Dell Wyse DataCenter for Microsoft VDI Reference Architecture

Dell Engineering

April 2016

A Dell Reference Architecture

Revisions

Date	Description
December 2014	Initial release
August 2015	Major overhaul, add DSMS
September 2015	Formatting, architectural additions
April 2016	Document overhaul with new Broadwell processors and other adjustments to hardware configuration.

Table of contents

1 Introduction

1.1 Purpose

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

1.2 Scope

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

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

1.3 What's New

- Support for Intel Broadwell CPUs (E5-2600v4) and DDR4-2400 RAM
- Support for new Dell Networking options
- Introduce the Dell Appliance for Wyse as base platform for Hyper-V solutions
- Support for all-flash for Local Tier 1 rack solution

2 Solution Architecture Overview

2.1 Introduction

Dell Wyse Datacenter solutions provide a number of deployment options to meet your desktop virtualization requirements. Our solution is able to provide a compelling desktop experience to a range of employees within your organization from task workers to knowledge workers to power users. The deployment options for Dell Wyse Datacenter include:

- Pooled Virtual Desktops (Non-persistent)
- Full Clone Virtual Desktops (Persistent)
- Shared Sessions (RDSH)

Additionally, our solution includes options for users who require:

- Graphics Acceleration
- Unified Communications

2.2 Physical Architecture Overview

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

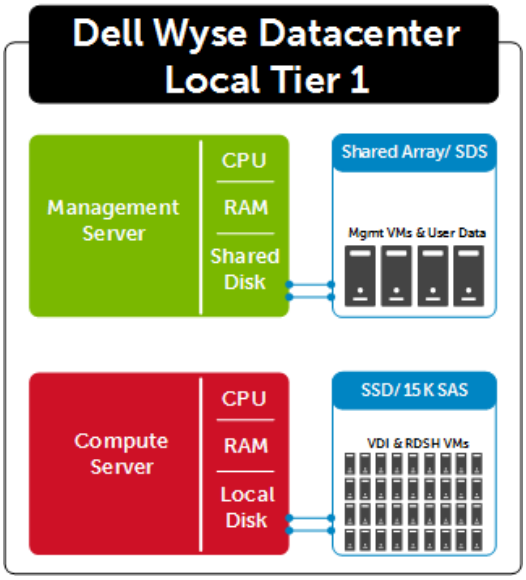

In the Shared Tier 1 solution model, an additional high-performance shared storage array is added to handle the execution of the VDI sessions. All compute and management layer hosts in this model are diskless.

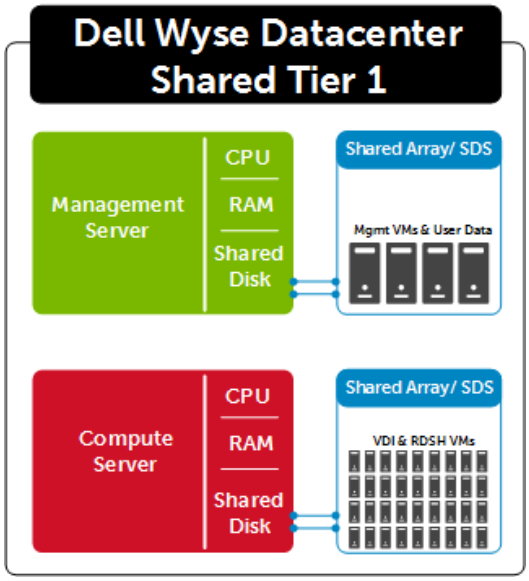

2.3 Solution Layers

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

- Networking Layer
- Compute Server Layer
- Management Server Layer
- Storage Layer
- Thin Client Layer (please refer to the Dell Wyse Thin Clients 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 48-port switch is required to get started in the network layer. This switch hosts all solution traffic consisting of 1Gb iSCSI and LAN sources for smaller stacks. Above 500 users we recommend that 10Gb LAN be used with iSCSI traffic separated into discrete switching Fabrics. Additional switches are added and stacked as required to provide High Availability for the Network layer.

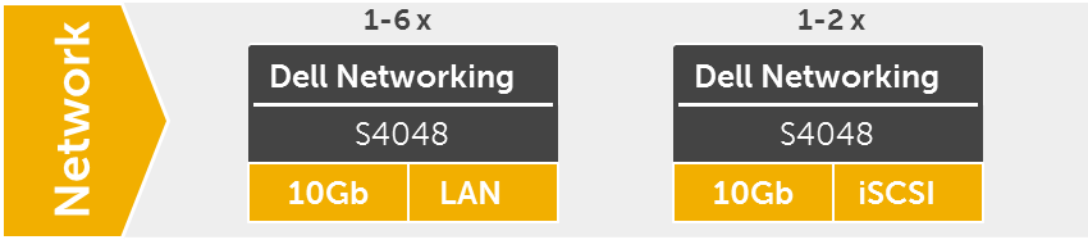

2.3.2 Compute

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

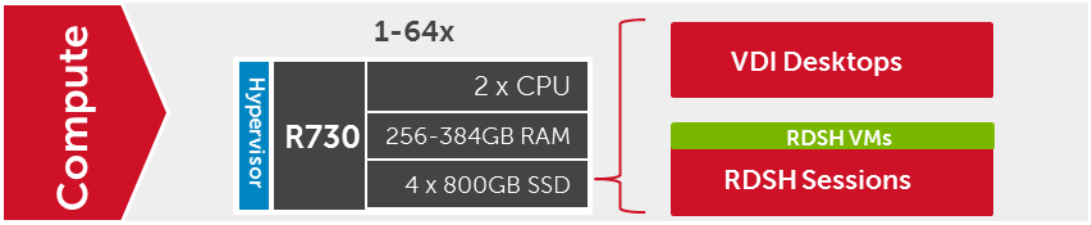

2.3.3 Management

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

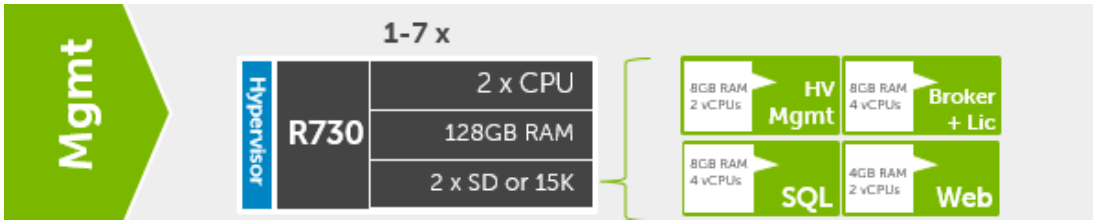

NOTE: Wyse vWorkspace integrates optionally with Microsoft System Center Virtual Machine Manager (SCVMM) for Hyper-V deployments and requires VMware vCenter for vSphere deployments, denoted by "HV Mgmt" in the image above and subsequent design diagrams in the following sections.

2.3.4 Storage

The storage layer consists of options provided by EqualLogic (EQL) for iSCSI or Compellent (CML) arrays for Fiber Channel to suit your Tier 1 (T1) and Tier 2 (T2) scaling and capacity needs. Software-defined storage options such as the Dell Storage with Microsoft Storage Spaces (DSMS) also exist. CML pictured below.

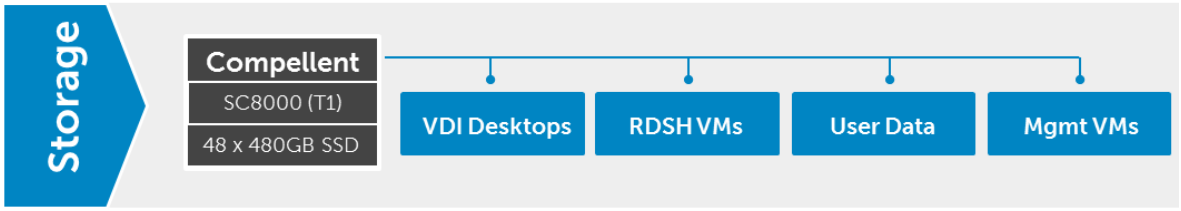

2.4 Local Tier 1

2.4.1 Base Offerings

2.4.1.1 Dell Appliance for Wyse

For pilot or small deployments, the Dell Appliance for Wyse provides an optimized server base with a Quick Start Tool (QST) utility to speed initial deployments. This architecture is non-distributed with all VDI, Management, and storage functions on a single host running **Hyper-V** (vSphere hypervisor not currently supported). If additional scaling is desired, you can grow into a larger distributed architecture seamlessly. Dell QST will help you configure the host, deploy the management infrastructure, and create a pooled desktop or RDSH. For more information on the Dell Appliance for Wyse, please visit: LINK.

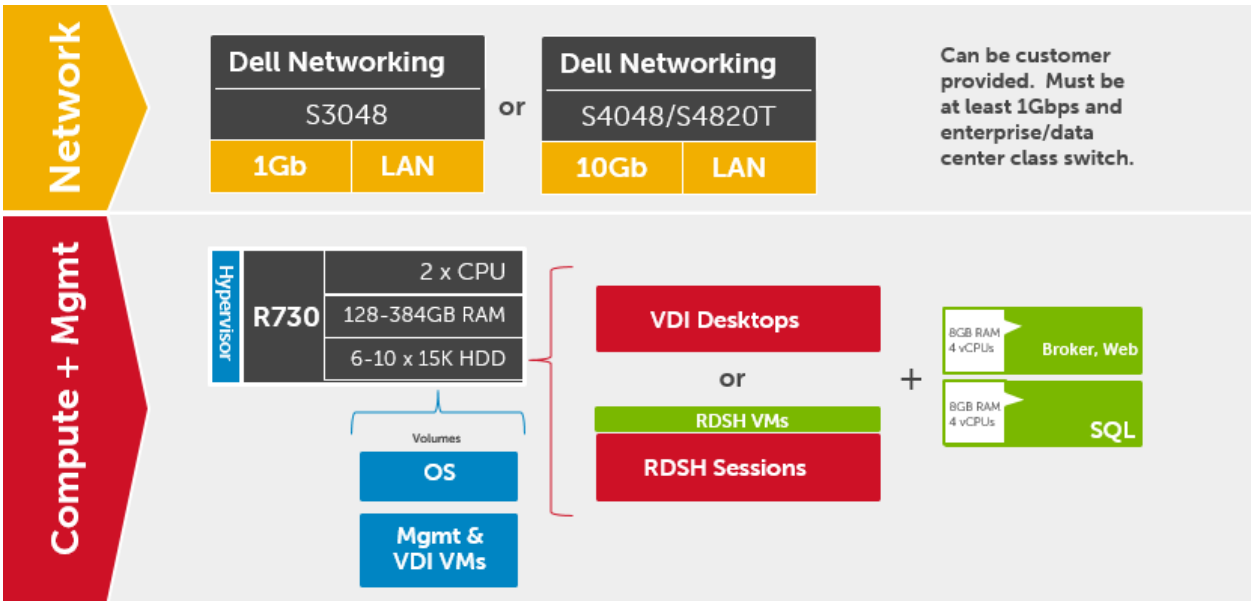

NOTE: The Dell Appliance for Wyse is also available in a tower form factor.

2.4.1.2 Dell Appliance for Wyse (Scale Out)

The Dell Appliance for Wyse can be scaled easily by adding incremental compute nodes for desktop or RDSH VMs and separating the management infrastructure onto its own physical servers protected by external EqualLogic Tier 2 storage. The compute nodes use the large appliance configuration, the mgmt nodes make use of the small appliance configuration.

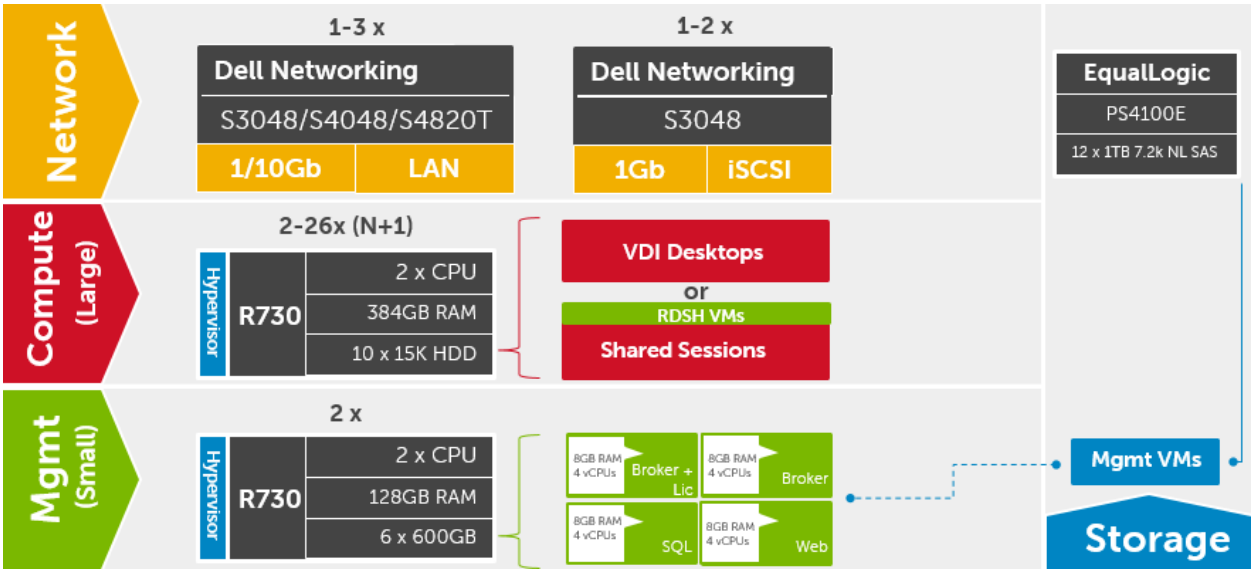

2.4.2 Local Tier 1 for Rack Servers (iSCSI)

The Local Tier 1 solution model provides a high-performance scalable rack-based configuration that hosts user VDI sessions on local SSD or spinning disk in the compute layer. vSphere or Hyper-V based solutions are available and scale based on the chosen hypervisor. All-flash pictured below, if spinning disk is desired substitute the SSDs with 10x 600GB 15K SAS HDDs for vSphere, 12 HDDs for Hyper-V.

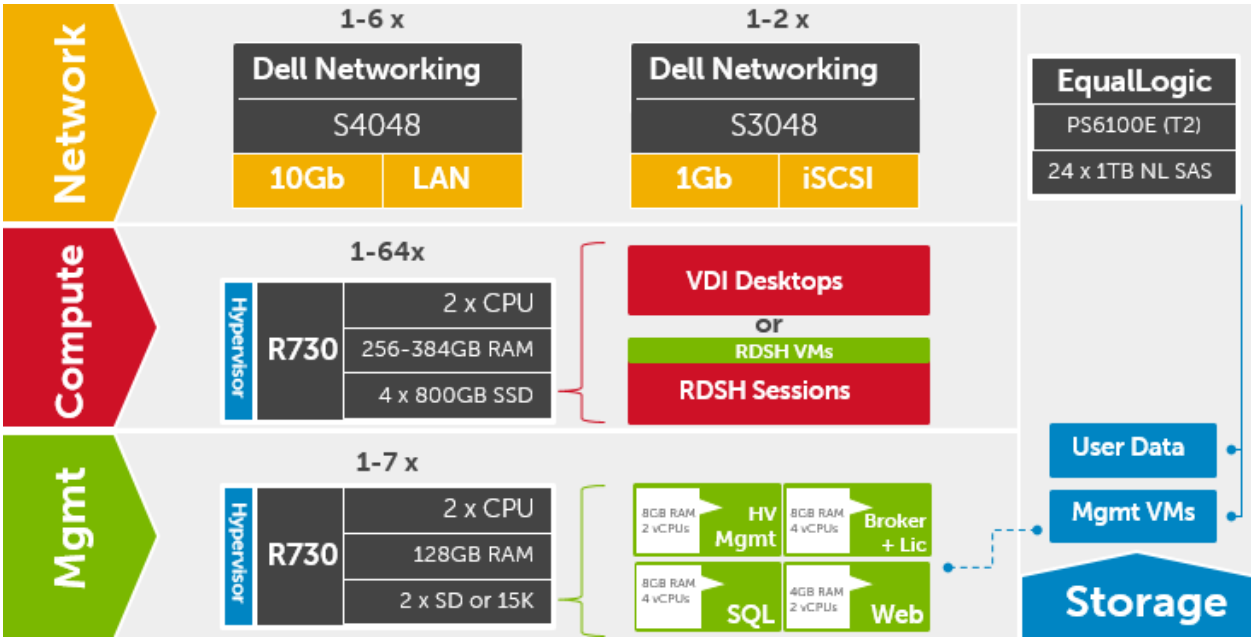

2.4.2.1 Local Tier 1 – Network Architecture

In the Local Tier 1 architecture, a single Dell Networking switch is shared among all network connections for both management and compute, up to 500 users. Over 500 users Dell Wyse Solutions Engineering recommends separating the network Fabrics to isolate iSCSI and LAN traffic as well as making each switch stack redundant. Only the management servers connect to iSCSI storage in this model. All Top of Rack (ToR) traffic is layer 2 (L2)/ switched locally, with all layer 3 (L3)/ routable VLANs trunked from a core or distribution switch. The following diagrams illustrate the logical data flow in relation to the core switch.

