Remote Desktop Connection Broker (RDCB)**

Remote Desktop Connection Broker (RD Connection Broker) allows users to connect to virtual desktops (RDVH) or RemoteApp programs and session-based desktops (RDSH) while evenly distributing the load to their respective collections.

- **Remote Desktop Gateway (RDG)**

Remote Desktop Gateway (RD Gateway) enables authorized users to connect to virtual desktops, RemoteApp programs, and session-based desktops on an internal corporate network from any Internet-connected device.

- **Remote Desktop Web Access (RDWA)**

Remote Desktop Web Access (RD Web Access) enables users to access RemoteApp and Desktop Connections through the Start menu on a computer that is running Windows 7 through Windows 10, or through a web browser. RemoteApp and Desktop Connection provides a customized view of RemoteApp programs and session-based desktops in a session collection, and RemoteApp programs and virtual desktops in a virtual desktop collection.

- **Remote Desktop Virtualization Host (RDVH)**

Remote Desktop Virtualization Host (RD Virtualization Host) integrates with Hyper-V to deploy pooled or personal virtual desktop collections within your organization.

- **Remote Desktop Session Host (RDSH)**

Remote Desktop Session Host (RDSH) enables a server to host RemoteApp programs or session-based (personal or shared)

desktops. Users can connect to RD Session Host servers in a session collection to run programs, save files, and use resources on those servers.

- **Database**

  Although RDS can be configured using an internal database, a Microsoft SQL Server database is recommended to store configuration and session information. To implement broker high availability, SQL Server is required. Alternatively, an Azure SQL database can be used. The RDCB must have a persistent connection to the database as it stores data collected and managed by the RDS services.

- **Remote Desktop Licensing**

  Remote Desktop Licensing (RD Licensing) manages the licenses required to connect to a Remote Desktop Session Host server or a virtual desktop. You can use RD Licensing to install, issue, and track the availability of licenses.

RDS deployment terms and options include:

- **Collections**

  As the name suggests, a collection is a group or pool of virtual machines that the broker will connect users to. A collection can be **managed** meaning the RDCB creates and maintains the collection using a template VM including recreating VMs as needed or a collection can be **unmanaged** meaning the RDCB can broker connections to the pool but there is no template VM. The RDCB does not create or manage the VMs in an unmanaged collection.

- **Pooled Virtual Desktops**

  Pooled virtual desktops are based off of a template VM and are reverted back to a pristine state when users log off making them non-persistent. They are thin provisioned, checkpointed, and only consume a fraction of the storage used by the original template VM. User Profile Disks (UPD) can be used to store user profile data so that it's available the next time a user logs on to a pooled desktop.

- **Personal Virtual Desktops**

  Personal virtual desktops are based off of a template VM but changes made by the user are stored to the VM virtual disk and available after their session ends. Users are reconnected to the same VM each time creating a persistent experience.

- **Session-based virtualization**

  Users connect to RDSH servers and run their applications or desktops in Windows Server 2016 sessions (multiple users connecting to the same server). Session-based desktops can be configured as **shared** where no session info is retained after logging off or as **personal** where users are reconnected to the same RDSH server allowing them to make changes to and save their session info.

For additional information about RDS, please visit: LINK

4.1.2 Licensing

Several licensing aspects must be taken in to account when deploying a Microsoft RDS environment including:

- Windows Server 2016 Hyper-V
- Windows Server CALs
- RDS CALs
- Software Assurance (SA) / Virtual Desktop Access (VDA) client licensing

Windows Server 2016 Hyper-V

Starting with Windows Server 2016, all physical CPU cores in a host must be licensed. Once all cores are licensed, customers are entitled to use 2 virtual machines with Standard Edition or unlimited virtual machines with Datacenter Edition. For Standard Edition, 2 additional VMs can be entitled by licensing all CPU cores again. For example, on a server with a total of 40 cores, 10 virtual machines can be used with Standard Edition by purchasing 200 core licenses (10 divided by 2 = 5, 5 x 40 cores = 200 core licenses). For VDI use where large numbers of VMs per host are possible, Datacenter Edition will typically be the most economical choice.

For additional details, Windows Server 2016 Licensing Datasheet found here: LINK

Windows Server and RDS CALs

Windows 2016 licensing follows a core + CAL licensing model. Any user/device accessing a Datacenter or Standard edition of Windows Server requires a Windows Server CAL. The Server CALs are considered the "base" CAL while RDS CALs are additive. Therefore, for a Microsoft RDS deployment, all users/devices accessing RDS require a Server CAL + a RDS CAL.

SA/VDA Client Licenses

Remotely accessing a VM running a Windows desktop OS requires that the user/device have a Software Assurance (SA) or Virtual Desktop Access (VDA) license to access the virtualized Windows desktop OS. Endpoints with an OEM licensed Windows desktop operating system are generally not licensed to access a VDI environment and require an appropriate SA/VDA license to do so. Certain endpoints, such as those running a non-Windows OS, will require a VDA license. Some endpoints, such as Windows embedded Thin Clients, may already include a VDA license with their purchase. Neither license type is required when accessing session-based desktops as the underlying OS for RDSH session-based desktops is Windows Server 2016 and not an actual desktop OS.

NOTE: For exact licensing details including pricing, please contact a Microsoft Software Licensing specialist at Dell.

4.1.3 RDSH

Compute hosts with the RDSH role provide easy access to a densely shared session environment. Each RDP-based session shares the total available server resources with all other sessions logged in concurrently on the server. An RDS CAL and Server CAL is required for each user/device accessing this type of

environment but neither SA nor VDA licenses are required since the underlying OS for RDSH VMs is Server based instead of Desktop based.

Benefits of hosted desktop sessions and applications:

- Management of applications (single instance)
- Management of simple desktop images (no applications installed)
- Scalability of compute hosts: CPU and IOPS reduction via application offload
- Shared storage scalability: less IOPS = more room to grow

RDSH 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. Starting with Windows Server 2016, RDSH now provides the ability to use personal session desktops where users always connect to the same RDSH server giving them a persistent desktop experience.

4.1.3.1 RDSH Integration into Dell Wyse Datacenter Architecture

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

4.1.3.2 NUMA Architecture Considerations

Best practices and testing has showed that aligning RDSH design to the physical Non-Uniform Memory Access (NUMA) architecture of the server CPUs results in increased and optimal performance. NUMA alignment ensures that a CPU can access its own directly-connected RAM banks faster than those banks of the adjacent 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 RDSH servers do not span physical NUMA nodes will ensure the greatest possible performance benefit.

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

4.1.3.3 NUMA Alignment

20 physical cores per CPU given the 2698v4 part and 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 RDSH VMs as shown below will ensure that no physical NUMA node spanning occurs which could lower performance for an effected VM.

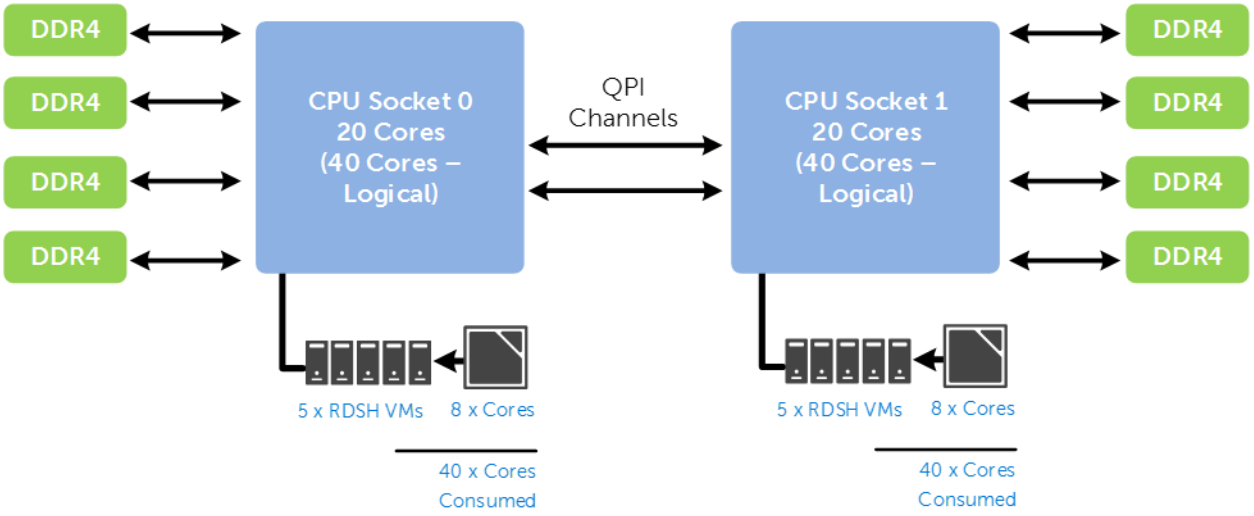

To ensure no NUMA spanning, remove the checkmark for "Allow virtual machines to span physical NUMA nodes" in the Hyper-V settings of the compute host(s). Also ensure "Maximum NUMA nodes allowed on a socket" is set to a value of 1. This setting is found under Processor > NUMA for each RDSH VM.