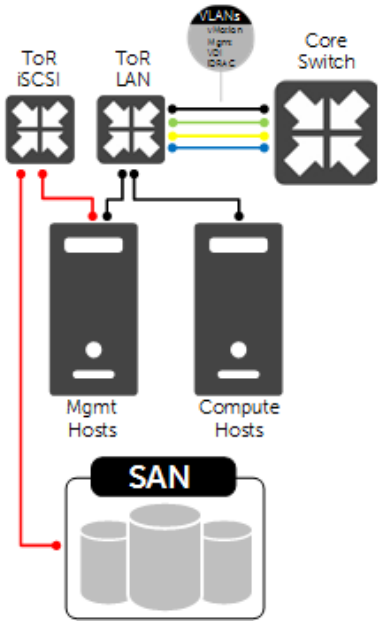

2.4.2.2 Local Tier 1 – Cabling (HA)

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

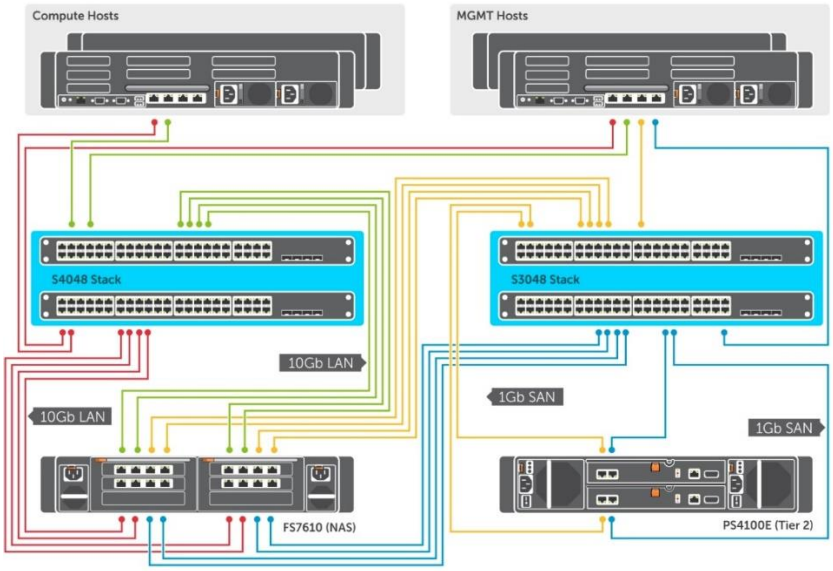

2.4.2.3 Local Tier 1 – Rack Scaling Guidance

Local Tier 1 HW Scaling (iSCSI)				
User Scale	**ToR LAN**	**ToR 1Gb iSCSI**	**EQL T2**	**EQL NAS**
0-1000	S4048	S3048	4100E	
0-1000 (HA)	S4048	S3048	4100E	FS7610
0-6000+	S4048	S3048	6100E	FS7610

2.4.3 Local Tier 1 for Blade Servers

The Local Tier 1 solution model for blade servers provides a high-performance 800GB SSD configuration that does not require shared storage but Tier 2 is added to provide HA to the management layer infrastructure. User VDI sessions are hosted locally on SSDs in each blade server. A pair of PowerEdge M I/O Aggregator switches are required in the A Fabric. The B Fabric, ToR iSCSI switches and Tier 2 storage are optional for HA or additional bandwidth.

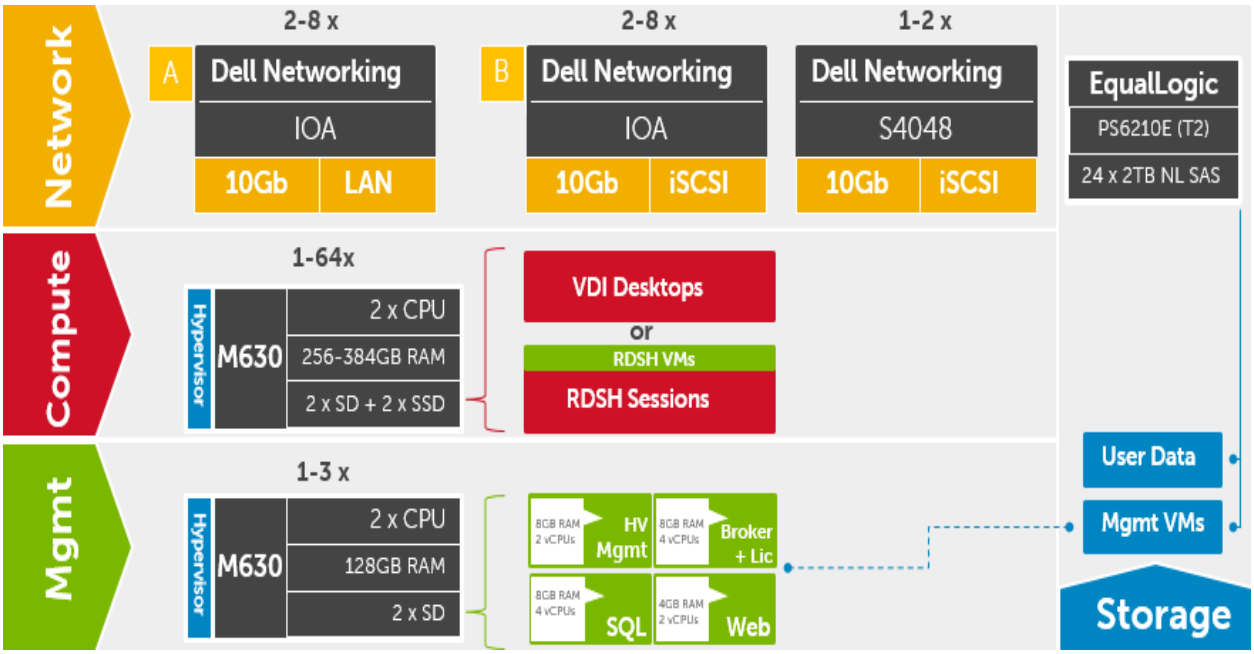

2.4.3.1 Local Tier 1 – Network Architecture

In the Local Tier 1 architecture for blades, there is no need for a ToR switch unless T2 iSCSI is present for HA. The A Fabric IOA switches can connect directly to the core or distribution network layer. Both Management and Compute servers connect to all VLANs in this model via a single vSwitch. The following diagram illustrates the server NIC to chassis switch connections, vSwitch assignments, as well as logical VLAN flow in relation to the core switch.

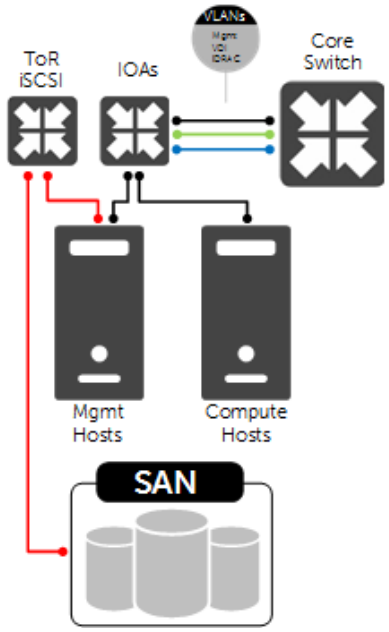

2.4.3.2 Local Tier 1 – Cabling (HA)

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

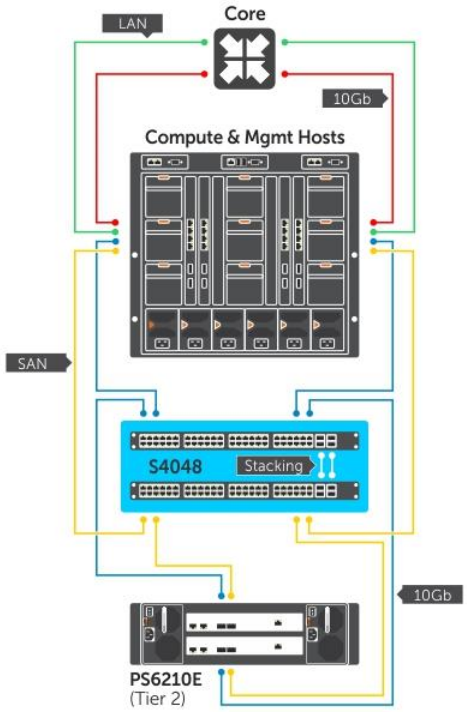

2.4.3.3 Local Tier 1 – Blade Scaling Guidance

Local Tier 1 HW scaling (Blade)				
User Scale	**Blade LAN (A Fabric)**	**Blade iSCSI (B Fabric)**	**ToR 10Gb iSCSI**	**EQL T2**
0-10000 (no HA)	IOA	-	-	-
0-1000 (HA)	IOA	IOA	S4048	4210E
0-6000+ (HA)	IOA	IOA	S4048	6210E

2.5 Shared Tier 1 Rack

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

2.5.1 Shared Tier 1 for Rack Servers (iSCSI)

For POCs or small deployments, Tier1 and Tier2 can be combined on a single EqualLogic PS6210XS storage array. For over 500 users the storage layers are separated into discrete arrays, as depicted in the figure below. Additional arrays are added for Tier 1 or Tier 2 as the user count grows.

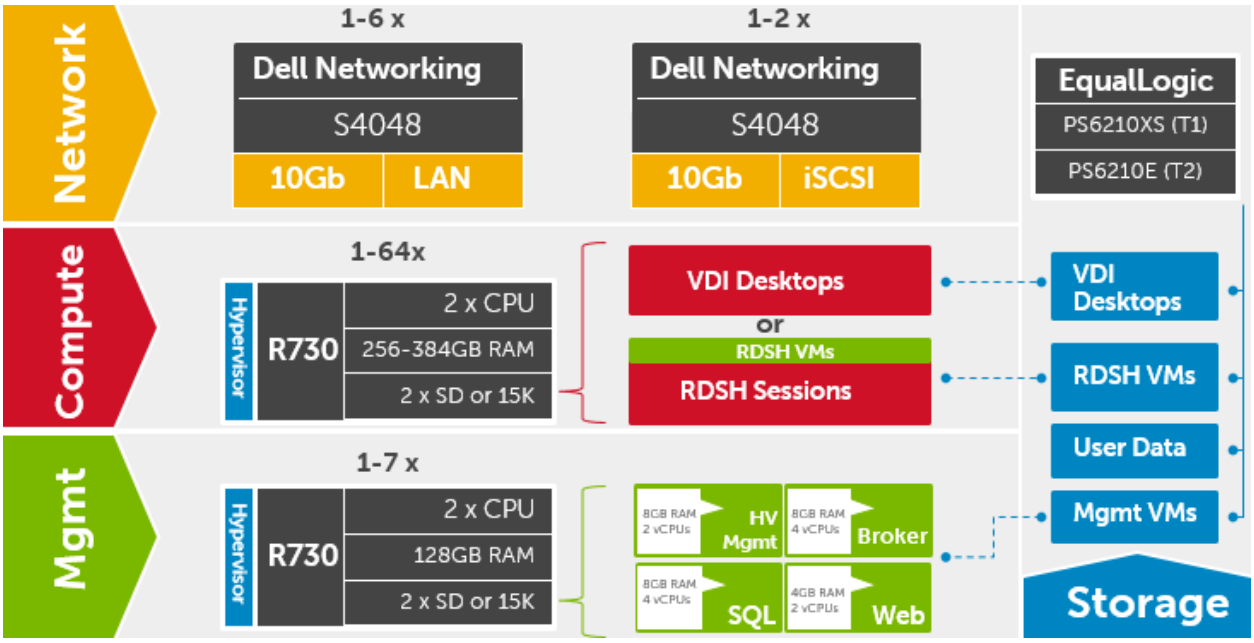

2.5.1.1 Shared Tier 1 – Network Architecture

In the Shared Tier 1 architecture for rack servers, both management and compute servers connect to shared storage. All ToR traffic is layer 2 switched locally, with all layer 3 routable VLANs routed through a core or distribution switch. The following diagrams illustrate the server NIC to ToR switch connections, vSwitch assignments, as well as logical VLAN flow in relation to the core switch.

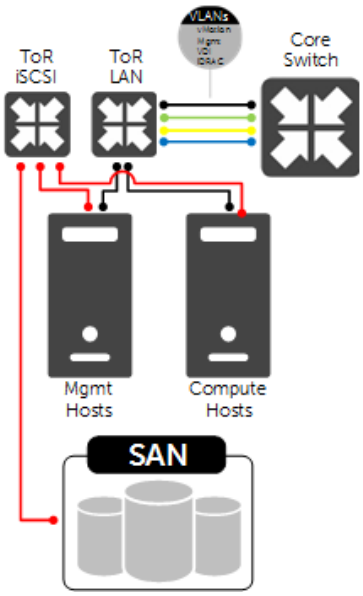

2.5.1.2 Shared Tier 1 – Cabling (HA)

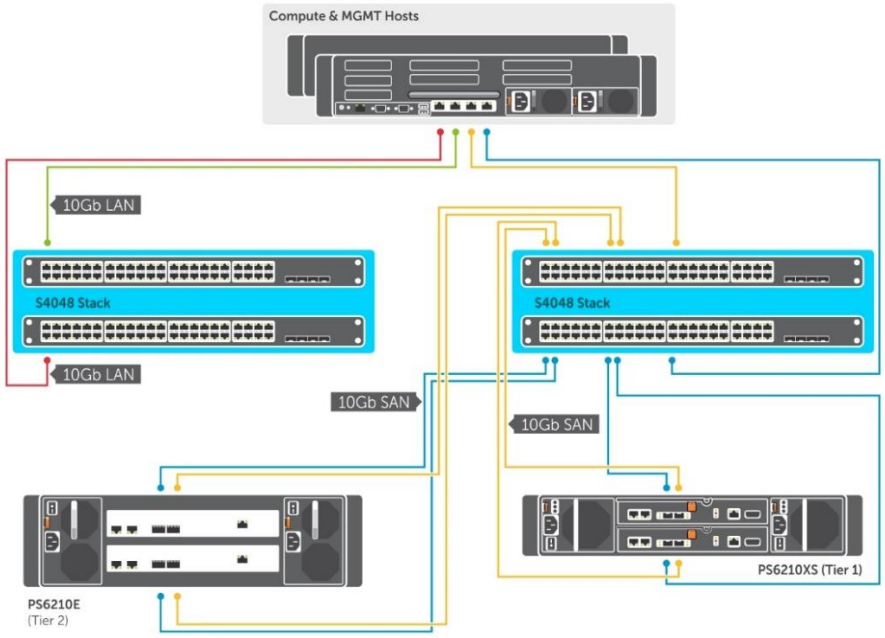

2.5.1.3 Shared Tier 1 – Rack Scaling Guidance

Shared Tier 1 HW scaling (Rack – iSCSI)					
User Scale	ToR LAN	ToR 10Gb iSCSI	EQL T1	EQL T2	EQL NAS
0-500	S4048	S4048	6210XS	-	-
500-1000				6210E	-
0-6000+ (HA)					FS7610

2.5.2 Shared Tier 1 for Rack Servers (FC)

Utilizing Compellent storage for Shared Tier 1 provides a fiber channel (FC) solution where Tier 1 and Tier 2 can optionally be combined in a single array. Tier 2 functions (user data + Mgmt VMs) are removed from the array if the customer has another Tier 2 solution in place or if a Tier 2 Compellent array is used. Scaling this solution is very linear by predictably adding Compellent arrays for every 2000 basic users, on average. The image below depicts a 3000 user array. Please see the Storage section for more information.

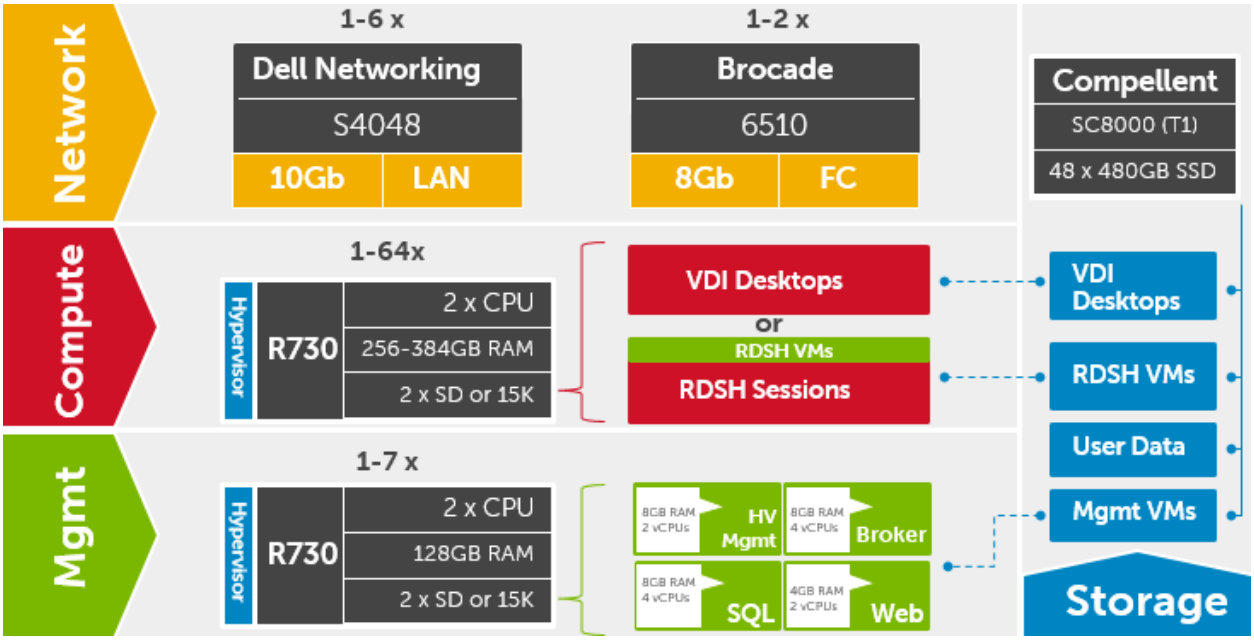

2.5.2.1 Shared Tier 1 – Network Architecture

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

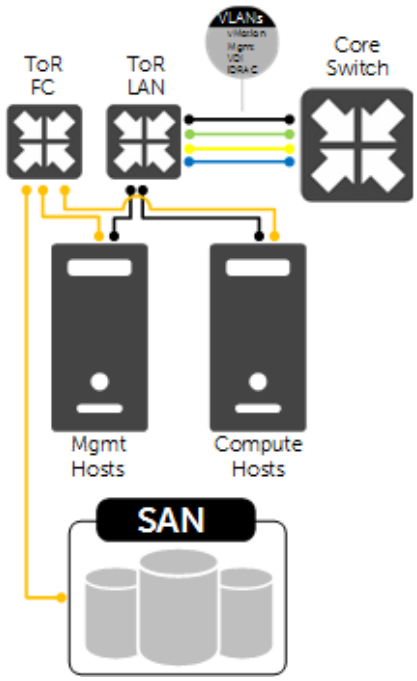

2.5.2.2 Shared Tier 1 – Cabling (HA)

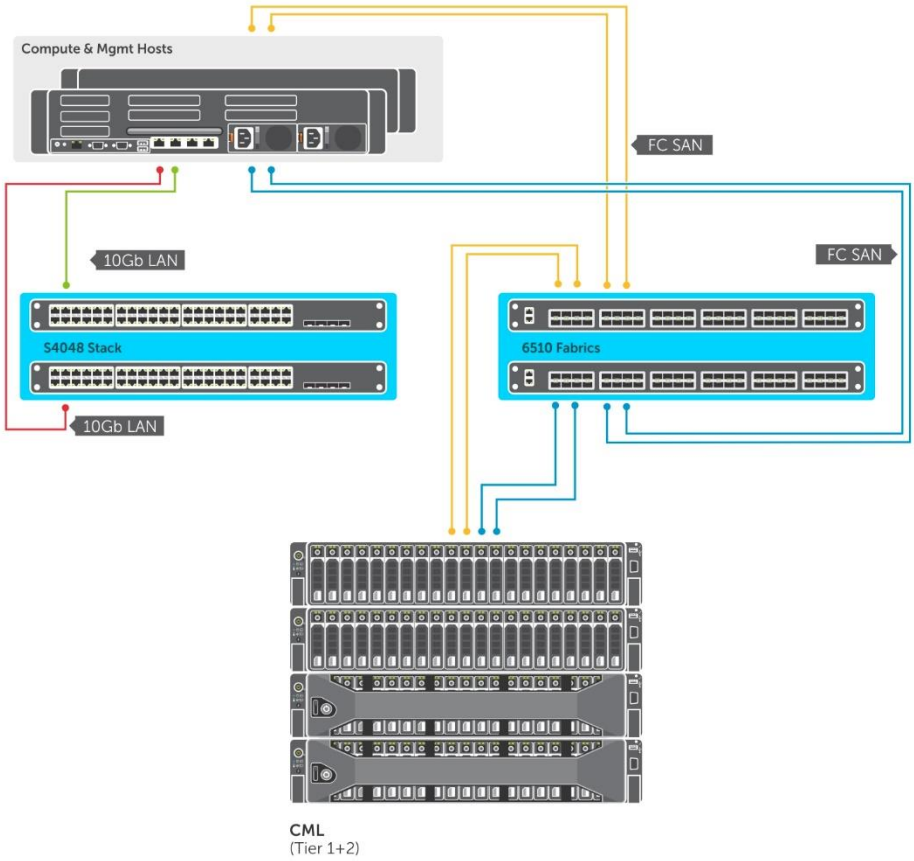

2.5.2.3 Shared Tier 1 – Rack Scaling Guidance

Shared Tier 1 HW scaling (Rack - FC)					
User Scale	**LAN Network**	**FC Network**	**CML T1**	**CML T2**	**CML NAS**
0-1000	S4048	6510	SSD	-	-
0-6000+ (HA)	S4048	6510	SSD	SSD or 15K	FS8600

2.5.3 Shared Tier 1 Rack (Dell Storage with Microsoft Storage Spaces)

As an additional shared tier 1 option for Hyper-V solutions, Microsoft Storage Spaces can be utilized to provide cost effective performance. The Compute hosts communicate with the Scale-Out File Servers via Microsoft's SMB 3.0 protocol over Chelsio RDMA network adapters. The SOFS server is attached to the Dell PowerVault MD1420 storage array via 12Gbps SAS.

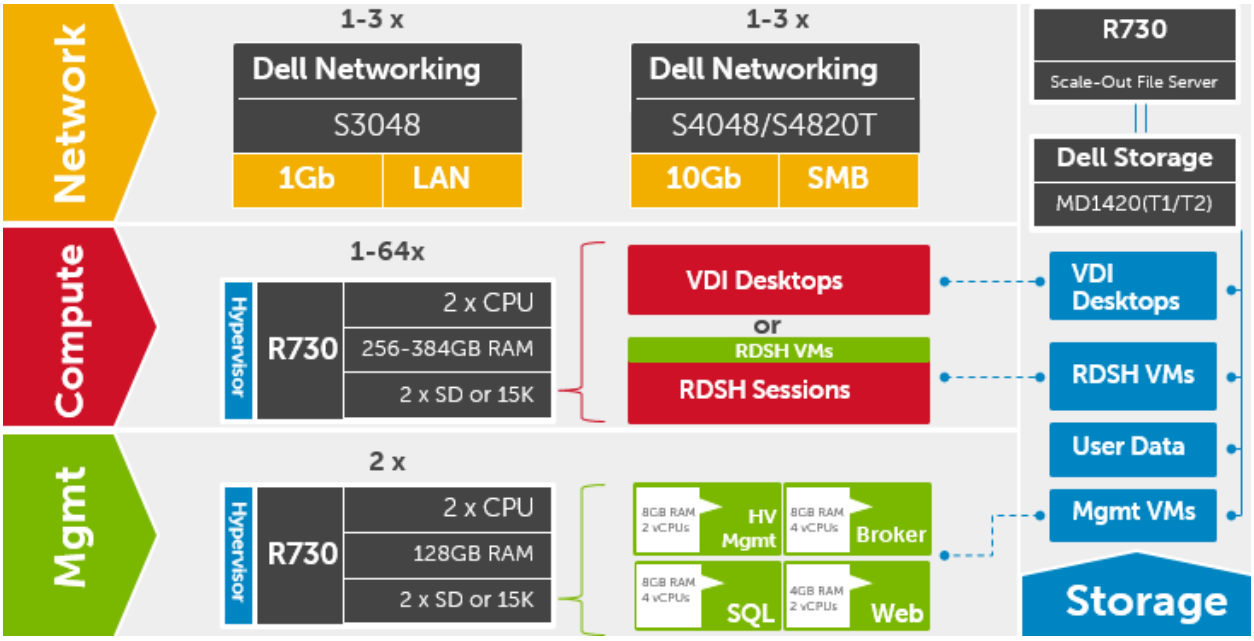

2.5.3.1 Shared Tier 1 Rack (DSMS) – Network Architecture

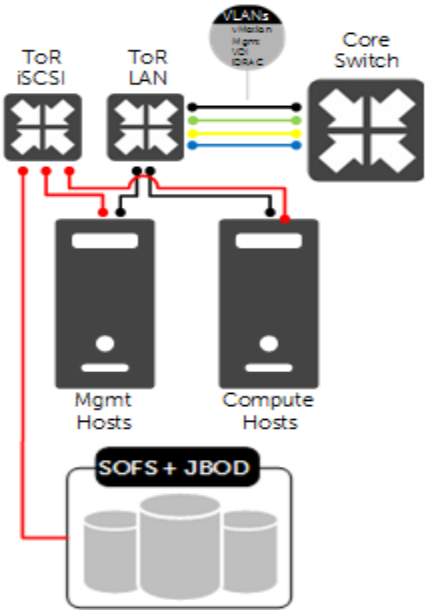

2.5.3.2 Shared Tier 1 Rack (DSMS) – Cabling

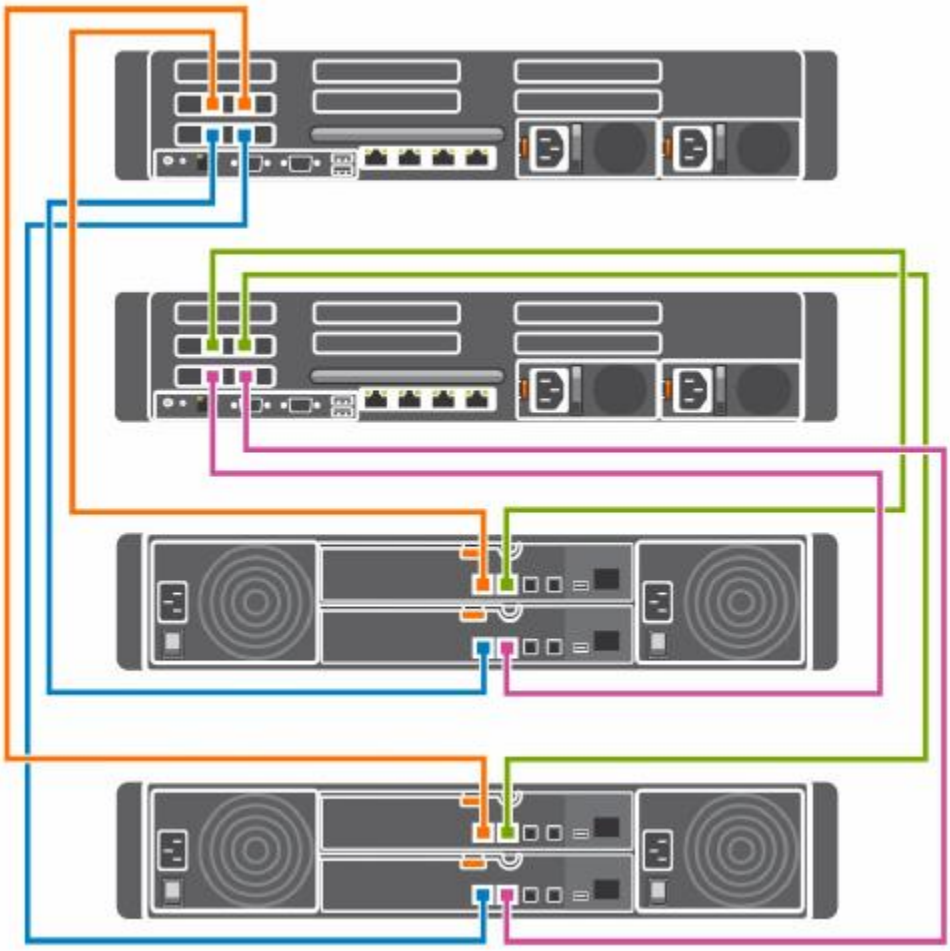

Server 0	Cable color coding		Server 1	Cable color coding	
Server HBA 0-0		Enclosure 0 EMM 0–Port 0	Server HBA 0-0		Enclosure 0 EMM 0–Port 1
Server HBA 0-1		Enclosure 1 EMM 0–Port 0	Server HBA 0-1		Enclosure 1 EMM 0–Port 1
Server HBA 1-0		Enclosure 0 EMM 1–Port 0	Server HBA 1-0		Enclosure 0 EMM 1–Port 1
Server HBA 1-1		Enclosure 1 EMM 1–Port 0	Server HBA 1-1		Enclosure 1 EMM 1–Port 1

2.5.4 Shared Infrastructure (VRTX)

The Shared Infrastructure model provides integrated network switching and internal Direct Attached Storage (DAS) along with up to four blades. The following solutions are based on Shared Infrastructure architecture designs.

Two blades and 15 total disks provide a pilot solution with combined Compute and Management layers for up to 250 pooled virtual desktops or 250 shared sessions. Virtual desktops or shared session VMs will execute on a 10 or 20 disk configuration for Tier 1. The management VMs are segmented on a five disks for Tier 2.

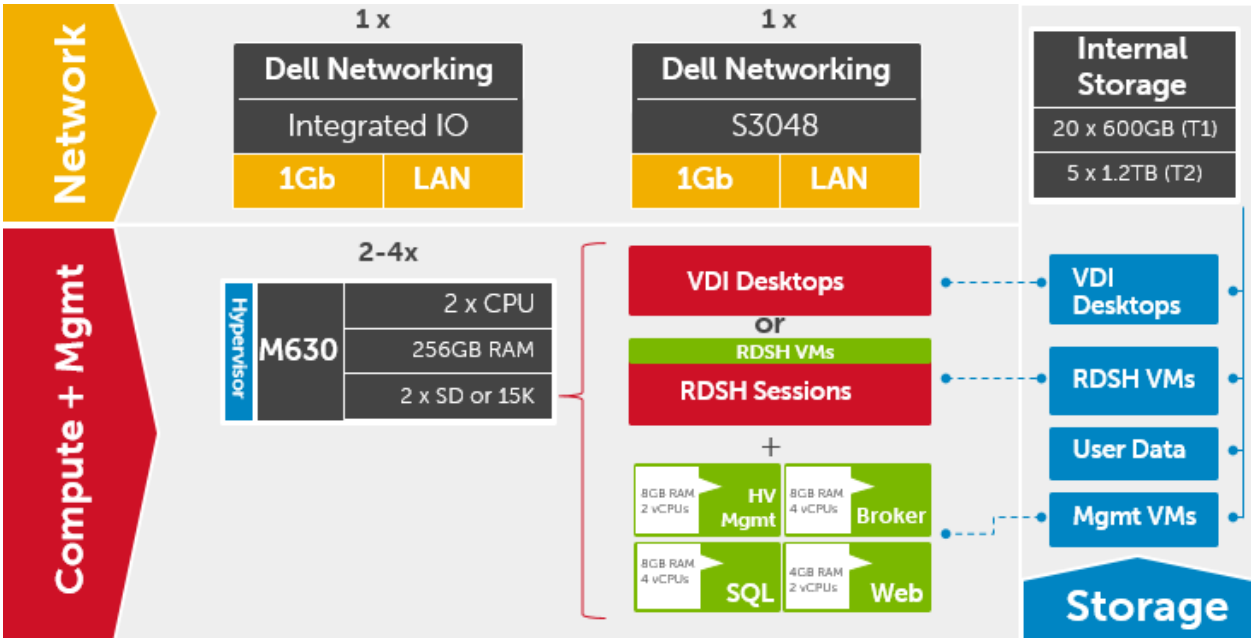

2.5.4.1 Shared Infrastructure – Network Architecture

All ToR traffic connecting to the VRTX integrated switch should be layer 2 switched locally, with all layer 3 routable VLANs trunked from a core or distribution switch. The following diagram illustrates the logical relationship of the VRTX chassis to the integrated switch connections, VLAN assignments, as well as logical VLAN flow in relation to the core switch.

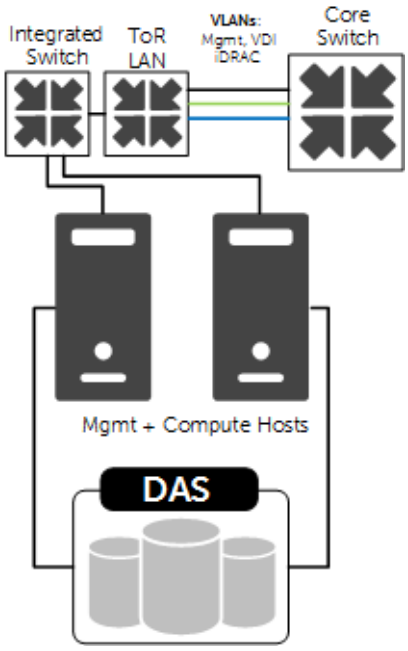

2.5.4.2 Shared Infrastructure – Cabling

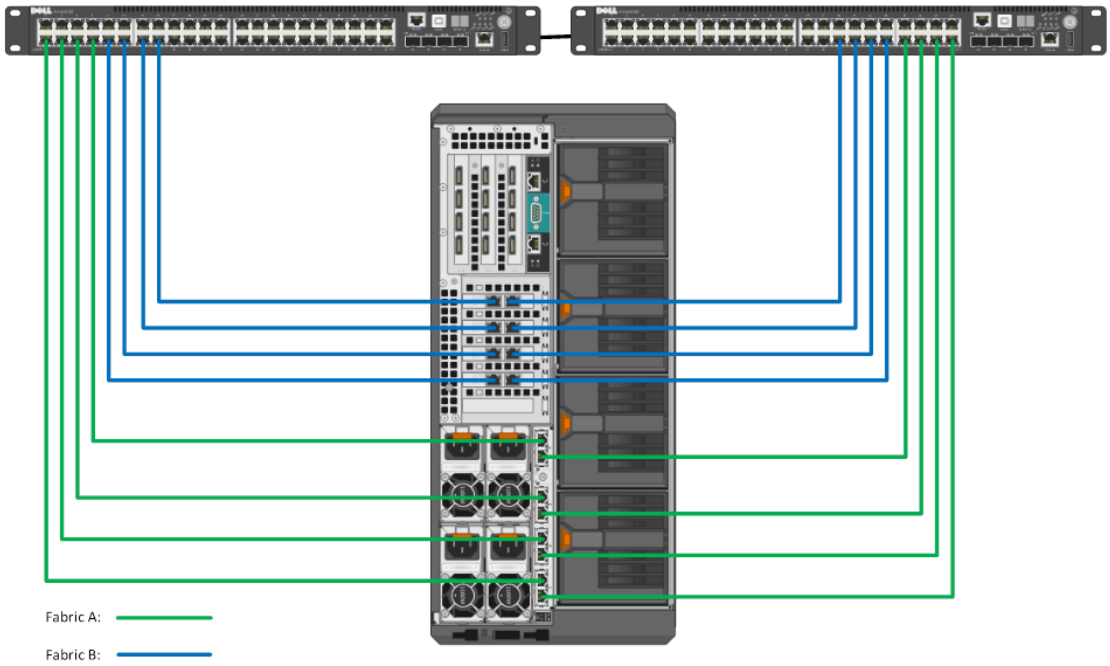

2.6 Graphics Acceleration

A graphics acceleration option can be added to the solution by enabling Microsoft RemoteFX vGPU support and adding up to three physical graphics cards to the Compute host. Based upon solution model and configuration, up to 75 virtual desktop users can be supported by this shared graphics configuration.