4.2 Hypervisor Platform

4.2.1 Microsoft Windows Server 2016 Hyper-V

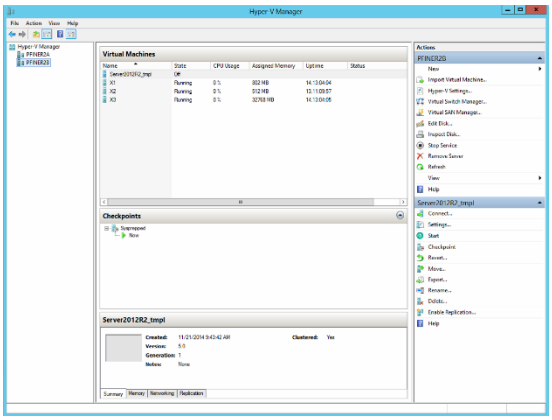

Hyper-V ™ is a powerful virtualization technology that enables businesses to leverage the benefits of virtualization. Hyper-V reduces costs, increases hardware utilization, optimizes business infrastructure, and improves server availability. Hyper-V works with virtualization-aware hardware to tightly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

4.3 Optional Management Platform

4.3.1 Microsoft System Center 2016 Virtual Machine Manager

An optional management platform for a Hyper-V based 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.

It’s important to note that SCVMM doesn’t integrate as seamlessly with RDS as one might expect. Service templates can be used with SCVMM to build out, maintain, and scale the RDS infrastructure components such as the connection broker (RDCB) and SQL. When it comes to compute VM collections, SCVMM can be used as a provisioning & management tool for unmanaged collections but doesn’t integrate directly with the RDS farm or RDCB. Therefore, SCVMM can’t be used to provision any VM intended to exist within a managed pool owned by the RDCB. As stated in the Remote Desktop Services section, a managed collection is a pool created and maintained by the RDCB by using a template VM in RDS which allows the RDCB to recreate and provision more VMs based on the template as needed.

SCVMM can be used with unmanaged collections in RDS but it will require some manual steps which can be handled via PowerShell. RDCB can broker connections to VMs in an unmanaged collection but it can’t be used to create/maintain them. However, SCVMM can manage the VMs in the collection and make use of features like Intelligent Placement as well as handling virtual MAC assignments which can become an issue on a server hosting a large number of VMs. In order to do so, create the unmanaged collection in RDS first. Next, create/clone the desired number of VMs using SCVMM. Cloning multiple copies of VMs is possible in SCVMM via PowerShell. Once the VMs are created, they must be manually added to the unmanaged collection via PowerShell or Server Manager. Alternatively, SCVMM service templates can be used to also provision the compute VMs. There are many online articles describing how to configure these templates for an RDSH deployment.

5 Solution Architecture for RDS

5.1 Management Role Configuration

5.1.1 Hyper-V

Role	vCPU	Startup RAM (GB)	Dynamic Memory			NIC	OS + Data vDisk	
			Min|Max	Buffer	Weight		Size (GB)	Tier2 Vol
RD Broker & Licensing	4	8	512MB | 10GB	20%	Med	1	60	-
RD Gateway & Web Access	4	4	512MB | 10GB	20%	Med	2	60	-
Primary SQL	8	8	512MB | 16GB	20%	Med	1	60	300 (VHDX)
SCVMM (optional)	8	8	512MB | 16GB	20%	Med	1	60	50 (VHDX)
File Server	2	4	512MB | 6GB	20%	Med	1	60	2048 (PTD)
Total	**30**	**32GB**	**2.5GB | 58GB**	**-**	**-**	**7**	**300GB**	**2398GB**

NOTE: SCVMM is optional but a recommended configuration is provided here in case it will be used.

5.1.2 RDSH VM Configuration

The recommended number of RDSH VMs and their configurations on 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.

RDSH 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
RDSH VM	10	8	32	16GB | 48GB	20%	Med	1	80	Tier 1

5.1.3 SQL Databases

The Microsoft databases are hosted by a single dedicated SQL 2016 (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 databases for:

- Microsoft RDS
- SCVMM (if using)

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 Microsoft are to be adhered to, to ensure optimal database performance.

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

5.1.4 DNS

DNS plays a crucial role in the environment not only as the basis for Active Directory but is used to control access to the various 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, RDS services) during the initial deployment. Use CNAMEs and the round robin DNS mechanism to provide a front-end “mask” to the back-end server actually hosting the service or data source.

5.1.4.1 DNS for SQL

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

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

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

5.2 Storage Architecture Overview

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

5.2.1 Local Tier 1 Storage

Selecting the local Tier 1 storage model means that the compute host servers use 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.

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. 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. For example, 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	NTFS
VDI-2	500	Tier 1	125 x desktop VMs	NTFS
VDI-3	500	Tier 1	125 x desktop VMs	NTFS
VDI-4	500	Tier 1	125 x desktop VMs	NTFS

5.2.3 Shared Tier 2 Storage

Tier 2 is shared file storage used to host the Management server VMs and user data. For the Local Tier 1 model, Unity 300 arrays are used for the Tier 2 data in order for the management server VMs to stay highly available. For smaller deployments in a Shared Tier 1 model (<=500 users), the XtremIO Starter X-Brick can host both Tier 1 and Tier 2 data. For Shared Tier 1 with more than 500 users, distinct Unity 300 arrays are used for Shared Tier 2. 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 is presented via file server VM(s). The

solution as designed presents all SQL disks using 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	300	Tier 2	RDS roles, SCVMM (optional), File & SQL	NTFS
User Data	2048	Tier 2	File Server	PTD
User Profiles	200	Tier 2	User profiles	NTFS
SQL DATA	100	Tier 2	SQL	NTFS
SQL LOGS	100	Tier 2	SQL	NTFS
TempDB Data	5	Tier 2	SQL	NTFS
TempDB Logs	5	Tier 2	SQL	NTFS
Templates/ISO	200	Tier 2	ISO storage (optional)	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.

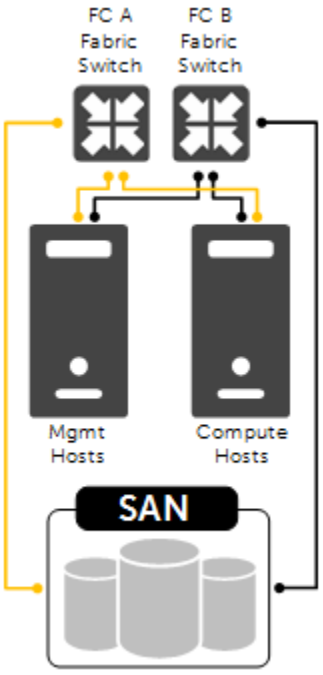

If Fiber Channel is the selected block storage protocol, then the XtremIO Integration for VMware vSphere client plug-in is installed on all hosts. This plugin enables all newly created data stores to be automatically aligned at the recommended 4MB offset. Although a single Fabric are configured to begin 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

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.

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.

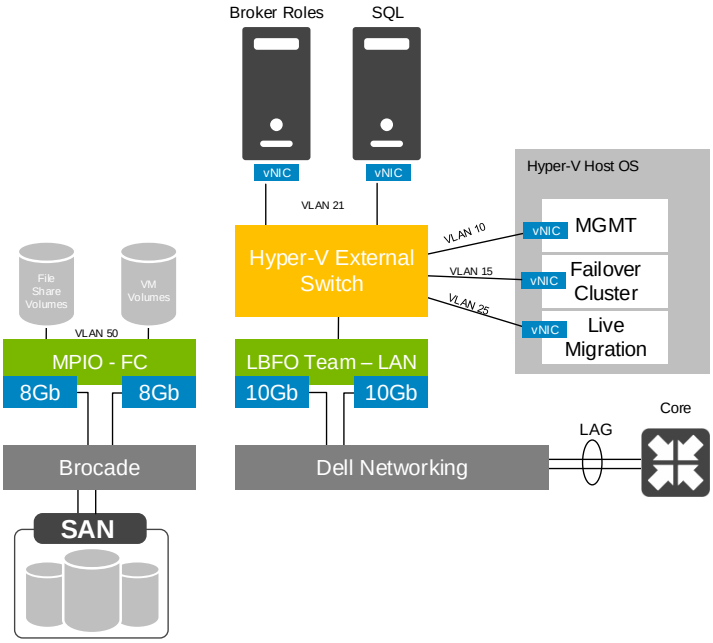

The dedicated compute host configuration is shown in the diagram below and configured similarly to the management host configuration but without the features enabled by shared storage.

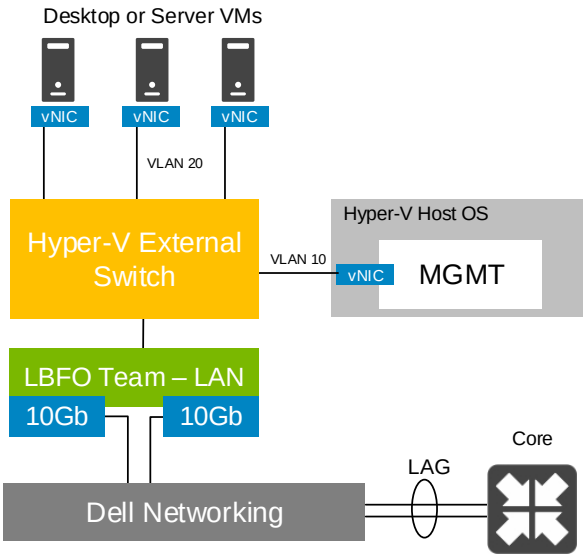

5.3.2 Shared Tier 1 – FC

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.

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

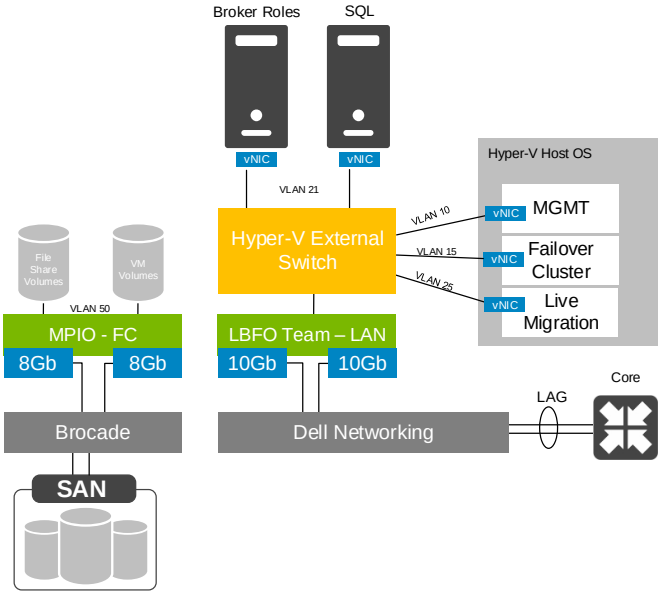

Compute Servers

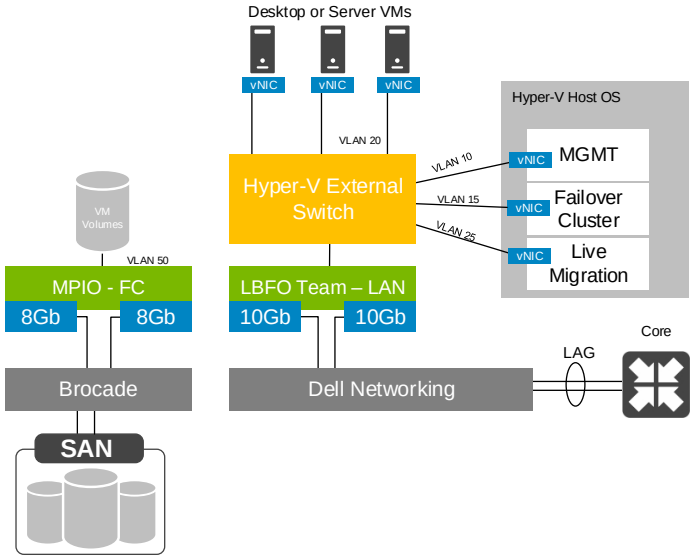

5.4 Scaling Guidance

- 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
Broker Servers	Desktops per instance (dependent on SQL performance as well)	Additional servers added to the farm	Additional virtual machine resources (RAM and CPU)
RDSH Servers	Apps/Desktops per instance	Additional virtual servers added to the collection	Additional physical servers to host virtual RDSH servers
RD Gateway/Web Access Servers	Logons/ minute	Additional servers added to the farm	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 an optional NAS device to provide high availability.	Additional RAM and CPU for the management nodes

Additional scalability details specific to Windows Server 2016 Hyper-V can be found here: LINK

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, Hyper-V clustering is introduced in both the management and compute layers, SQL is configured with AlwaysOn and native load balancing.

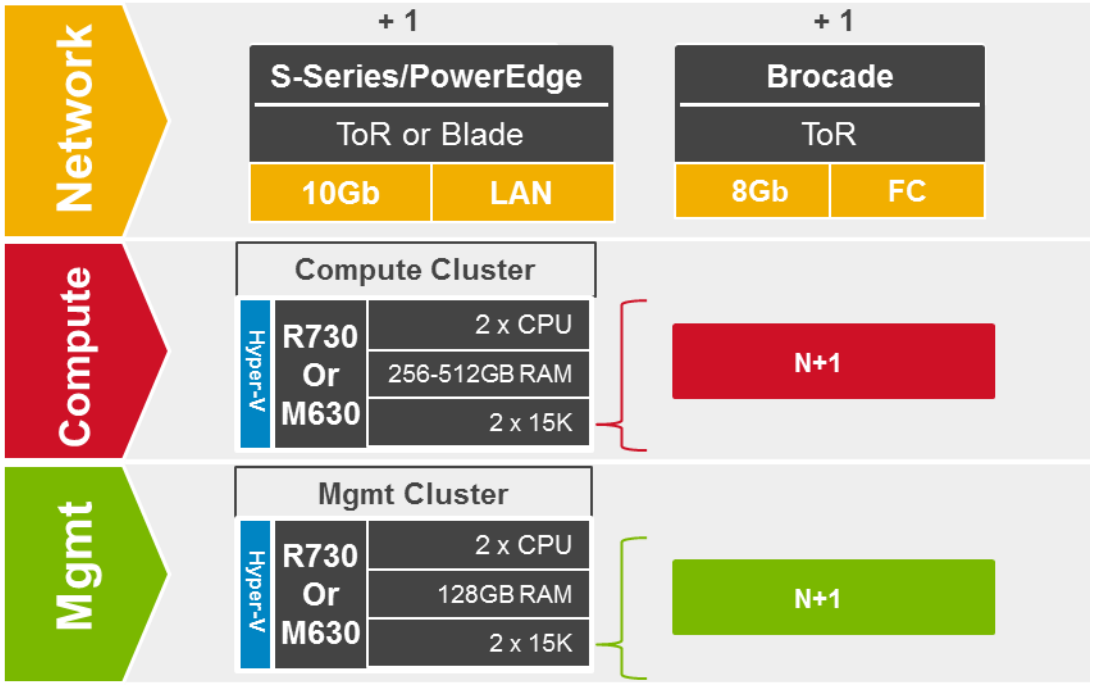

The HA options provide redundancy for all critical components in the stack while improving the performance and efficiency of the solution as a whole.

- Additional switches added to the existing thereby equally spreading each host's network connections across multiple switches.
- Additional Hyper-V hosts added in the compute or management layers to provide N+1 protection.
- Applicable RDS infrastructure server roles are duplicated and spread amongst management host instances where connections to each are load balanced.
- 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.