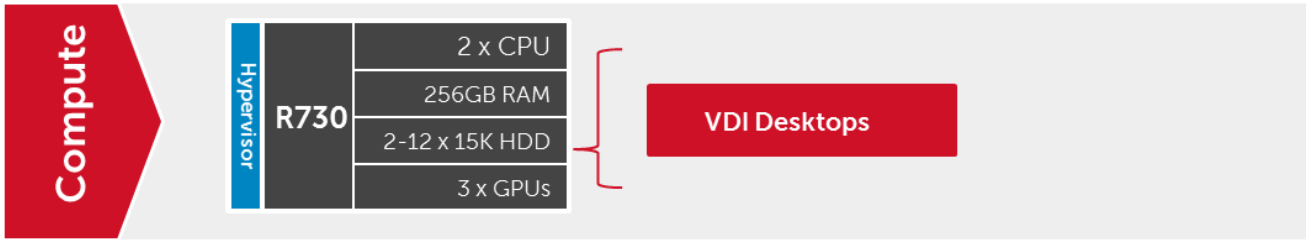

3 Hardware Components

3.1 Network

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

3.1.1 Dell Networking S3048 (1Gb ToR Switch)

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

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

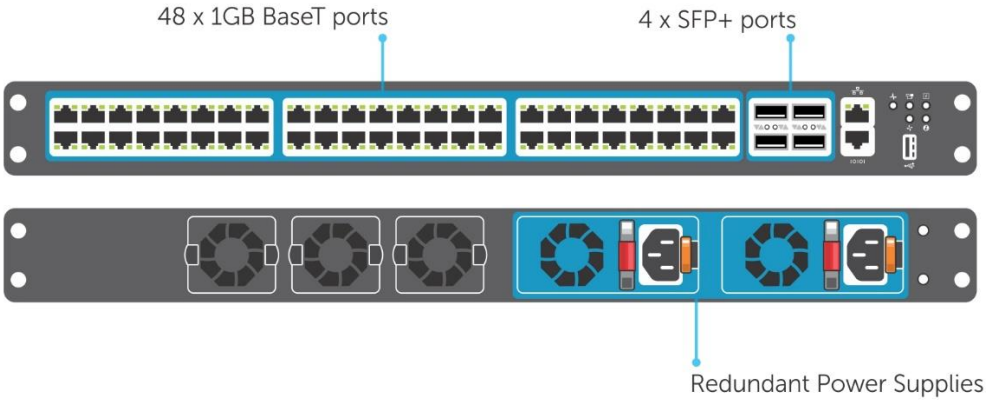

3.1.2 Dell Networking S4048 (10Gb ToR Switch)

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

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

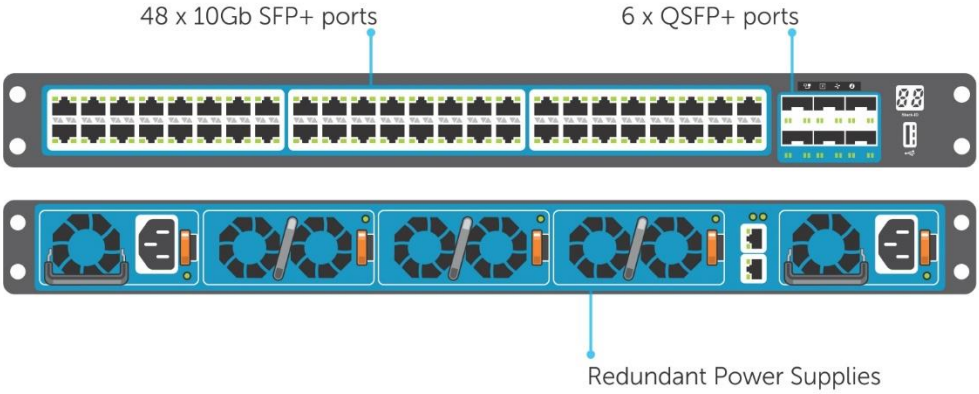

For more information on the S3048, S4048 switches and Dell Networking, please visit: LINK

3.1.3 Brocade 6510 (FC ToR Switch)

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

leading Storage Area Network (SAN) technology while providing "pay-as-you-grow" scalability to meet the needs of an evolving storage environment.

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

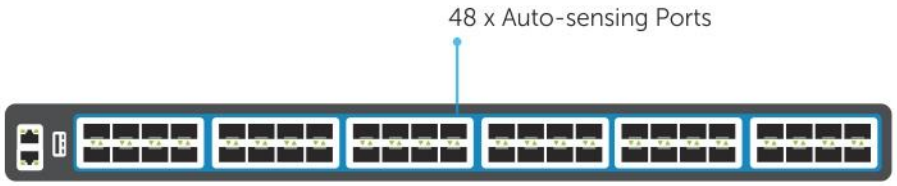

For more information on the Brocade 6510 switch, please visit: LINK

3.1.4 Brocade M5424 (FC Blade Interconnect)

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

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

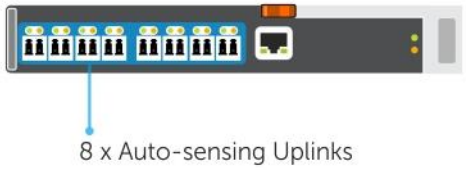

For more information on the Brocade M5424 switch, please visit: LINK

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

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

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

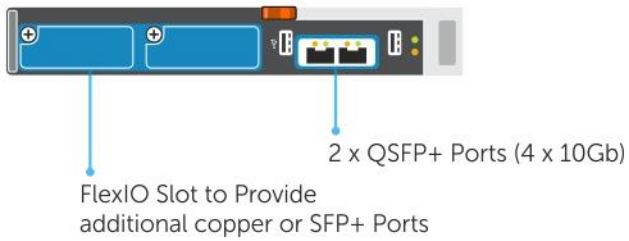

For more information on the Dell IOA switch, please visit: LINK

3.2 Servers

3.2.1 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 up to 22 cores (Broadwell), can host up to 24 DIMMs of DDR4-2400 RAM, supports up to 16 x 2.5" SAS disks and can be outfitted with 2 double-wide GPU accelerators from AMD or NVIDIA. The Dell PowerEdge R730 offers uncompromising performance and scalability in a 2U form factor.

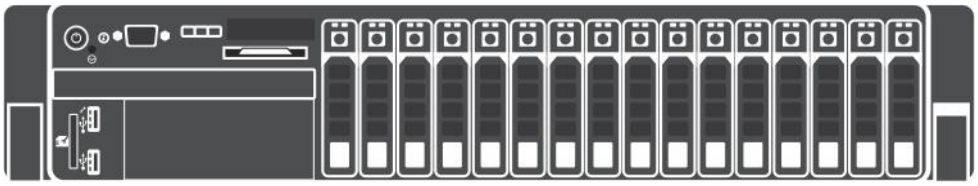

For more information on the R730, please visit: Link

3.2.1.1 Dell Appliance for Wyse

The Dell Appliance for Wyse can be ordered in two optimized configurations, **Large** and **Small** as shown below, or customized as needed. For all-in-one POCs or as a basis for a larger deployment, the **Large** model provides the foundation of the Local Tier 1 model on Hyper-V. The **Small** model can be used as a standalone appliance in a smaller all-in-one deployment, such as a pilot project, or as the dedicated management node in a scale out deployment.

Dell Appliance for Wyse – PowerEdge R730	
Large	**Small**
2 x Intel Xeon E5-2660v4 14C CPU (2Ghz)	2 x Intel Xeon E5-2630v4 10C CPU (2.2Ghz)
384GB Memory (24 x 16GB RDIMMs, (2400MT/s)	128GB Memory (8 x 16GB RDIMMs, 2400MT/s)
Microsoft Hyper-V on 10 x 600GB 15K SAS	**Microsoft Hyper-V** on 6 x 600GB 15K SAS
PERC H730 Integrated RAID Controller – RAID10	PERC H730 Integrated RAID Controller – RAID10
Embedded 2 x 10Gb + 2 x 1Gb NDC	Embedded 2 x 10Gb + 2 x 1Gb NDC
iDRAC8 Enterprise	iDRAC8 Enterprise
2 x 750W PSUs	2 x 750W PSUs

3.2.1.2 Local Tier 1 Rack

In the Local Tier 1 model, VDI desktops or RDSH sessions execute from local storage on each Compute server. The hypervisors used in this solution are vSphere or Hyper-V. In this model, only the Management server hosts access iSCSI storage to support the solution's Management role VMs using the 1Gb NICs on the embedded Network Daughter Card (NDC). 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. All-flash shown for compute below, the SSDs can be optionally substituted for 10-12 15K SAS.

Local Tier 1 Compute Host – PowerEdge R730	
2 x Intel Xeon E5-2698v4 20C CPU (2.2Ghz)	2 x Intel Xeon E5-2698v4 20C CPU (2.2Ghz)
384GB Memory (24 x 16GB RDIMMs, (2400MT/s)	384GB Memory (24 x 16GB RDIMMs, 2400MT/s)
VMware vSphere on internal 8GB Dual SD	**Microsoft Hyper-V on dual 15K SAS HDDs**
4 x 800GB SSD or 10 x 15K SAS	4 x 800GB SSD or 10 x 15K SAS
PERC H730 Integrated RAID Controller – RAID10	PERC H730 Integrated RAID Controller – RAID10
Embedded 2 x 10Gb + 2 x 1Gb NDC	Embedded 2 x 10Gb + 2 x 1Gb NDC
iDRAC8 Enterprise	iDRAC8 Enterprise
2 x 750W PSUs	2 x 750W PSUs

Local Tier 1 Management Host – PowerEdge R730	
2 x Intel Xeon E5-2660v4 14C CPU (2Ghz)	2 x Intel Xeon E5-2660v4 14C CPU (2Ghz)
256GB Memory (16 x 16GB RDIMMs, 2400MT/s)	128GB Memory (8 x 16GB RDIMMs, 2400MT/s)
VMware vSphere on 2 x internal 8GB Dual SD	**Microsoft Hyper-V** on 2 x 600GB 15K SAS
Embedded 2 x 10Gb + 2 x 1Gb NDC	Embedded 2 x 10Gb + 2 x 1Gb NDC
iDRAC8 Enterprise	iDRAC8 Enterprise
2 x 750W PSUs	2 x 750W PSUs

3.2.1.3 Shared Tier 1 Rack (iSCSI)

In the Shared Tier 1 model, VDI desktops or RDSH sessions execute on shared storage so there is no need for local disk on each server to host VMs. To provide server-level network redundancy using the fewest physical NICs possible, both the Compute and Management servers use a QP NDC: 4 x 10Gb ports 2 x 10Gb ports for iSCSI, 2 x 1Gb ports for LAN. All configuration options (beyond the hypervisor differences noted in the table below) are identical except for CPU and RAM which are reduced on the Management host.

Shared Tier 1 Compute Host PowerEdge R730	
2 x Intel Xeon E5-2698v4 20C CPU (2.2Ghz)	2 x Intel Xeon E5-2698v4 20C CPU (2.2Ghz)
384GB Memory (24 x 16GB RDIMMs, 2400MT/s)	384GB Memory (24 x 16GB RDIMMs, 2400MT/s)
VMware vSphere on internal 8GB Dual SD	**Microsoft Hyper-V** on 2 x 300GB 15K SAS
Embedded 4 x 10Gb NDC	Embedded 4 x 10Gb NDC
iDRAC8 Enterprise	iDRAC8 Enterprise
2 x 750W PSUs	2 x 750W PSUs

Shared Tier 1 Management Host PowerEdge R730	
2 x Intel Xeon E5-2660v4 14C CPU (2Ghz)	2 x Intel Xeon E5-2660v4 14C CPU (2Ghz)
256GB Memory (16 x 16GB RDIMMs, 2400MT/s)	128GB Memory (8 x 16GB RDIMMs, 2400MT/s)
VMware vSphere on internal 8GB Dual SD	**Microsoft Hyper-V** on 2 x 300GB 15K SAS
Embedded 4 x 10Gb NDC	Embedded 4 x 10Gb NDC
iDRAC8 Enterprise	iDRAC8 Enterprise
2 x 750W PSUs	2 x 750W PSUs

3.2.1.4 Shared Tier 1 Rack (FC)

Fiber Channel is optionally leveraged as the block storage protocol for Compute and Management hosts with Compellent Tier 1 and Tier 2 storage. Aside from the use of FC HBAs to replace the 10Gb NICs used for iSCSI, the rest of the server configurations are the same.

Shared Tier 1 Compute Host PowerEdge R730	
2 x Intel Xeon E5-2698v4 20C CPU (2.2Ghz)	2 x Intel Xeon E5-2698v4 20C CPU (2.2Ghz)
384GB Memory (24 x 16GB RDIMMs, 2400MT/s)	384GB Memory (24 x 16GB RDIMMs, 2400MT/s)
VMware vSphere on internal 8GB Dual SD	**Microsoft Hyper-V** on 2 x 300GB 15K SAS
Embedded 4 x 10Gb NDC	Embedded 4 x 10Gb NDC
2 x QLogic 2562 8Gb DP FC HBA	2 x QLogic 2562 8Gb DP FC HBA
iDRAC8 Enterprise	iDRAC8 Enterprise
2 x 750W PSUs	2 x 750W PSUs

Shared Tier 1 Management Host PowerEdge R730	
2 x Intel Xeon E5-2660v4 14C CPU (2Ghz)	2 x Intel Xeon E5-2660v4 14C CPU (2Ghz)
256GB Memory (16 x 16GB RDIMMs, 2400MT/s)	128GB Memory (8 x 16GB RDIMMs, 2400MT/s)
VMware vSphere on internal 8GB Dual SD	**Microsoft Hyper-V** on 2 x 300GB 15K SAS
Embedded 4 x 10Gb NDC	Embedded 4 x 10Gb NDC
2 x QLogic 2562 8Gb DP FC HBA	2 x QLogic 2562 8Gb DP FC HBA
iDRAC8 Enterprise	iDRAC8 Enterprise
2 x 750W PSUs	2 x 750W PSUs

3.2.2 Scale-Out File Server for Storage Spaces

The following is the hardware configuration for the Microsoft Scale-Out File Server (SOFS) used for Storage Spaces in conjunction with the Dell Storage MD1420 JBOD:

Microsoft Windows Scale-Out File Server – PowerEdge R730
2 x Intel Xeon E5-2630v3 Processor (2.3Ghz)
128GB Memory (8 x 16GB DIMMs, 2133Mhz)
Microsoft Windows Server 2012 on 2 x 300GB SAS disks
PERC H330P Integrated RAID Controller – RAID1
Intel X520 10Gb DP + I350 1Gb DP NDC
2 x 12Gbps SAS HBA Controllers
iDRAC8 Enterprise, integrated Dell Remote Access Controller
2 x 750W PSUs

3.2.3 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 24 cores, 2 socket Intel Xeon CPUs (Broadwell) and 24 DIMMs (768GB RAM) of DDR4 memory in an extremely compact half-height blade form factor.

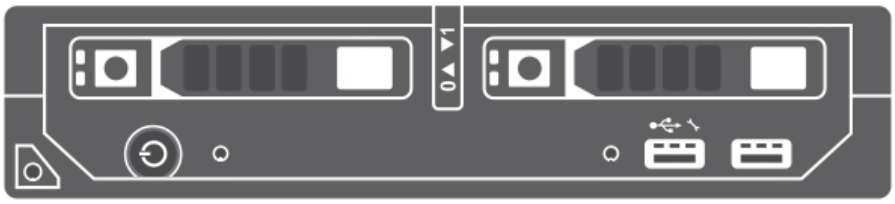

For more information on the PowerEdge M630, please visit: Link

3.2.3.1 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. vSphere is the supported hypervisor in this solution due to its ability to run from integrated SD freeing the pair of SSDs for VDI execution only. In this model, shared storage is not required for Tier 2 unless management host-level HA is required. All management and desktop VMs is hosted locally on their respective blades. The recommended provisioning method is non-persistent desktops. Refer to the Local Tier 1 for Blade Servers section for solution cabling implications.

Local Tier 1 Compute Host – PowerEdge M630
2 x Intel Xeon E5-2698v4 20C CPU (2.2GHz)
384GB Memory (24 x 16GB RDIMMs, 2400MT/s)
VMware vSphere on 2 x 8GB internal SD
QLogic 57810S-k 10Gb DP KR NDC (LAN)
iDRAC8 Enterprise w/ vFlash, 8GB SD

Local Tier 1 Management Host – PowerEdge M630
2 x Intel Xeon E5-2660v4 14C CPU (2GHz)
128GB Memory (8 x 16GB RDIMMs, 2400MT/s)
VMware vSphere on 2 x 8GB internal SD
QLogic 57810S-k 10Gb DP KR NDC (iSCSI-HA)
iDRAC8 Enterprise w/ vFlash, 8GB SD

3.2.3.2 Shared Tier 1 Blade (iSCSI)

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 iSCSI traffic, the B-Fabric dedicated to 10Gb LAN.

Shared Tier 1 Compute Host – PowerEdge M630		
2 x Intel Xeon E5-2698v4 20C CPU (2.2GHz)	Or	2 x Intel Xeon E5-2698v4 20C CPU (2.2GHz)
384GB RAM (24 x 16GB RDIMMs @ 2400MT/s)		384GB RAM (24 x 16GB RDIMMs @ 2400MT/s)
VMware vSphere on 2 x 8GB internal SD		**Microsoft Hyper-V** on 2 x 600GB 15K SAS
QLogic 57810S-k 10Gb DP KR NDC (iSCSI)		QLogic 57810S-k 10Gb DP KR NDC (iSCSI)
QLogic 57810-k 10Gb DP KR NDC (LAN)		QLogic 57810-k 10Gb DP KR NDC (LAN)
iDRAC8 Enterprise w/ vFlash, 8GB SD		iDRAC8 Enterprise w/ vFlash, 8GB SD

Shared Tier 1 Management Host – PowerEdge M630		
2 x Intel Xeon E5-2660v4 14C CPU (2GHz)	Or	2 x Intel Xeon E5-2660v4 14C CPU (2GHz)
256GB RAM (16 x 16GB RDIMMs @ 2400MT/s)		128GB RAM (8 x 16GB RDIMMs @ 2400MT/s)
VMware vSphere on 2 x 8GB internal SD		**Microsoft Hyper-V** on 2 x 300GB 15K SAS
QLogic 57810S-k 10Gb DP KR NDC (iSCSI)		QLogic 57810S-k 10Gb DP KR NDC (iSCSI)
QLogic 57810-k 10Gb DP KR NDC (LAN)		QLogic 57810-k 10Gb DP KR NDC (LAN)
iDRAC8 Enterprise w/ vFlash, 8GB SD		iDRAC8 Enterprise w/ vFlash, 8GB SD

3.2.3.3 Shared Tier 1 Blade (FC)

Fiber Channel is optionally leveraged as the block storage protocol for Compute and Management hosts with Compellent Tier 1 and Tier 2 storage. Aside from the use of FC HBAs to replace the 10Gb NICs used for iSCSI, the rest of the server configurations are the same. Please note that FC is only currently supported using vSphere.

Shared Tier 1 Compute Host – PowerEdge M630		
2 x Intel Xeon E5-2698v4 20C CPU (2.2GHz)	Or	2 x Intel Xeon E5-2698v4 20C CPU (2.2GHz)
384GB RAM (24 x 16GB RDIMMs @ 2400MT/s)		384GB RAM (24 x 16GB RDIMMs @ 2400MT/s)
VMware vSphere on 2 x 2GB internal SD		**Microsoft Hyper-V** on 2 x 600GB 15K SAS
QLogic 57810S-k 10Gb DP KR NDC (LAN)		QLogic 57810S-k 10Gb DP KR NDC (LAN)
1 x QLogic QME2572 8Gb FC mezz (FC)		1 x QLogic QME2572 8Gb FC mezz (FC)
iDRAC8 Enterprise w/ vFlash, 8GB SD		iDRAC8 Enterprise w/ vFlash, 8GB SD

Shared Tier 1 Management Host – PowerEdge M630		
2 x Intel Xeon E5-2660v4 14C CPU (2GHz)	Or	2 x Intel Xeon E5-2660v4 14C CPU (2GHz)
256GB RAM (16 x 16GB RDIMMs @ 2400MT/s)		128GB RAM (8 x 16GB RDIMMs @ 2400MT/s)
VMware vSphere on 2 x 8GB internal SD		**Microsoft Hyper-V** on 2 x 600GB 15K SAS
QLogic 57810S-k 10Gb DP KR NDC (LAN)		QLogic 57810S-k 10Gb DP KR NDC (LAN)
1 x QLogic QME2572 8Gb FC mezz (FC)		1 x QLogic QME2572 8Gb FC mezz (FC)
iDRAC8 Enterprise w/ vFlash, 8GB SD		iDRAC8 Enterprise w/ vFlash, 8GB SD

3.2.4 PowerEdge VRTX

The shared infrastructure platform for the Dell Wyse Datacenter solution is the all in one Dell PowerEdge VRTX. Configurable with up to four PowerEdge blade servers and 25 x 2.5" SAS disks in consolidated 5U form factor, the PowerEdge VRTX provides a flexible performance and capacity for small and midsize businesses as well as remote/branch offices of larger enterprises.

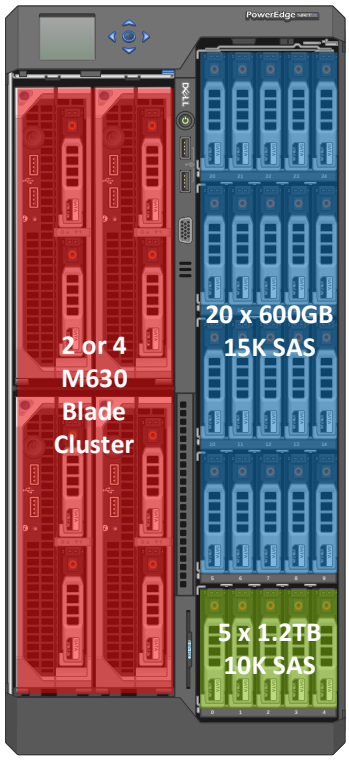

For additional information about the Dell PowerEdge VRTX, please visit: Link

3.2.5 PowerEdge M630 VRTX

The following is the hardware configuration for the M630 nodes used in the Shared Infrastructure model:

Shared Infrastructure – PowerEdge M630 VRTX
2 x Intel Xeon E5-2698v4 20C CPU (2.2GHz)
256GB RAM (16 x 16GB RDIMMs @ 2400MT/s)
Microsoft Hyper-V on 2 x 300GB SAS disks
PERC H710P Integrated RAID Controller – RAID1
QLogic 57810-k 1Gb/ 10Gb DP KR NDC
iDRAC8 Enterprise, integrated Dell Remote Access Controller
2 x 750W PSUs

3.3 Storage

3.3.1 EqualLogic Tier 1 – PS6210XS (iSCSI)

Implement both high-speed, low-latency solid-state disk (SSD) technology and high-capacity HDDs from a single chassis. The PS6210XS 10GbE iSCSI array is a Dell Fluid Data™ solution with a virtualized scale-out architecture that delivers enhanced storage performance and reliability that is easy to manage and scale for future needs. For more information please visit: LINK

Model	Features	Options	Uses
EqualLogic PS6210XS	24 drive hybrid array (SSD + 10K SAS), dual HA controllers, Snaps/Clones, Async replication, SAN HQ, 4 x 10Gb interfaces per controller (2 x SFP + 2 x 10GBT)	13TB – 7 x 400GB SSD + 17 x 600GB 10K SAS	Tier 1 array for Shared Tier 1 solution model (10Gb – iSCSI)
		26TB – 7 x 800GB SSD + 17 x 1.2TB 10K SAS	

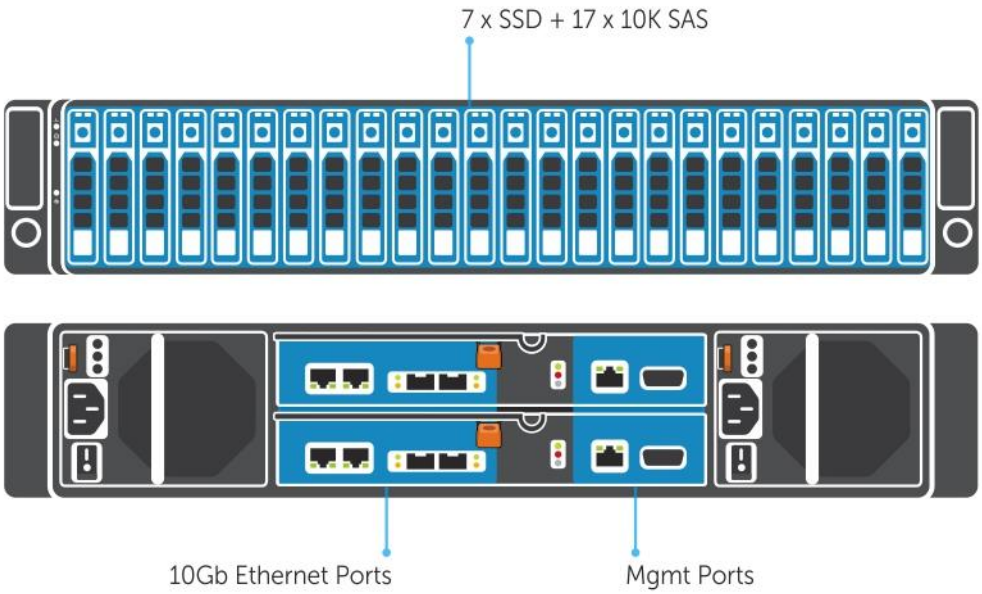

3.3.2 EqualLogic Tier 2 – PS4100E

The Dell EqualLogic PS4100 series enables you to choose the storage solution that best suits your small or midsize business or enterprise branch office. For small T2 deployments the PS4100E is the core of our solution stack.

Model	Features	Options	Uses
EqualLogic PS4100E	12 drive bays (NL-SAS/ 7200k RPM), dual HA controllers, Snaps/Clones, Async replication, SAN HQ, 1Gb	12TB – 12 x 1TB HDs	Tier 2 array for 1000 users or less in Local Tier 1 solution model (10Gb – iSCSI)
		24TB – 12 x 2TB HDs	
		36TB – 12 x 3TB HDs	

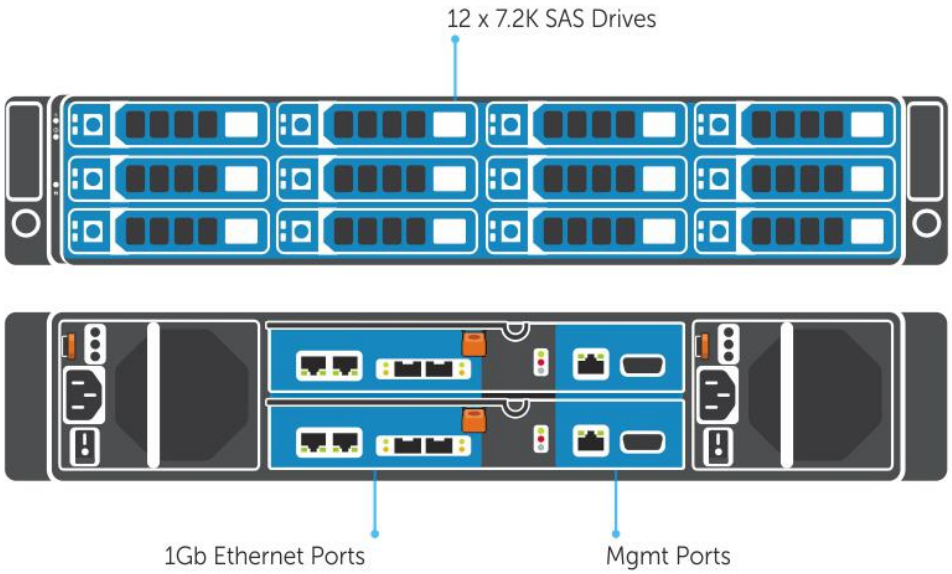

3.3.3 EqualLogic Tier 2 – PS6100E

Fulfill your high-capacity needs with the PS6100E 1Gb iSCSI array, a Dell Storage solution with a virtualized scale-out architecture that delivers enhanced storage performance and reliability that is easy to manage. For larger scale Local Tier 1 solutions, the PS6100E provides the performance and capacity required.

Model	Features	Options	Uses
EqualLogic PS6100E	24 drive bays (NL-SAS/ 7200k RPM), dual HA controllers, Snaps/Clones, Async replication, SAN HQ, 1Gb, 4U chassis	24TB – 24 x 1TB HDs	Tier 2 array for up to 1500 users, per array, in local Tier 1 solution model (1Gb)
		48TB – 24 x 2TB HDs	
		72TB – 24 x 3TB HDs	
		96TB – 24 x 4TB HDs	

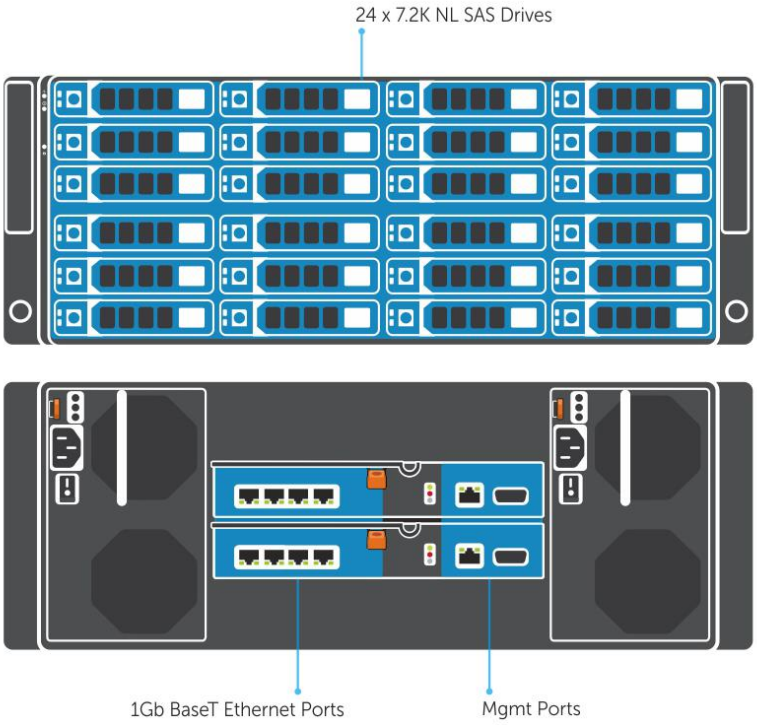

3.3.4 EqualLogic Tier 2 – PS6210E

EqualLogic PS6210 Array Series provide enhanced storage performance for the small-to-medium enterprise with the simplicity of the PS Series product line. Dell EqualLogic PS6210 arrays can drive up to approximately 2GB/sec of throughput per array for sequential, large-block workloads.