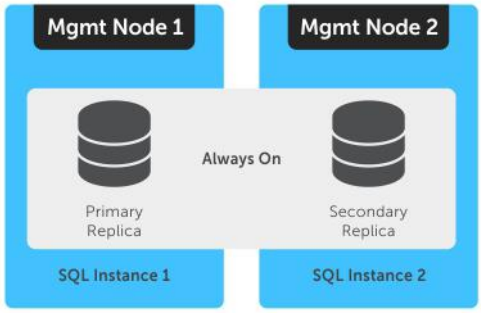

Please refer to these links for more information: LINK1 and LINK2

5.6 Microsoft RDS Communication Flow

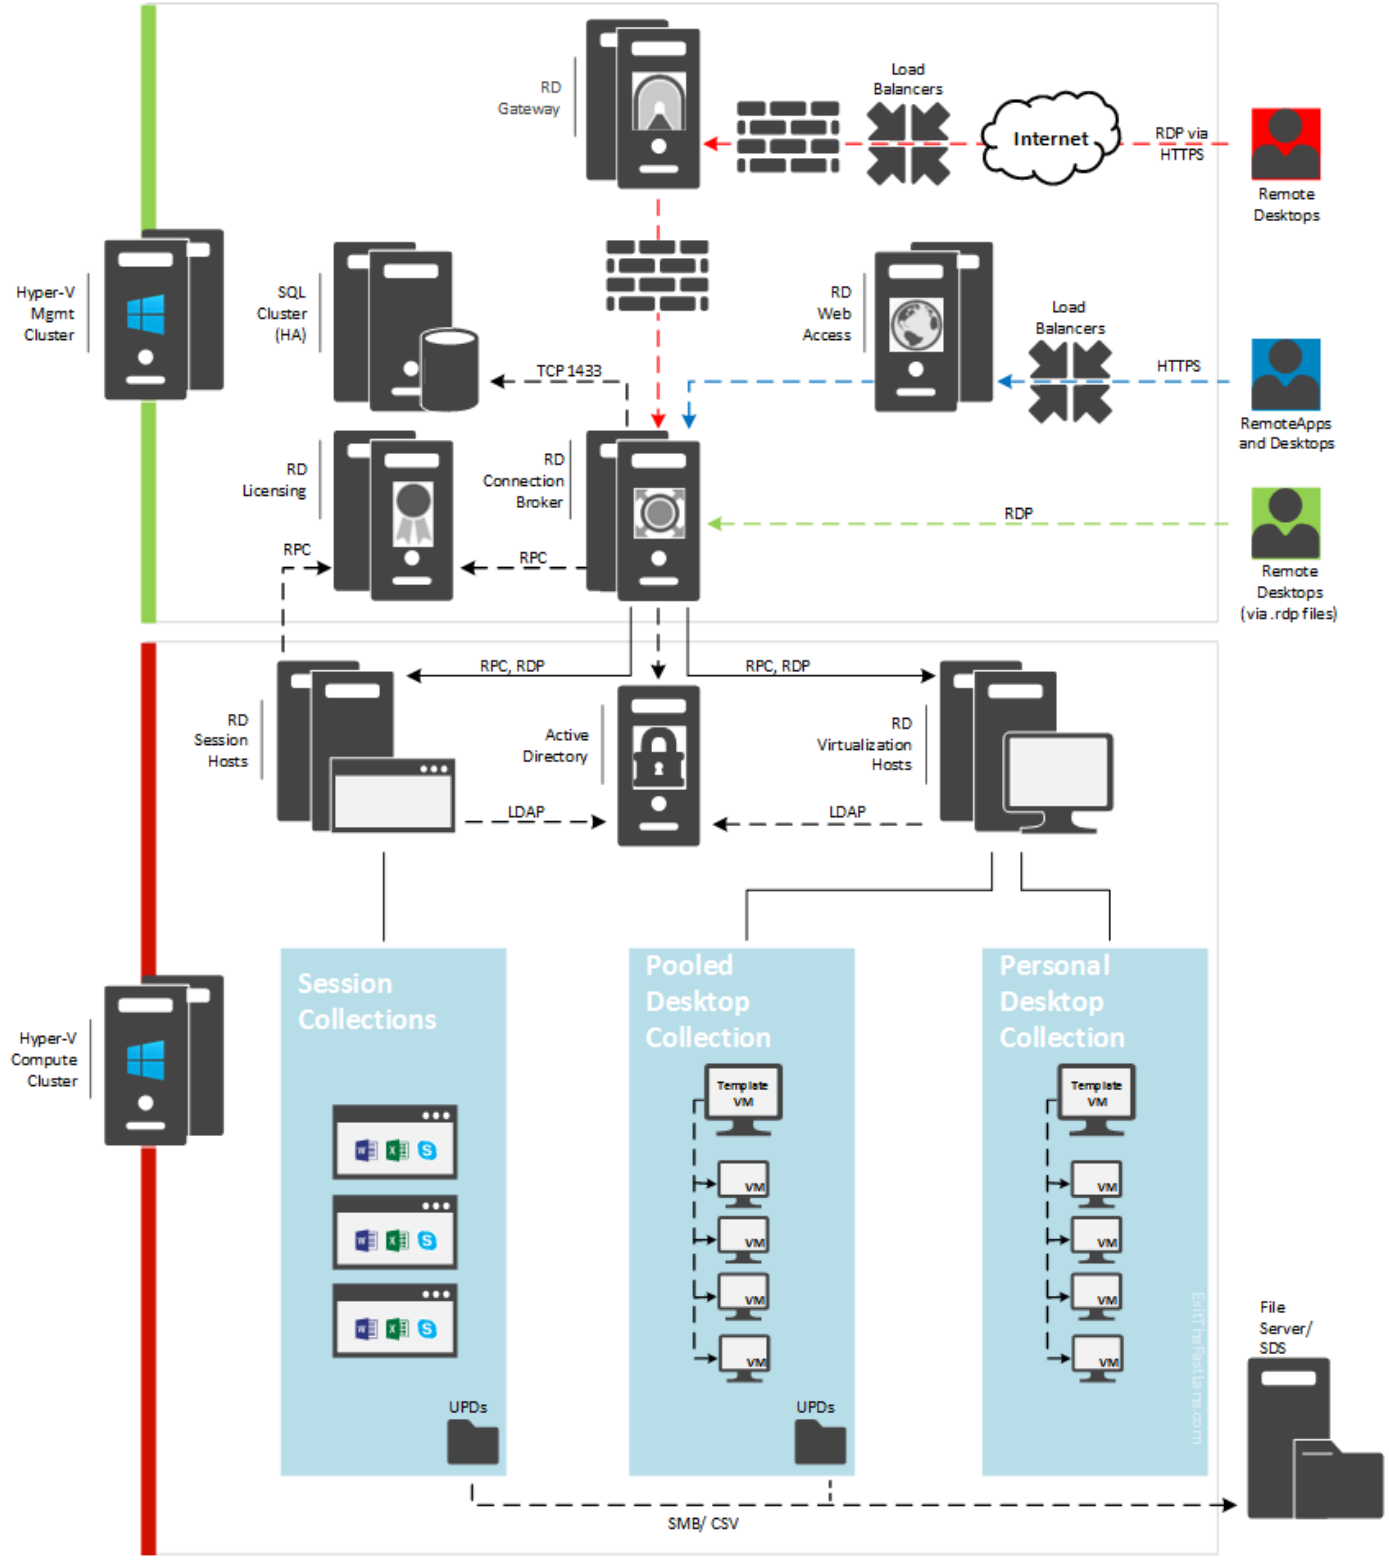

6 Customer-provided solution components

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	Deployment / Provisioning	Workload	Template	User Density
PER730 w/2698v4 CPUs	Hyper-V 2016	RDVH – Pooled Desktops	Task Worker	Windows 10 x64 & Office 2016	325
PER730 w/2698v4 CPUs	Hyper-V 2016	RDVH – Pooled Desktops	Knowledge Worker	Windows 10 x64 & Office 2016	270
PER730 w/2698v4 CPUs	Hyper-V 2016	RDVH – Pooled Desktops	Power Worker	Windows 10 x64 & Office 2016	250
PER730 w/2698v4 CPUs	Hyper-V 2016	Session-based Desktops (RDSH)	Knowledge Worker	Windows 2016 x64 & Office 2016	350

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 (graphics currently not tested with this solution). 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

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.

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 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.

7.1.3 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 (Hyper-V)	85%
Physical Host Memory Utilization	90%
Network Throughput	85%
Storage IO Latency	20ms
LVSI Failed/Unresponsive Sessions	5%

NOTE: The CPU utilization threshold is set to 85% since the additional headroom provided by the processor Turbo Boost feature is reported by Performance Monitor.

NOTE: Since system memory is 512GB, memory threshold has been adjusted to 90% as this still leaves sufficient room for use.

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)	Windows Server 2016 Hyper-V
Broker technology	RDS
Broker database	Microsoft SQL 2016
Management VM OS	Windows Server 2016
Virtual desktop OS	Windows 10 Enterprise 64-bit
Office application suite	Office Professional 2016
Login VSI test suite	Version 4.1.25

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	Hyper-V Startup Memory	Hyper-V Min | Max Dynamic	Screen Resolution	Operating System
Task Worker	1	1GB	1GB | 2GB	1280 X 720	Windows 10 Enterprise 64-bit
Knowledge Worker	2	1.5GB	1GB | 3GB	1920 X 1080	Windows 10 Enterprise 64-bit
Power Worker	2	2GB	1GB | 4GB	1920 X 1080	Windows 10 Enterprise 64-bit

RDSH VM specifications

Platform Config	vCPUs	Hyper-V Startup Memory	Hyper-V Min | Max Dynamic	Operating System
RDSH	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

Note: 2 x 10Gb networking was used for VDI traffic.