Model	Features	Options	Uses
EqualLogic PS6210E	24 drive bays (NL-SAS/ 7.2K RPM), dual HA controllers, Snaps/Clones, Async replication, SAN HQ, 4 x 10Gb ports, 4U chassis	24TB – 24 x 1TB HDs	Tier 2 array for up to 1500 users, per array, in shared Tier 1 solution model (10Gb)
		48TB – 24 x 2TB HDs	
		72TB – 24 x 3TB HDs	
		96TB – 24 x 4TB HDs	

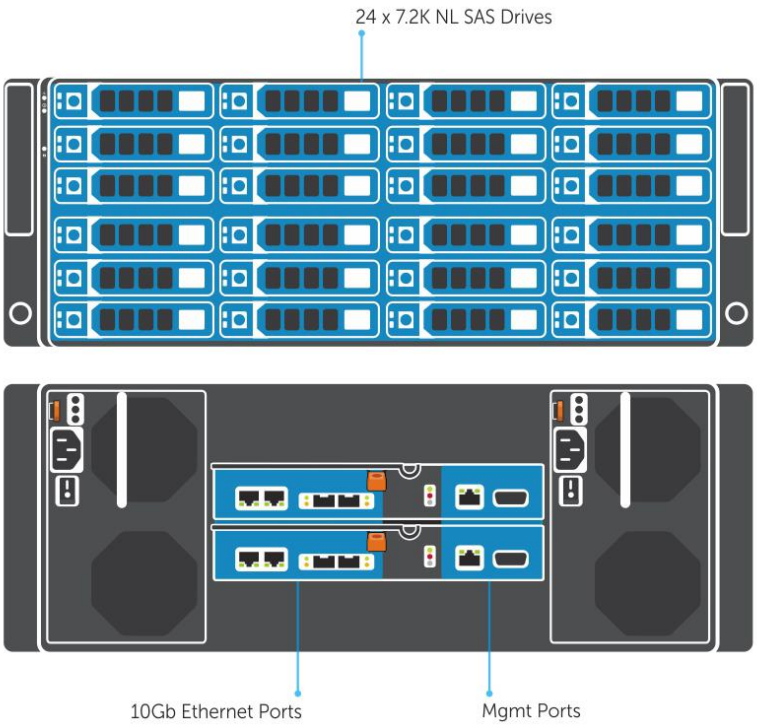

3.3.5 EqualLogic Configuration

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

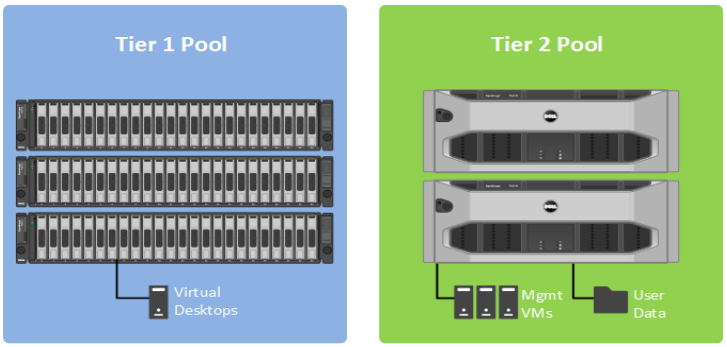

3.3.6 Compellent Tier 1 – SC8000

Experience enhanced storage performance, efficiency and scalability with the next-generation Dell Compellent SC8000 Storage Center Controller. Dell Wyse Solutions Engineering recommends that all Compellent storage arrays be implemented using two controllers in an HA cluster. Fiber Channel is the

preferred storage protocol for use with this array, but Compellent is fully capable of supporting iSCSI as well. Key Storage Center applications used strategically to provide increased performance include:

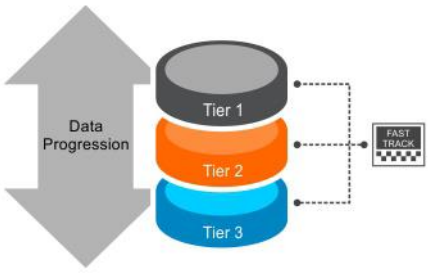

- Simultaneous iSCSI, Fibre Channel (FC) and Fibre Channel over Ethernet (FCoE) front-end interconnects
- Controls SAS and FC drive enclosures, including those with write-intensive and read-intensive SSDs in SLC, MLC and TLC formats
- Scales up to 960 SAS drives in multiple enclosures per dual-controller system and scales-out to many systems across numerous sites, monitored by a single console
- Seamlessly integrates with the FS8600 NAS appliance for fully interoperable block and file storage solutions

Compellent Tier 1 storage consists of a standard dual controller configuration and scales upward by adding disks/ shelves and additional discrete arrays. A single pair of SC8000 controllers will support Tier 1 and Tier 2 for up to 2000 knowledge worker users, as depicted below, utilizing all 15K SAS disks. If Tier 2 is to be separated then an additional 30% of users are added per Tier 1 array. Scaling above this number, additional arrays will need to be implemented. Additional capacity and performance capability is achieved by adding larger disks or shelves, as appropriate, up to the controller's performance limits. Each disk shelf requires 1 hot spare per disk type. RAID is virtualized across all disks in an array (RAID10 or RAID6). Please refer to the test methodology and results for specific workload characteristics in section 7. SSDs are added for use in scenarios where boot storms or provisioning speeds are an issue.

Controller	Front-End IO	Back-End IO	Disk Shelf	Disks	SCOS (min)
2 x SC8000 (16GB)	2 x dual-port 8Gb FC cards (per controller)	2 x quad-port SAS cards (per controller)	2.5" SAS shelf (24 disks each)	2.5" 300GB 15K SAS (~206 IOPS each) or 480GB SSD	6.3

3.3.7 Compellent Tier 2 – SC8000 or SC4020

Compellent Tier 2 storage is optional if a customer wishes to deploy discrete arrays for each Tier. The guidance below is provided for informational purposes. The optional Compellent Tier 2 array consists of a standard dual controller configuration and scales upward by adding disks and shelves. A single pair of SC8000 or SC4020 controllers is able to support Tier 2 for 10,000 basic users. Additional capacity and performance capability is achieved by adding disks and shelves, as appropriate. Each disk shelf requires 1 hot spare per disk type. When designing for Tier 2, capacity requirements will drive higher overall array performance capabilities due to the amount of disk that is on hand. Our base Tier 2 sizing guidance is based on 1 IOPS and 5GB per user.

Controller	Front-End IO	Back-End IO	Disk Shelf	Disks
2 x SC8000 (16GB)	2 x dual-port 8Gb FC cards (per controller)	2 x quad-port SAS cards (per controller)	2.5" SAS shelf (24 disks each)	2.5" 1TB NL SAS (~76 IOPS each)
2 x SC4020 (16GB)	2 x dual-port 8Gb FC cards (per controller)	2 x quad-port SAS cards (per controller)	2.5" SAS shelf (24 disks each)	2.5" 300GB 15K SAS (~206 IOPS each)

Sample Tier 2 Scaling Guidance (based on 1 IOPS per user + 5GB minimum):

Users	Controllers	Disk Shelves	Disks	RAW Capacity
500	2	1	7	7TB
1000	2	1	14	14TB
5000	2	3	66	66TB
10000	2	6	132	132TB

3.3.8 Network Attached Storage (NAS)

3.3.8.1 EqualLogic FS7610

Model	Features	Scaling	Uses
EqualLogic FS7610	Dual active-active controllers, 24GB cache per controller (cache mirroring), SMB & NFS support, AD-integration. Up to 2 FS7610 systems in a NAS cluster (4 controllers). 10Gb iSCSI via 8 x Ethernet ports.	Each controller can support 1500 concurrent users, up to 6000 total in a 2 system NAS cluster.	Scale out NAS for Local Tier 1 to provide file share HA.

3.3.8.2 Compellent FS8600

Model	Features	Options	Uses
Compellent FS8600	Dual active-active controllers, 24GB cache per controller (cache mirroring), SMB & NFS support, AD-integration. Up to 4 FS8600 systems in a NAS cluster (8 controllers). FC only.	Each controller can support 1500 concurrent users, up to 12000 total in a 4 system NAS cluster.	Scale out NAS for Shared Tier 1 on Compellent, to provide file share HA (FC Only).

3.3.9 Dell Storage with Microsoft Storage Spaces Configuration

For pools that span multiple enclosure configurations, ensure that the physical disks are spread evenly across all the enclosures by using multiple pools, rather than assigning a pool to one enclosure.

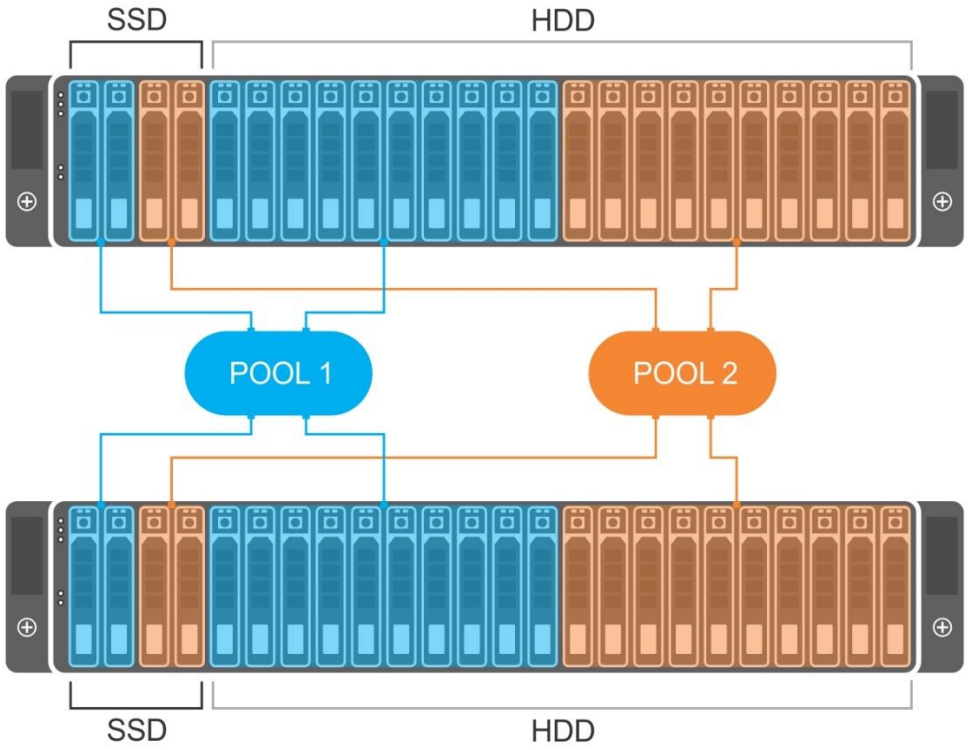

3.4 Dell Wyse Thin Clients

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

3.4.1 Wyse 3020 (ThinOS)

The 3020 sets the standard for thin clients. Providing an exceptional user experience, it features the incredibly fast Dell Wyse ThinOS, for environments in which security is critical—there's no attack surface to put your data at risk. The 3020 delivers outstanding performance based on its dual core system-on-a-chip (SoC) design, and a built-in media CPU 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 3020 is one of the only affordable thin clients to support dual monitors with monitor rotation, enabling increased productivity by providing an extensive view of task work. Designing smooth playback of high bit-rate HD video and graphics in such a small box hasn't been at the expense of energy consumption and heat emissions either. Using less than 7 watts of electricity, the Wyse 3020 small size enables discrete mounting options: under desks, to walls, and behind monitors, creating cool workspaces in every respect.

3.4.2 Wyse 5010 (ThinOS)

Designed for knowledge workers and power users, the Wyse 5010 (ThinOS) is a high-performance thin client based on Dell Wyse ThinOS, the virus-immune firmware base designed for optimal thin client security, performance, and ease-of-use. Highly secure, compact and powerful, the Wyse 5010 combines Dell Wyse ThinOS with a dual-core AMD processor and a revolutionary unified graphics engine for an outstanding user experience. The Wyse 5010 (ThinOS) addresses the performance challenges of processing-intensive applications like computer-aided design, multimedia, HD video and 3D modeling. Scalable enterprise-wide on-premise or cloud-based management provides simple deployment, patching and updates. Take a unit from box to productivity in minutes with auto configuration. Delivering outstanding processing speed and power, security and display performance, the Wyse 5010 (ThinOS) offers a unique combination of performance, efficiency, and affordability. The Wyse 5010 (ThinOS) is Citrix HDX, Microsoft® RemoteFX, and VMware® Horizon View certified. It also supports legacy peripherals via an optional USB adapter. For more information, please visit: Link

3.4.3 Wyse 5010 (Windows Embedded Standard 8)

WE8S

In addition to Dell Wyse ThinOS, the Dell Wyse 5010 thin client is available with Windows Embedded Standard 8 and packs dual-core processing power into a compact form factor for knowledge workers who need performance for demanding virtual Windows® desktops and cloud applications. It's also great for kiosks, and multi-touch displays in a wide variety of environments, including manufacturing, hospitality, retail, and healthcare. It features dual-core processing power and an integrated graphics engine for a fulfilling Windows® 8 user experience. Knowledge workers will enjoy rich content creation and consumption as well as everyday multimedia. Kiosk displays will look great on a thin client that is Microsoft RemoteFX®, Citrix® HDX, VMware PCoIP, and HD video-enabled. Operating with less than 9 watts of energy, the Dell Wyse 5010 (Windows) offers cool, quiet operations, potentially lowering your overall carbon footprint.

3.4.4 Wyse 7010 (Windows Embedded Standard 8)

The versatile Dell Wyse 7010 thin client runs Windows Embedded Standard 8 and gives people the freedom to mix and match a broad range of legacy and cutting edge peripheral devices. Ports for parallel, serial, and USB 3.0 offer fast, flexible connectivity. Like all Dell Wyse thin clients, the new Dell Wyse 7010 (Windows) is one cool operator. Its energy efficient processor – which out-performs other more power-hungry alternatives – and silent fan-less design, all contribute to help lower an organization's carbon footprint through power requirements that are a fraction of traditional desktop PCs.

WE8S

3.4.5 Wyse 5040 AIO

The Dell Wyse 5040 AIO all-in-one (AIO) 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 Cloud Client Manager can help lower your total cost of ownership as you grow from just a few thin clients to tens of thousands.

3.4.6 Dell Venue 11 Pro

Meet the ultimate in productivity, connectivity and collaboration. Enjoy full laptop performance in an ultra-portable tablet that has unmatched flexibility for a business in motion. This dual purpose device works as a tablet when you're out in the field but also enables you to work on your desktop in the office thanks to an optional dock. For more information, please visit: Link

3.4.7 Dell Chromebook 13

The lightweight, easy-to-use Dell Chromebook 13 helps turn education into exploration - without the worries of safety or security. Priced to make 1:1 computing affordable today, Chromebook 13 is backed by Dell support services to make the most of your budget for years to come. The Chrome OS and Chrome browser get students online in an instant and loads web pages in seconds. A high-density battery supported by a 5th Gen Intel® CPU provides up to 12 hours of power. Encourage creativity with the Chromebook 13 and its multimedia features that include a 13.3" screen, stereo sound and webcam. For more information, please visit: Link

4 Software components

4.1 Dell Wyse vWorkspace

Wyse vWorkspace is an enterprise class desktop virtualization management solution which enables blended deployment and support of virtual desktops, shared sessions and virtualized applications. The core components of vWorkspace are:

- **Connection broker**: The vWorkspace Connection Broker helps users connect to their virtual desktops, applications, and other hosted resource sessions. The user's endpoint sends a request to the connection broker to access their virtual environment. The connection broker processes the request by searching for available desktops, and then redirects the user to the available managed desktop or application.
- **Management database**: The vWorkspace Management Database is required to perform administrative functions. The management database stores all the information relevant to a vWorkspace farm, such as configuration data, administrative tasks and results, and information regarding client connections to virtual desktops and RDSH environments.
- **Management console**: The vWorkspace Management Console is an integrated graphical interface that helps you perform various management and administrative functions and can be installed on any workstation or server.
- **Terminal server role**: A set of executables, dynamic link libraries, and device drivers that provides features and management functionality for RDSH deployments in a vWorkspace environment.
- **Data collector service**: The vWorkspace Data Collector service is a Windows service on RDSH servers, virtual desktops, and Hyper-V hosts in a vWorkspace farm that sends a heartbeat signal and other information to the connection broker.
- **Monitoring and Diagnostics**: Built on Dell Software's Foglight platform, vWorkspace Monitoring and Diagnostics provides real-time and historical data for user experience, hypervisor performance, RDSH servers/applications, virtual desktops, Connection Broker servers, Web Access servers, Secure Access servers, profile servers, and farm databases.
- **User profile management**: vWorkspace User Profile Management uses virtual user profiles as an alternative to roaming profiles in a Microsoft Windows environment including virtual desktops and RD Session Hosts. The virtual user profiles eliminate potential profile corruption and accelerate logon and logoff times by combining the use of a mandatory profile with a custom persistence layer designed to preserve user profile settings between sessions.
- **Web access**: vWorkspace Web Access is a web application that acts as a web-based portal to a vWorkspace farm. It helps users to retrieve the list of available applications and desktops by using their web browser. After successful authentication, their published desktops and applications are displayed in the web browser.
- **Secure access**: vWorkspace Secure Access is an SSL gateway that simplifies the deployment of applications over the Internet and can provide proxy connections to vWorkspace components such as RDP sessions, the Web Access client, and connection brokers.

For additional information about Wyse vWorkspace, please visit: LINK

4.1.1 Wyse vWorkspace Deployment Options

Wyse vWorkspace provides a number of delivery options to meet your needs, all within a single, simple, wizard-driven environment that is easy to set up and manage.

- **RD Session Host (RDSH) Sessions**– Provide easy access to a densely shared session environment. vWorkspace RD Session Hosts can deliver full desktops or seamless application sessions from Windows Server Virtual Machines running Windows Server 2003 R2 (32 or 64 Bit), 2008 (32 or 64 bit), 2008 R2, 2012, and 2012 R2. RD Session Host Sessions are well-suited for task based workers using office productivity and line of business applications, without needs for supporting complex peripheral devices or applications with extreme memory or CPU requirements.
- **Computer Groups Types** – Computer Groups can be for virtual or physical computers running Windows XP Pro to Windows 8 Enterprise or Server 2003 R2 to 2012 R2. Additionally there is limited support for Linux computer groups, but Linux is outside of the scope of this reference architecture.
    - **Desktop Cloud** – provides users with access to a single virtual machine from a pool of available virtual machines on one or more non-clustered Hyper-V Servers with local storage. Desktop Clouds are elastic in nature and automatically expand as additional Hyper-V Compute Hosts are added to vWorkspace. New Compute Hosts automatically receive instant copies of the virtual machine templates, from which they provision new virtual machines locally. Desktop Cloud virtual machines are temporarily assigned to a user or device at logon, and at logoff are re-provisioned from the parent VHDX (instant copy of the virtual machine template). Desktop Cloud virtual machines are well suited for task based workers using office productivity and line of business applications.
    - **Temporary Virtual Desktop** – are the non-persistent user desktop VMs traditionally associated with VDI. Each desktop VM is assigned a dedicated portion of the host server's resources to guarantee the performance of each desktop. The desktop VM is dedicated to a single user or device while in use then returned to the computer group at logoff, or rebooted and reset to a pristine gold image state for the next user. Applications can be built into gold images or published via RemoteApp. A Microsoft VDA license is required for each non-Microsoft Software Assurance covered device accessing this type of environment.
    - **Persistent Virtual Desktop Groups** – 1-to-1 desktop VMs assigned to a specific entitled user or device. All changes made by Personal VM users will persist through logoffs and reboots making this a truly personalized computing experience. A Microsoft VDA license is required for each non- Microsoft Software Assurance covered device accessing this type of environment.
    - **Physical Computers** – Like Virtual Desktop Computer Groups, Physical Computers can be persistently or temporarily assigned to users or devices. Common use cases for connections to physical computers are remote software development and remote access to one's office PC.