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	Deployment / Provisioning	Login VSI Workload	Density per Host	Avg CPU	Avg Mem Consumed	Avg IOPS / User	Avg Net Mbps / User
R730	Hyper-V 2016	RDVH – Pooled Desktops	Task Worker	325	50%	450GB	3.1	.5
R730	Hyper-V 2016	RDVH – Pooled Desktops	Knowledge Worker	270	60%	450GB	5.5	2.2
R730	Hyper-V 2016	RDVH – Pooled Desktops	Power Worker	250	80%	460GB	7.2	4
R730	Hyper-V 2016	Session-based Desktops (RDSH)	Knowledge Worker	350	76%	450GB	1.5	2.2

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.

Avg Consumed Memory: The amount of physical memory used by a host during the steady state phase. For clusters, this is the average consumed memory 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, 325 Users, Hyper-V 2016, RDVH Desktops

The graph below shows the CPU utilization for 325 Task Worker sessions per host. The CPU reaches a steady state average of 50% during the test cycle when 325 users are logged on. A very aggressive boot phase raised the CPU to 100%. However the low CPU usage during the steady state indicates that the CPU is more than capable of the load and there is plenty of room for peaks in usage.

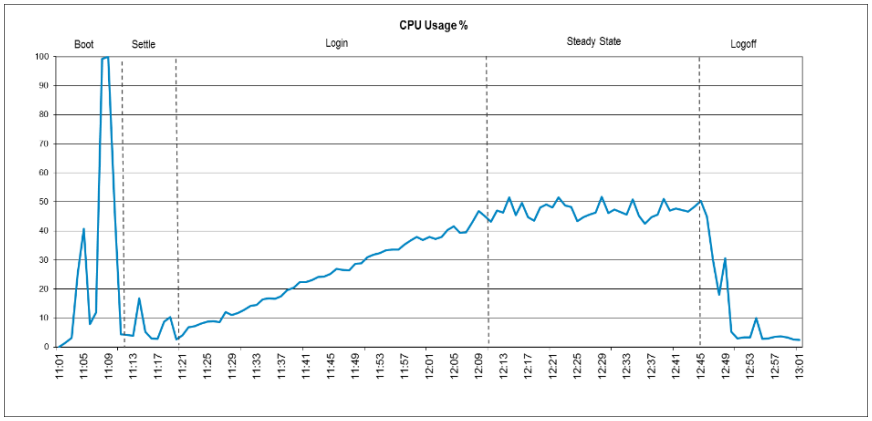

450GB of memory was used during the test to stay within the 90% threshold limit. Although CPU usage and IO latency are low, the system is limited by the memory available.

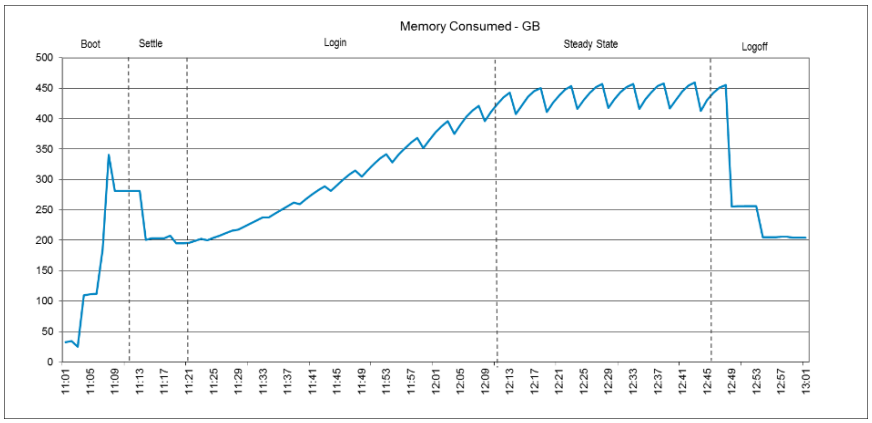

The host reached a maximum of 45,000 IOPs during the boot phase with the latency on the SSD disks remaining less than 1 ms. IOPS peaked at 5,000 during the login phase and dropped to 1,000 in steady state resulting in about 3.1 IOPS per user on average during this phase. In all phases, the latency remained below 2 ms during the test.

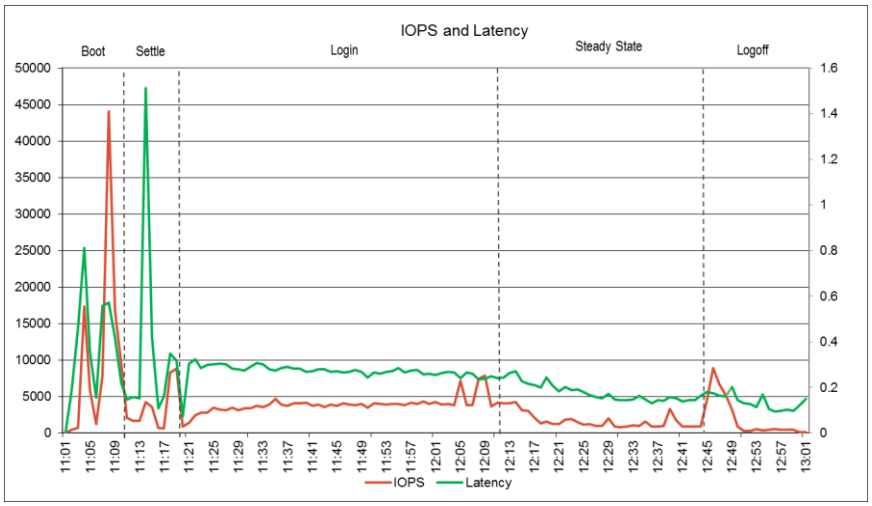

As seen below, network bandwidth is not an issue with a steady state peak of approximately 250 Mbps. Network usage is affected by the Login VSI content and profile and home folder redirection. The output queues on the network adapters remained at zero (0) throughout the testing.

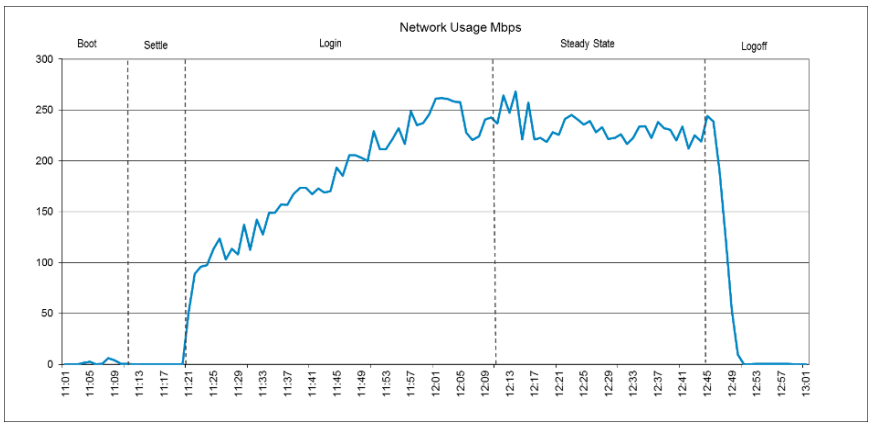

Login VSI VSIMAX was not reached during the test; however, no additional sessions were used since memory utilization hit the threshold limit at 325 users.

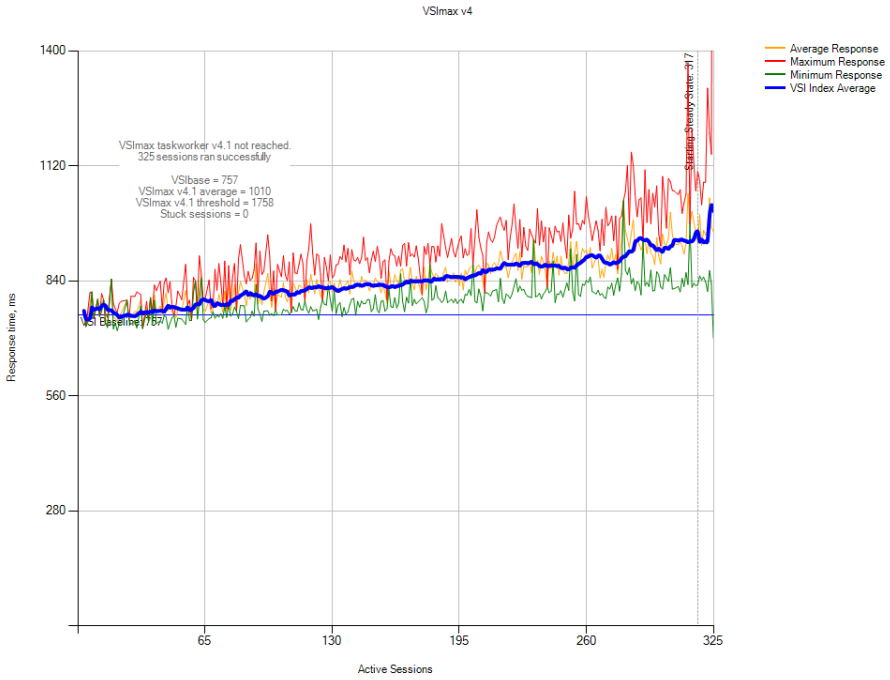

7.3.1.2 Knowledge Worker, 270 Users, Hyper-V 2016, RDVH Desktops

The graph below shows the CPU utilization for 270 Knowledge Worker sessions per host. The CPU reaches a steady state average of 60% during the test cycle when 270 users are logged on. A very aggressive boot phase raised the CPU to 100%. However, the low CPU usage during the steady state indicates that the CPU is more than capable of the load and there is plenty of room for peaks in usage.