Please contact your Dell sales rep for more information on licensing requirements for VDI.

4.1.2 Hyper-V Catalyst

Hyper-V Catalyst Components - vWorkspace Hyper-V Catalyst Components increase the scalability and performance of virtual computers on Hyper-V Hosts. Hyper-V catalyst components consist of two components: HyperCache and HyperDeploy.

HyperCache provides read Input/Output Operations per Second (IOPS) savings and improves virtual desktop performance through selective RAM caching of parent VHDs. This is achieved through the following:

- Reads requests to the parent VHD are directed to the parent VHD cache.
- Requests data that is not in cache is obtained from disk and then copied into the parent VHD cache.
- Provides a faster virtual desktop experience as child VMs requesting the same data find it in the parent VHD cache.
- Requests are processed until the parent VHD cache is full. The default size is 800 MB, but can be changed through the Hyper-V virtualization host property.

HyperDeploy manages parent VHD deployment to relevant Hyper-V hosts and enables instant cloning of Hyper-V virtual computers. HyperDeploy uses the following techniques to minimize the time used to deploy a virtual computer.

- Smart copying that only copies to the Hyper-V hosts the parent VHD data that is needed.
- Instant provisioning allows the child VHDs to be cloned while the parent VHD is still being copied to the Hyper-V host.
- Copy status is displayed on the Parent VHDs tab to allow for monitoring of the progress and completion.

HyperDeploy is a core component and requires no configuration.

4.1.3 RDSH Integration

The Remote Desktop Session Host role within Windows Server 2012 R2 enables a physical or virtual server to host RemoteApp programs or session-based desktops. Users can connect to RD Session Host servers in a session collection to run programs, save files, and use resources on those servers. RDSH provides a cost effective method for task or knowledge workers to access full desktop sessions or published applications on a shared host. RDSH is very effectively virtualized allowing several instances to be run in parallel on a single host providing access to hundreds of simultaneous users.

For this solution, the RDSH servers exist as virtualized instances of Windows Server 2012 R2 Standard edition.

4.1.3.1 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

ensures that a CPU can access its own directly-connected RAM banks faster than those banks of the other processor which are accessed via the Quick Path Interconnect (QPI). Ensuring that your virtual RDSH servers do not span physical CPUs will ensure the greatest possible performance benefit.

4.1.3.2 NUMA Alignment

20 physical cores per CPU given the 2698v4 part, 40 logical cores with Hyperthreading 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.

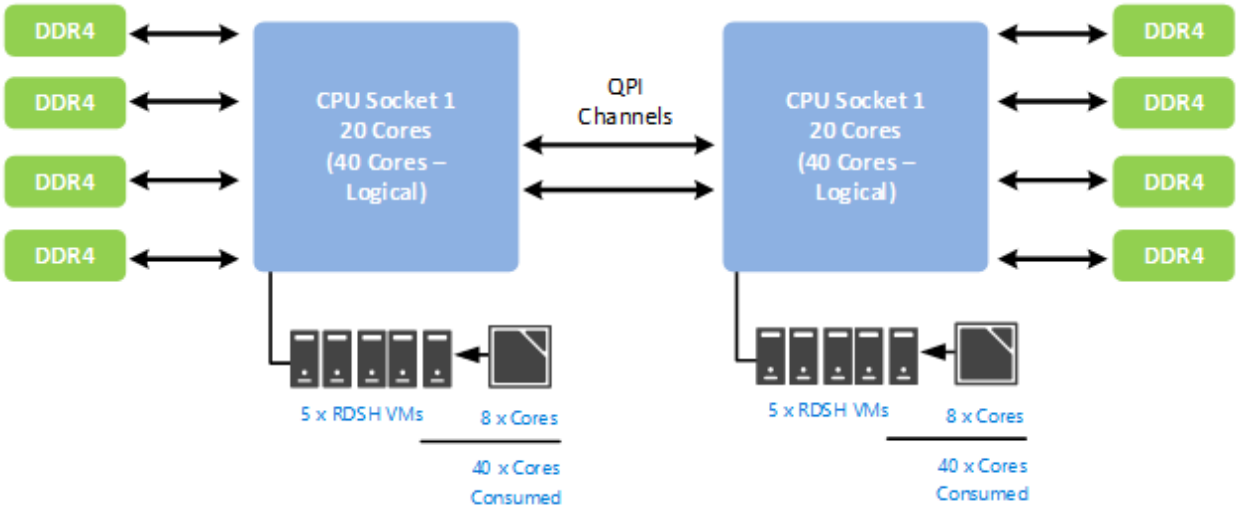

4.2 Hypervisor Platforms

4.2.1 Microsoft Windows Server 2012 R2 Hyper-V

Windows Server 2012 R2 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.

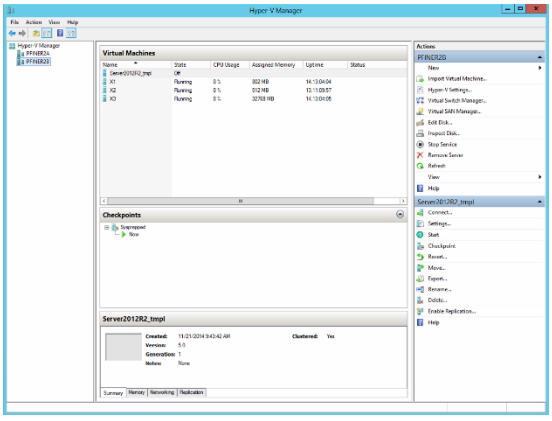

From a network management standpoint, virtual machines are much easier to manage than physical computers. To this end, Hyper-V includes many management features

designed to make managing virtual machines simple and familiar, while enabling easy access to powerful VM-specific management functions.

4.2.2 VMware vSphere 6

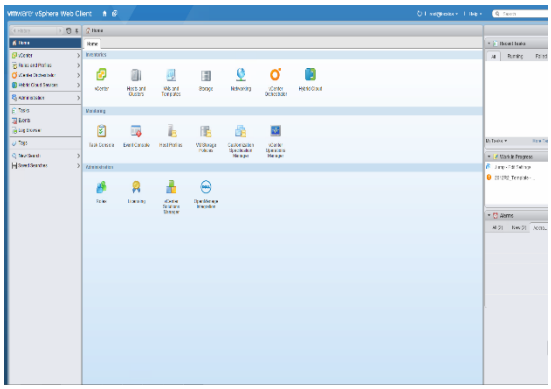

The vSphere hypervisor also known as ESXi is a bare-metal hypervisor that installs directly on top of your physical server and partitions it into multiple virtual machines. Each virtual machine shares the same physical resources as the other virtual machines and they can all run at the same time. Unlike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

VMware vSphere 6 includes three major layers: Virtualization, Management and Interface. The Virtualization layer includes infrastructure and application services. The Management layer is central for configuring, provisioning and managing virtualized environments. The Interface layer includes the vSphere web client.

Throughout the Dell Wyse Datacenter solution, all VMware and Microsoft best practices and prerequisites for core services are adhered to (NTP, DNS, Active Directory, etc.). The vCenter 6 VM used in the solution is a single Windows Server 2012 R2 VM or vCenter virtual appliance, residing on a host in the management layer. SQL server is a core component of the Windows version of vCenter and is hosted on another VM also residing in the management layer. It is recommended that all additional vWorkspace components be installed in a distributed architecture, one role per server VM.

4.3 Microsoft Storage Spaces

Storage Spaces is a storage virtualization feature included in Windows 2012 R2 that uses SAS HBAs and storage enclosures to provide highly available storage for hosted and virtualized deployments. Storage Spaces is based on a disk pooling model. Storage pools can be created flexibly, using affordable hardware, depending on the needs of the deployment. A storage space is carved out of an existing disk pool and presented as a virtual disk in Windows. A storage space integrates well with the rest of the features available in Windows, helping Storage Spaces provide a cost-effective platform with feature sets designed for a variety of scenarios.

Storage tiers combine the best attributes of SSDs and hard disk drives (HDDs) by letting you create virtual disks with two tiers of storage – an SSD tier for frequently accessed data, and a HDD tier for less-frequently accessed data. New data is generally written to the HDD tier and Storage Spaces transparently moves data at a sub-file level between the two tiers based on how frequently data is accessed. As a result, storage tiers can dramatically increase performance for the most used ("hot") data by moving it to SSD storage, without sacrificing the ability to store large quantities of data on inexpensive HDDs.

For additional information about Microsoft Storage Spaces, please visit: LINK

4.3.1 RDMA and SMB Direct

Remote direct memory access or RDMA technology enables removal of data copy operations and reduction in latencies by allowing one computer to directly place information in another computer's memory with minimal demands on memory bus bandwidth and CPU processing overhead, while preserving memory protection semantics. Today, communications over TCP/IP typically require copy operations, which add latency and consume significant CPU and memory resources.

There are a number of implementations of RDMA. Currently DSMS configurations support two RDMA adapters, the Chelsio T520-CR 10GbE adapter that utilizes iWARP and the Mellanox ConnectX-3 10GbE adapter which uses RDMA over Converged Ethernet (RoCE).

SMB Direct is feature already included and configured in Windows Server 2012 R2 that utilizes capabilities of RDMA adapters. With the increased throughput, low latency, and lower CPU utilization provided by RDMA, SMB Direct allows the Scale-Out File Server deployment of this DSMS configuration to resemble similar if not better VDI performance when compared to more traditional storage platforms.

4.4 Application Virtualization

Microsoft Application Virtualization (App-V) provides multiple methods to deliver virtualized applications to RDS environments, virtual desktops, physical desktops, connected as well as disconnected clients. App-V can help reduce the costs and time associated with managing gold master VM and PC images with integrated applications. App-V also removes potential conflicts, such as legacy application compatibility, since virtualized applications are never installed on an end point. Once an application has been packaged using the Microsoft Application Virtualization Sequencer, it can be saved to removable media, streamed to desktop clients or presented to session-based users on a RDSH host. App-V provides a scalable framework that can be managed by System Center Configuration Manager for a complete management solution.

To learn more about application virtualization and how it integrates into a RDS environment please visit: LINK

For more information about vWorkspace and App-V integration, reviews the administration guide: LINK

4.5 Unified Communications

Unified communications is offered as a deployment option to support instant messaging and user conferencing enabled by Microsoft Lync and the VDI Plugin. The Compute host configuration does not change for this deployment option. Pooled and Personal virtual desktops are supported.

For additional information about the Microsoft Lync VDI 2013 plugin, please visit: http://technet.microsoft.com/en-us/library/jj204683.aspx

5 Solution Architecture for Dell Wyse vWorkspace

5.1 Management Role Configuration

The Management role requirements for the base solution are summarized below for each hypervisor. Depending on the scale of the deployment, the number of these services must be adjusted. Use data disks for role-specific application files such as data, logs and IIS web files in the Management volume.

5.1.1 Hyper-V

Role	vCPU	Startup RAM (GB)	Dynamic Memory			NIC	OS + vDisk	
			Min|Max	Buffer	Weight		Size (GB)	Tier 2 Volume (GB)
Broker & Licensing	4	4	1GB|4GB	20%	Med	1	40	-
Foglight	2	4	1GB|4GB	20%	Med	1	60	-
Web access	2	4	1GB|4GB	20%	Med	1	40	-
Secure Gateway	2	4	1GB|4GB	20%	Med	1	40	-
Profiles/Universal Print Server	2	4	1GB|4GB	20%	Med	1	60	-
File Server	1	4	1GB|4GB	20%	Med	1	40	2048 (PTM)
SQL	4	8	2GB|8GB	20%	Med	1	40	200 (VHDX)
SCVMM*	2	8	2GB|8GB	20%	Med	1	40	50 (VHDX)
Total	**19**	**40**	**10GB|40Gb**	**-**	**-**	**8**	**360**	**2298**

***NOTE:** SCVMM is optional with Wyse vWorkspace deployments

5.1.2 vSphere

Role	vCPU	RAM (GB)	NIC	OS + Data vDisk (GB)	Tier 2 Volume (GB)
Broker & Licensing	4	4	1	40 + 5	-
Foglight	2	4	1	40 + 20	-
vCenter	2	8	1	40 + 5	50 (VMDK)
Web access	2	4	1	40 + 5	-
Secure Gateway	2	4	1	40 + 5	-
Profiles/Universal Print Server	2	4	1	40 + 20	-
SQL Server	4	8	1	40 + 5	200 (VMDK)
File Server	1	4	1	40 + 5	2048 (RDM)
TOTALS	**19**	**40**	**8**	**390**	**2298**

5.1.3 SQL Databases

The vWorkspace databases will be hosted by a single dedicated SQL Server 2014 VM in the Management layer. In non-production environments with fewer than 2500 seats SQL Server Express can be used to minimize licensing costs. For production or deployments requiring HA, SQL Server Standard Edition should be used. This architecture provides configuration guidance using a dedicated SQL Server VM to serve the environment. Use caution during database setup to ensure that SQL data, logs, and TempDB are properly separated onto their respective volumes and auto-growth is enabled for each database. Initial placement of all databases into a single SQL instance is fine unless performance becomes an issue, in which case the database needs to be separated into separate named instances. Best practices defined by Dell and Microsoft are to be adhered to, to ensure optimal database performance.

5.1.4 DNS

DNS plays a crucial role in the environment not only as the basis for Active Directory but will be used to control access to the various Dell and Microsoft software components. All hosts, VMs, and consumable software components need to have a presence in DNS, this includes forward and reverse lookups and 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 during the initial deployment. CNAMEs and the round robin DNS mechanism should be employed 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 would be to connect to <CNAME>\<instance name>.

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

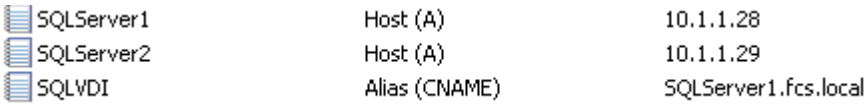

5.2 Storage Architecture Overview

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

5.2.1 Local Tier 1 Storage

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

5.2.2 Shared Tier 1 Storage

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

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

5.2.3 Shared Tier 2 Storage

Tier 2 is shared iSCSI or FC storage used to host the Management server VMs and user data. EqualLogic 4100 or 6100 series 1Gb arrays are used for smaller scale deployments (Local Tier 1 only), the 6210 or 65x0 series for larger deployments (up to 16 in a group), or a single CML array scaled up to 10K users. The 10Gb iSCSI variants are intended for use in Shared Tier 1 solutions. The Compellent Tier 2 array, as specified in section 3.3.2 scales simply by adding disks. The table below outlines the volume requirements for Tier 2. Larger disk sizes are chosen to meet the capacity needs of the customer. The user data are presented either via a file server VM using RDM/ PTD for small scale deployments or via NAS for large scale or HA deployments. The solution as designed presents all SQL disks using VMDK or VHDX formats. RAID 50 are used in smaller deployments but is **not** recommended for critical environments. The recommendation for larger scale and mission critical deployments with higher performance requirements is to use RAID 10 or RAID 6 to maximize performance and recoverability. The following depicts the component volumes required to support a 500 user environment. Additional Management volumes are created as needed along with size adjustments as applicable for user data and profiles.

Volumes	Size (GB)	Storage Array	Purpose	File System
Management	350	Tier 2	vCenter/ SCVMM, XD roles, File & SQL	VMFS/ NTFS
PVS	100	Tier 2	PVS vDisks	VMFS/ NTFS
User Data	2048	Tier 2	File Server/ NAS	RDM/PTD
User Profiles	20	Tier 2	User profiles	VMFS/ NTFS
SQL DATA	100	Tier 2	SQL	VMFS/ NTFS
SQL LOGS	100	Tier 2	SQL	VMFS/ NTFS
TempDB Data	5	Tier 2	SQL	VMFS/ NTFS
TempDB Logs	5	Tier 2	SQL	VMFS/ NTFS
Templates/ISO	200	Tier 2	ISO storage (optional)	VMFS/ NTFS

5.2.4 Storage Networking – EqualLogic iSCSI

Dell's iSCSI technology provides compelling price/performance in a simplified architecture while improving manageability in virtualized environments. Specifically, iSCSI offers virtualized environments simplified deployment, comprehensive storage management and data protection functionality, and seamless VM mobility. Dell iSCSI solutions give customers the "Storage Direct" advantage – the ability to seamlessly integrate virtualization into an overall, optimized storage environment.

If iSCSI is the selected block storage protocol, then the Dell EqualLogic MPIO plugin or MEM (Multipathing Extension Module) is installed on all hosts that connect to iSCSI storage. This module is added via a command line using a Virtual Management Appliance (vMA) from VMware. This plugin allows for easy configuration of iSCSI on each host. The MPIO plugin allows for the creation of new or access to existing data stores and handle IO load balancing. The plugin will also configure the optimal multi-pathing settings for the data stores as well. Some key settings to be used as part of the configuration:

- Specify 2 IP Addresses for iSCSI on each host
- Specify NICs
- Specify Jumbo Frames at 9000 MTU
- Initialize iSCSI initiator
- Specify IP for the EqualLogic Storage group.

5.2.5 Storage Networking – Compellent Fiber Channel (FC)

Based on Fluid Data architecture, the Dell Compellent Storage Center SAN 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.