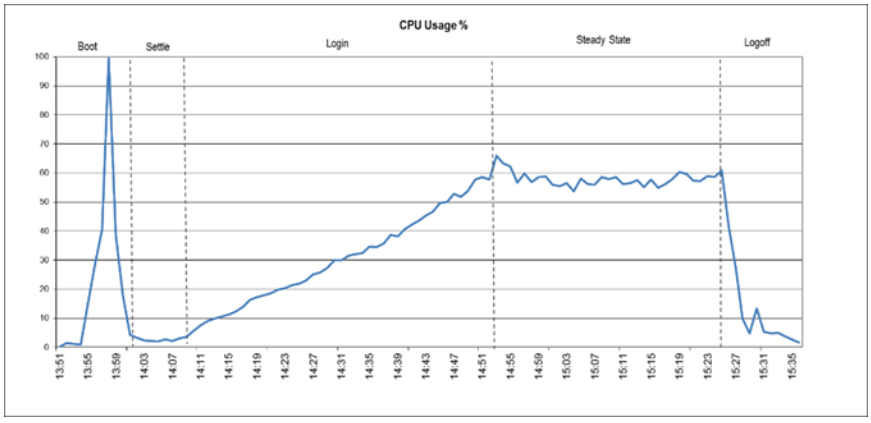

460GB of memory was used during the test to stay within the 90% threshold limit. Although CPU usage and IO latency are low, the system is limited by available memory.

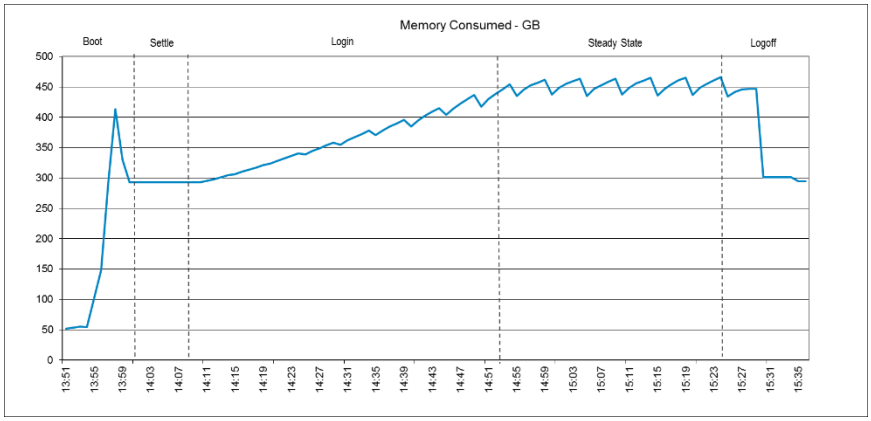

The host reached a maximum of 35,000 IOPs during the boot phase with the latency on the SSD disks remaining less than 1 ms. IOPS peaked at 5,000 during the login phase and dropped to 1,000 - 1,500 in steady state resulting in about 5.5 IOPS per user on average during this phase. In all phases, the latency remained below 1 ms during the test.

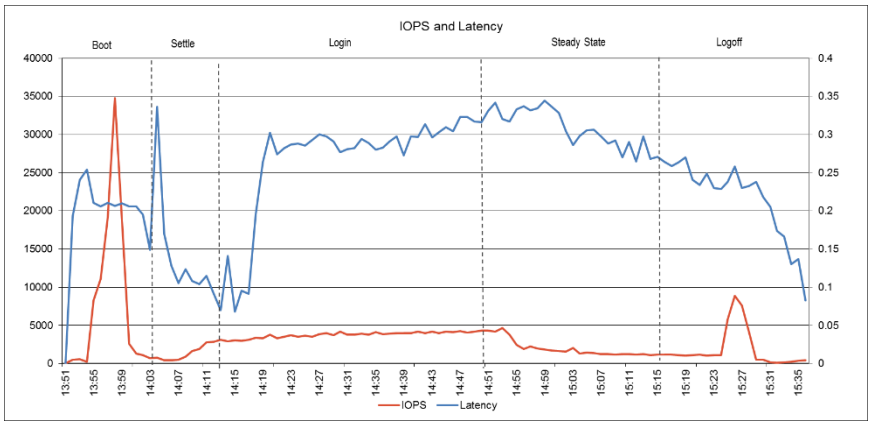

Network bandwidth is not an issue on this test run with a steady state peak of approximately 600 Mbps which is well within the abilities of the 10Gbps network cards. Network usage is affected by the Login VSI content and profile and home folder redirection. The output queues on the network adapters remained at zero (0) throughout the testing.

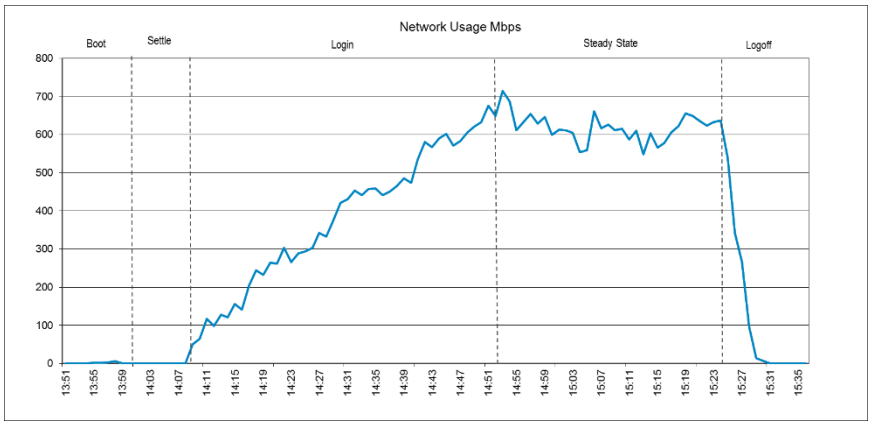

Login VSI VSIMAX was not reached during the test; however, no additional sessions were used since memory utilization hit the threshold limit at 270 users.

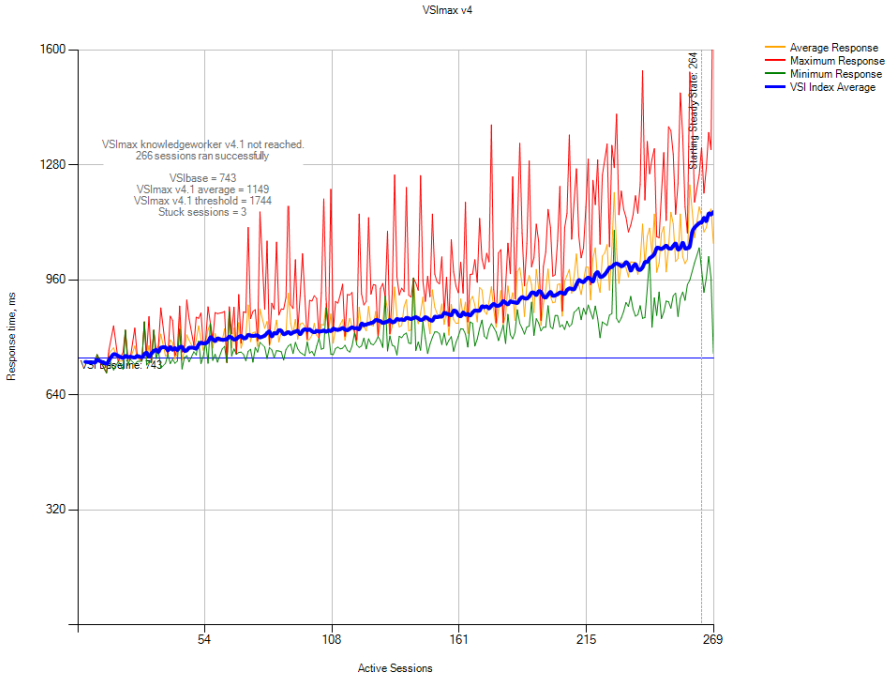

7.3.1.3 Knowledge Worker, 350 Users, Hyper-V 2016, RDSH Desktops

The graph below shows the CPU utilization for 350 Knowledge Worker sessions per host. The CPU reaches a steady state average of 70% during the test cycle when 350 users are logged on.

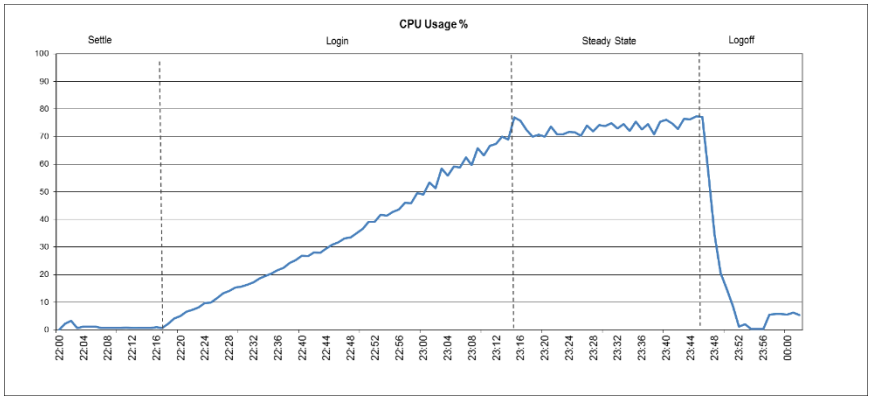

450GB of memory was used during the test stay within the 90% threshold limit. Although CPU usage and IO latency are low, the system is limited by available memory.

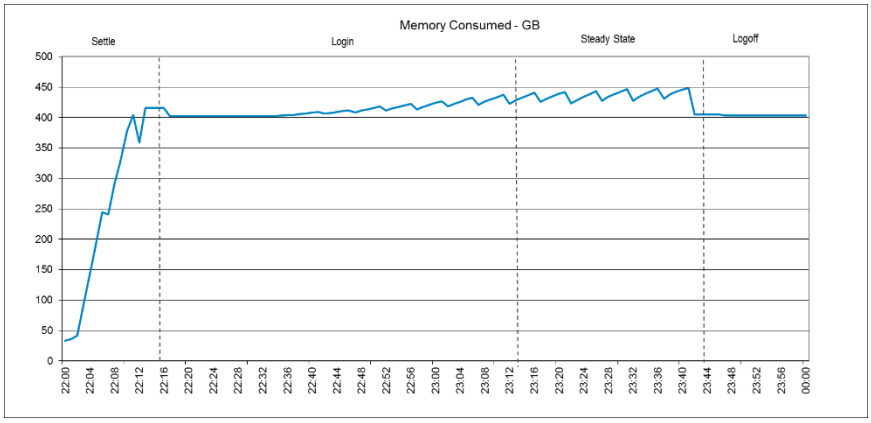

The host reached a maximum of 3,500 IOPs during the boot phase, 2,000 during the login phase and dropped to 1,000 during the steady state resulting in about 1.5 IOPS per user on average during this phase. In all phases, the latency remained below 1ms during the test.

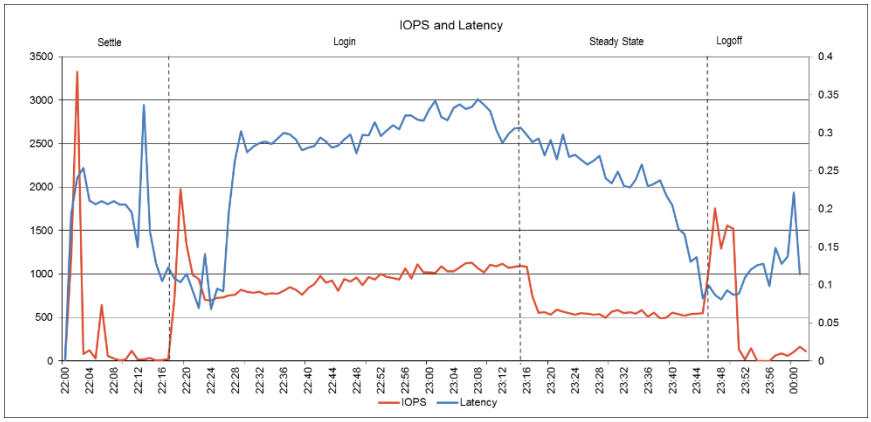

All temporary data was stored on local disk of the compute host. Latency was low throughout the logon steady state and logoff periods.

Network bandwidth is not an issue with a steady state peak of approximately 450 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.

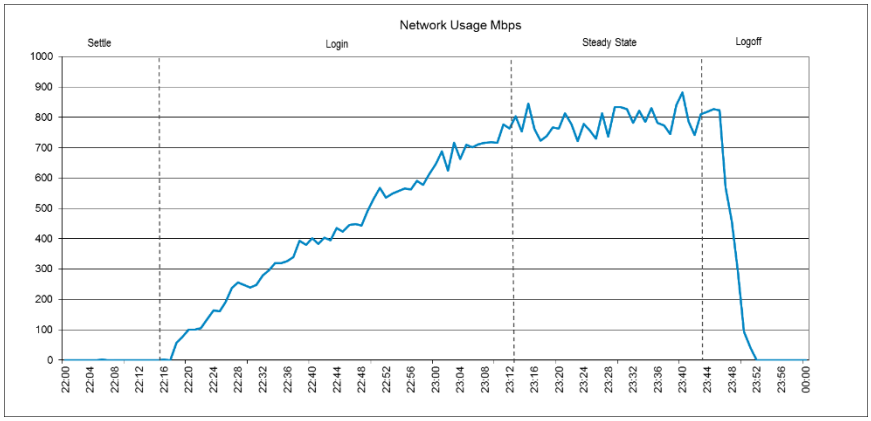

The Login VSI Max user experience score shown below for this test was reached at 345 sessions indicating that the system was correctly sized for this number of users. Subjective user experience showed very good responses during the test. Other indicators form the Login VSI Analyzer show that Logon times remained low as did disk IO latency and application load times. The VSI base line was very good at 701.

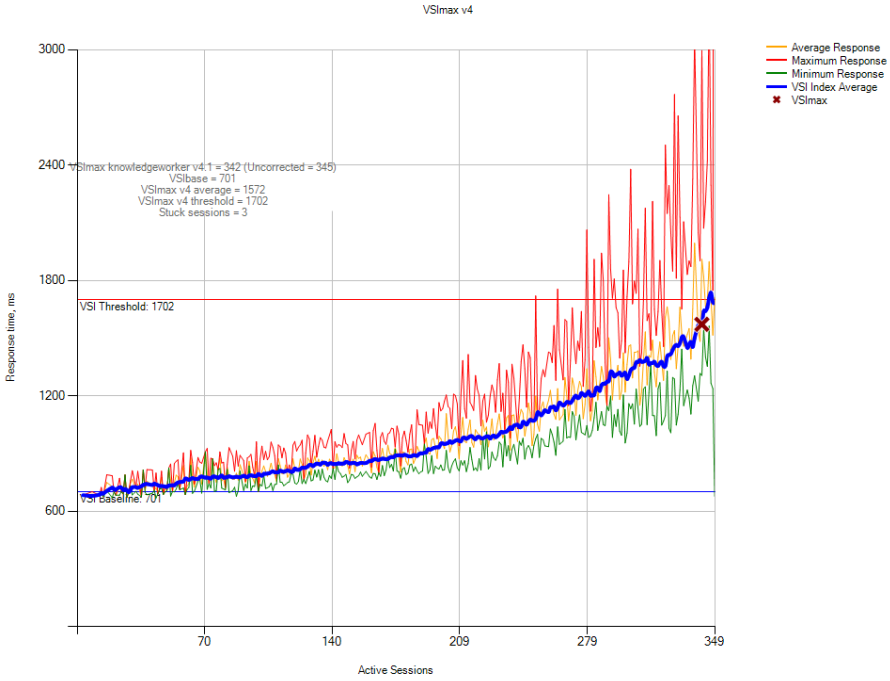

7.3.1.4 Power Worker, 250 Users, Hyper-V 2016, RDVH Desktops

The charts below show the performance of the Power Worker test for 250 users. Again, memory is the limiting factor with CPU still below 80%, VSIMAX not reached, and disk latency less than 1 ms.

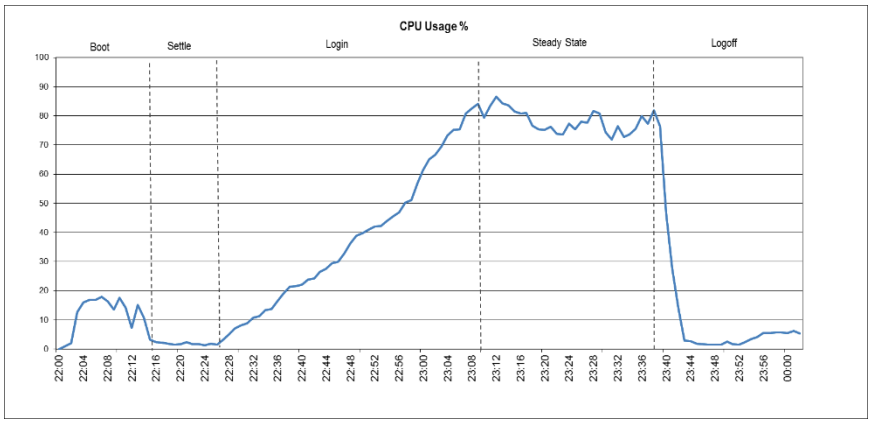

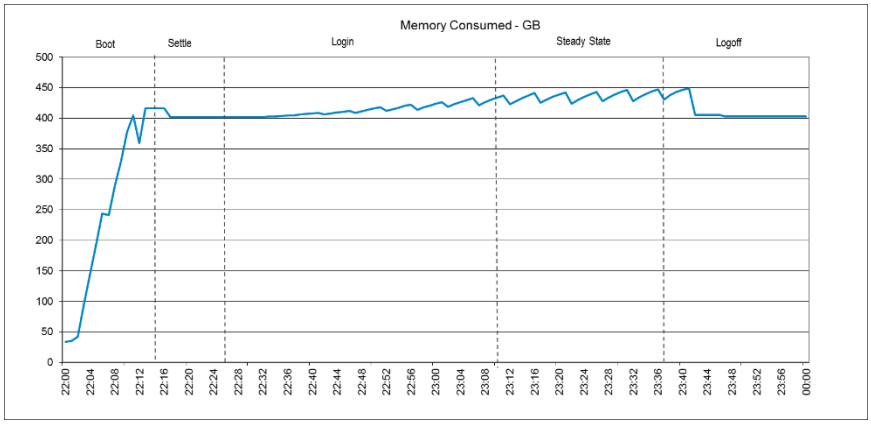
Memory Consumed - GB
Boot
Settle
Login
Steady State
Logoff
500
450
400
350
300
250
200
150
100
50
0
22:00
22:04
22:08
22:12
22:16
22:20
22:24
22:28
22:32
22:36
22:40
22:44
22:48
22:52
22:56
23:00
23:04
23:08
23:12
23:16
23:20
23:24
23:28
23:32
23:36
23:40
23:44
23:48
23:52
23:56
00:00

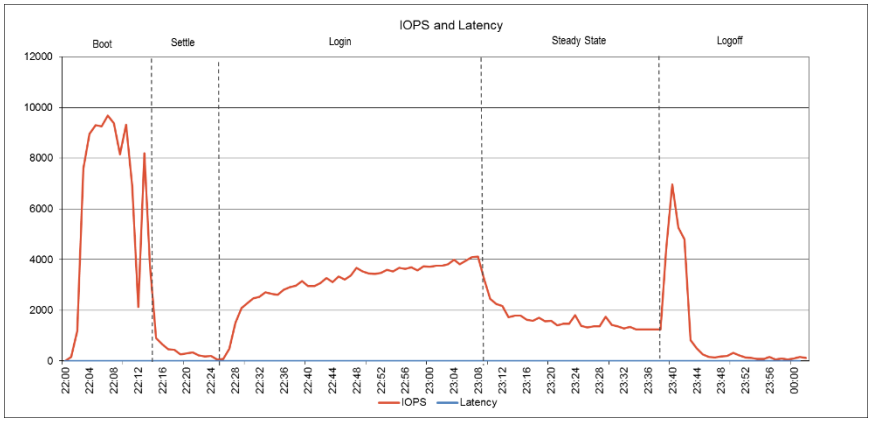
IOPS and Latency
Boot
Settle
Login
Steady State
Logoff
12000
10000
8000
6000
4000
2000
0
22:00
22:04
22:08
22:12
22:16
22:20
22:24
22:28
22:32
22:36
22:40
22:44
22:48
22:52
22:56
23:00
23:04
23:08
23:12
23:16
23:20
23:24
23:28
23:32
23:36
23:40
23:44
23:48
23:52
23:56
00:00
IOPS
Latency

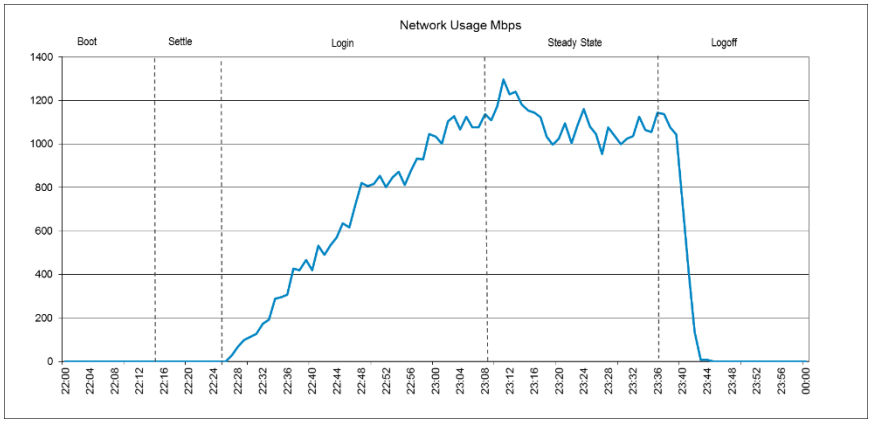
Network Usage Mbps
Boot
Settle
Login
Steady State
Logoff
1400
1200
1000
800
600
400
200
0
22:00
22:04
22:08
22:12
22:16
22:20
22:24
22:28
22:32
22:36
22:40
22:44
22:48
22:52
22:56
23:00
23:04
23:08
23:12
23:16
23:20
23:24
23:28
23:32
23:36
23:40
23:44
23:48
23:52
23:56
00:00

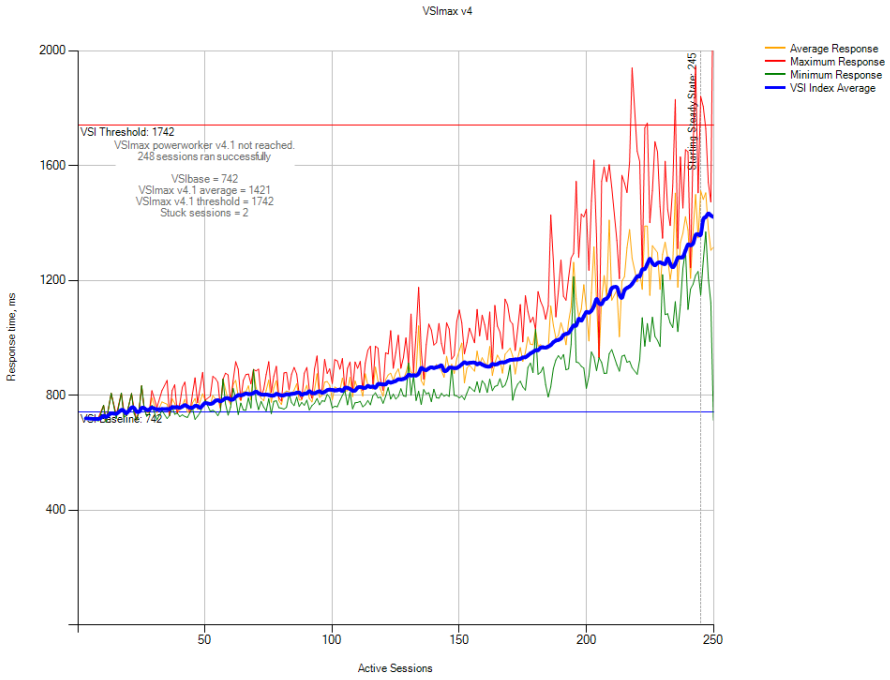
VSImax v4
Average Response
Maximum Response
Minimum Response
VSI Index Average
VSI Threshold: 1742
VSImax powerworker v4.1 not reached.
248 sessions ran successfully
VSIbase = 742
VSImax v4.1 average = 1421
VSImax v4.1 threshold = 1742
Stuck sessions = 2
Starting Steady State: 245
VSI Baseline: 742
Response time, ms
2000
1600
1200
800
400
50
100
150
200
250
Active Sessions

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	File Server VM(s) 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	4 x XtremIO X-Bricks	4 x Unity 300		

Tier 2 storage sizing is based upon 5GB allocated per user for profile and user data. Recommendations are based upon IOPS, not capacity. 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

NOTE: Density is based on empirical data and has not been validated with Windows Server 2016 with Hyper-V.

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 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 Nicolas Cuendet of Dell Wyse Product Marketing for his support and assistance with Wyse Datacenter Appliances at Dell and for his contribution to the Dell Wyse Endpoints section of this document.

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.