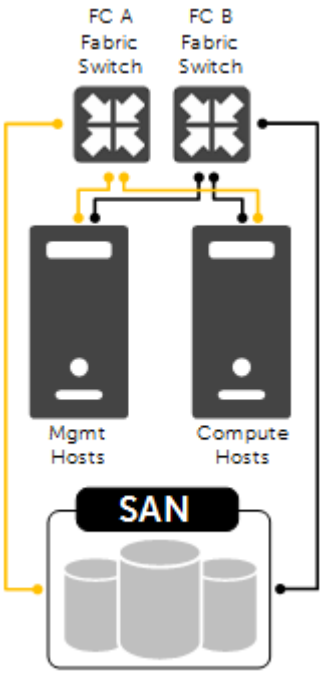

Storage Center actively manages data at a block level using real-time intelligence, providing fully virtualized storage at the disk level. Resources are pooled across the entire storage array. All virtual volumes are thin-provisioned. And with sub-LUN tiering, data is automatically moved between Tiers and RAID levels based on actual use.

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

The following QLogic HBA BIOS settings are used:

- Set the "connection options" field to 1 for point to point only
- Set the "login retry count" field to 60 attempts
- Set the "port down retry" count field to 60 attempts

- Set the "link down timeout" field to 30 seconds
- Set the "queue depth" (or "Execution Throttle") field to 255
- This queue depth are set to 255 because the ESXi VMkernel driver module and DSNRO can more conveniently control the queue depth

5.2.5.1 FC Zoning

Zone at least one port from each server HBA to communicate with a single Compellent fault domain. 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 Compellent Virtual Ports to minimize port consumption as well as simplify deployment. Zone each controller's front-end virtual ports, within a fault domain, with at least one ESXi initiator per server.

5.2.6 Shared Infrastructure

The VRTX chassis contains up to 25 x 2.5" SAS disks to be shared amongst all blade servers in the cluster.

Solution Model	Features	Tier 1 Storage (VDI disks)	Tier 2 Storage (mgmt. + user data)
2 Blade	Up to 250 desktops	10 x 600GB 2.5" 15K SAS	5 x 1.2TB 2.5" 10K SAS
4 Blade	Up to 500 desktops	20 x 600GB 2.5" 15K SAS	5 x 1.2TB 2.5" 10K SAS

VRTX solution volume configuration:

Volumes	Size (GB)	RAID	Disk Pool	Purpose	File System
VDI	1024	10	Tier 1	VDI Desktops	VMFS/ NTFS
Management	200	6	Tier 2	Mgmt VMs, File Server	VMFS/ NTFS
User Data	2048	6	Tier 2	File Server	VMFS/ NTFS
User Profiles	20	6	Tier 2	User profiles	VMFS/ NTFS
SQL DATA	100	6	Tier 2	SQL	VMFS/ NTFS
SQL LOGS	100	6	Tier 2	SQL	VMFS/ NTFS
TempDB Data	5	6	Tier 2	SQL	VMFS/ NTFS
TempDB Logs	5	6	Tier 2	SQL	VMFS/ NTFS
Templates/ ISO	200	6	Tier 2	ISO storage (optional)	VMFS/ NTFS

5.3 Virtual Networking

5.3.1 Local Tier 1 – iSCSI

The network configuration in this model will vary between the Compute and Management hosts. The Compute hosts do not need access to iSCSI 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)
    - iSCSI VLAN: Configured for iSCSI traffic – L2 switched only via ToR switch
    - 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. Each Local Tier 1 Compute host will have a quad port NDC which includes both 10Gb and 1Gb interfaces. Configure the LAN traffic from the server to the ToR switch as a LAG.

5.3.1.1 vSphere

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

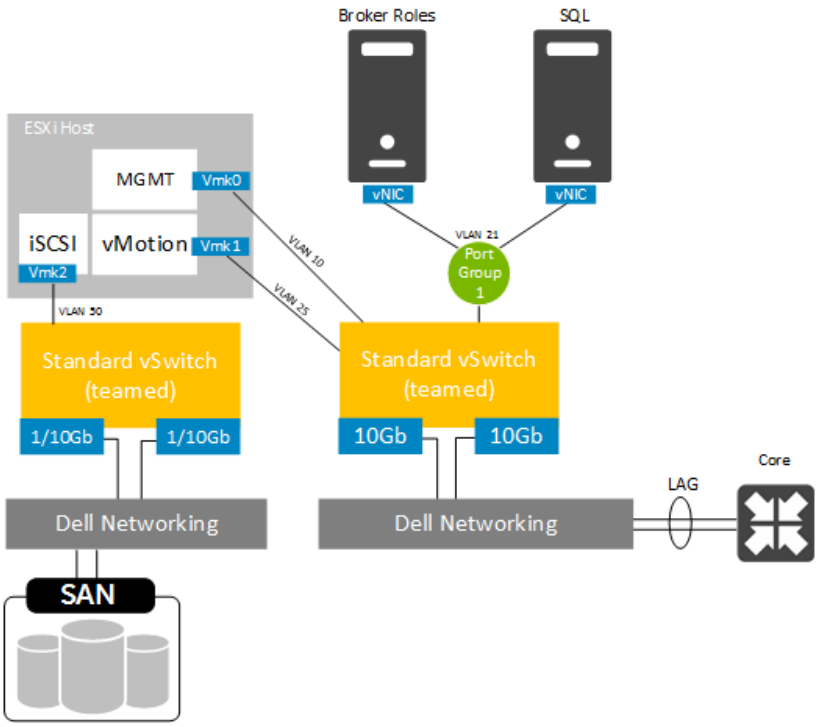

The Compute hosts are configured in the same basic manner, minus the shared storage, with the desktop VMs connecting to the primary port group on the external vSwitch.

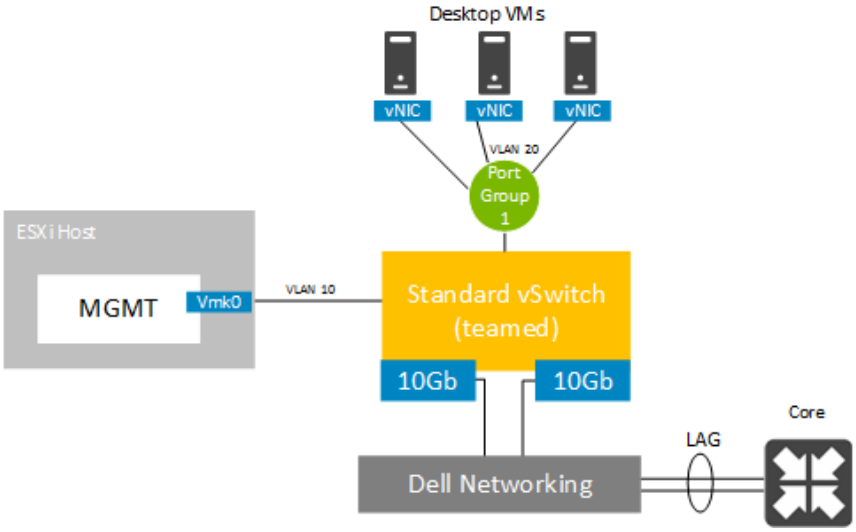

5.3.1.2 Hyper-V

The Hyper-V configuration, while identical in core requirements and hardware, is executed differently due to how Hyper-V and Windows Server 2012 R2 implement networking and virtual switches. As shown in the diagram below, native Windows Server 2012 R2 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.

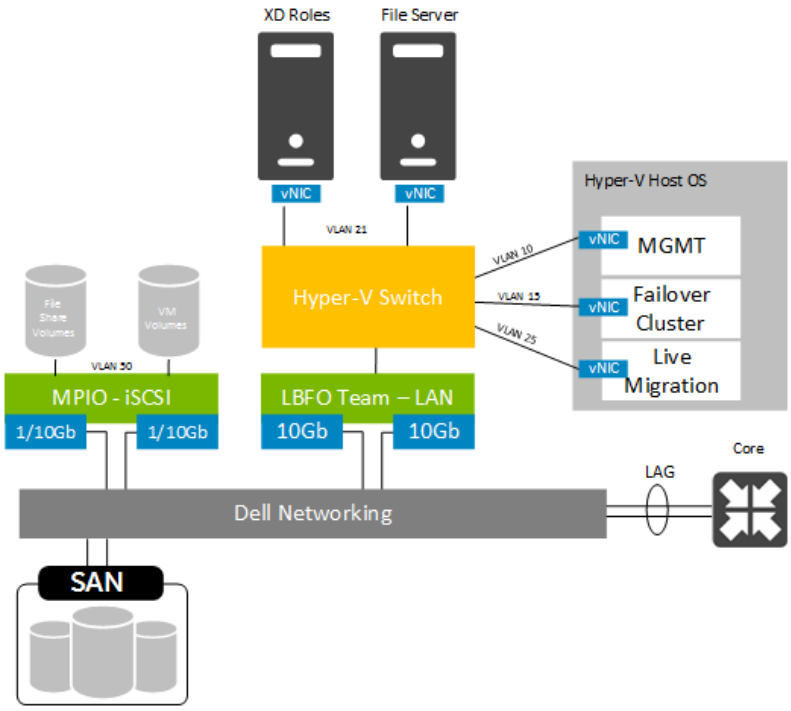

The dedicated compute host configuration is shown in the diagram below and configured similarly to the management host configuration but without the features enabled by shared storage.

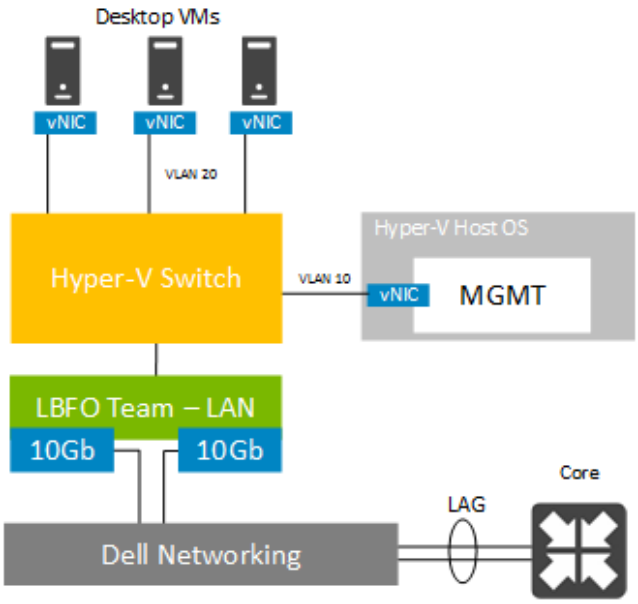

5.3.2 Shared Tier 1 – iSCSI

The network configuration in this model is identical between the Compute and Management hosts since all VMs are hosted on shared storage, whether T1 for desktops or T2 for Mgmt. 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)
  - iSCSI VLAN: Configured for iSCSI traffic – L2 switched only via ToR switch
  - 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)
  - iSCSI VLAN: Configured for iSCSI traffic – L2 switched only via ToR switch
  - VDI Management VLAN: Configured for VDI infrastructure traffic – L3 routed via core switch

- An VLAN for iDRAC is configured for all hardware management traffic – L3 routed via core switch

Following best practices, iSCSI and LAN traffic is physically separated into discrete Fabrics. Each Shared Tier 1 Compute and Management host have a quad port NDC with 4 x 10Gb SFP+ ports each. iSCSI is isolated onto its own vSwitch with redundant ports to guarantee storage bandwidth. Configure the LAN traffic from the server to the ToR switch as a LAG.

5.3.2.1 vSphere

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

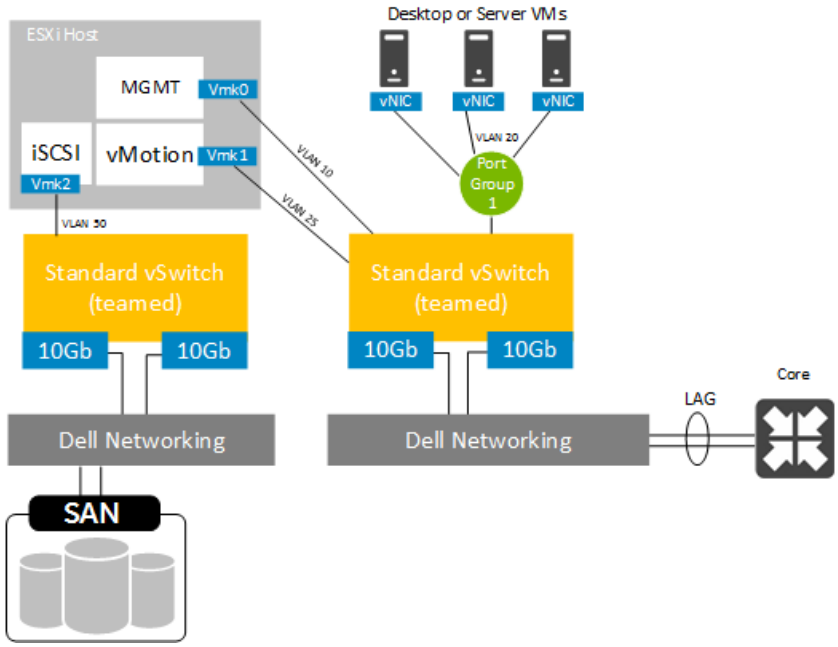

5.3.2.2 Hyper-V

The Hyper-V configuration, while identical in core requirements and hardware, is executed differently due to how Hyper-V and Windows Server 2012 R2 implement networking and virtual switches. As shown in the diagram below, native Windows Server 2012 R2 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.

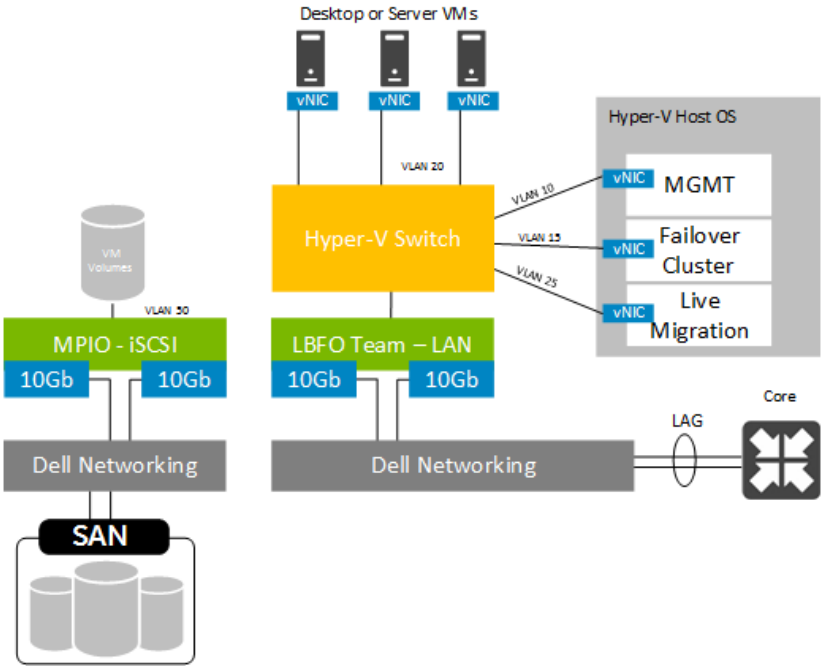

5.3.3 Shared Tier 1 – FC

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

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

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

5.3.3.1 vSphere

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

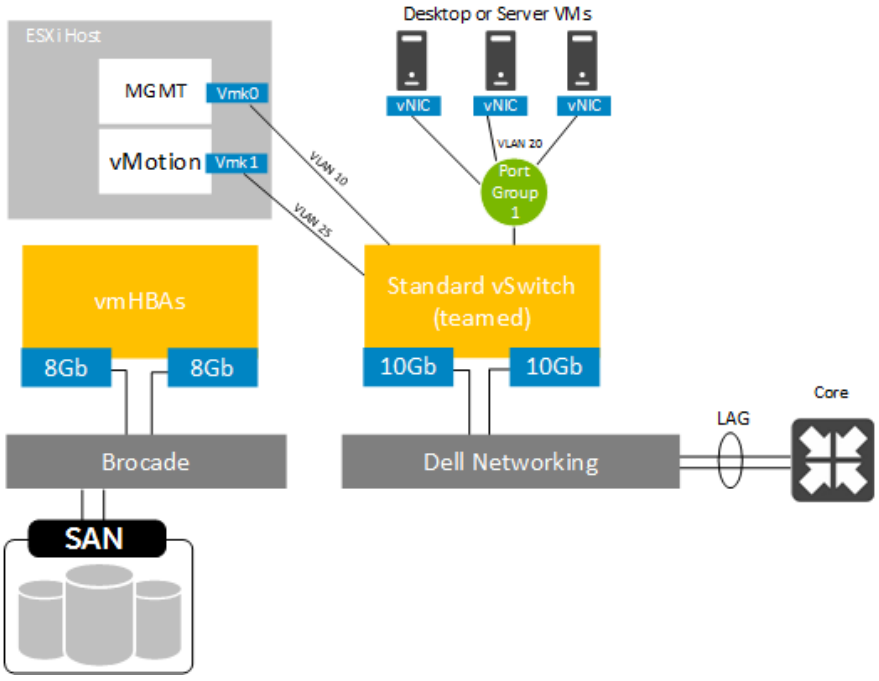

5.3.3.2 Hyper-V

The Hyper-V configuration, while identical in core requirements and hardware, is executed differently due to how Hyper-V and Windows Server 2012 R2 implement networking and virtual switches. As shown in the diagram below, native Windows Server 2012 R2 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.

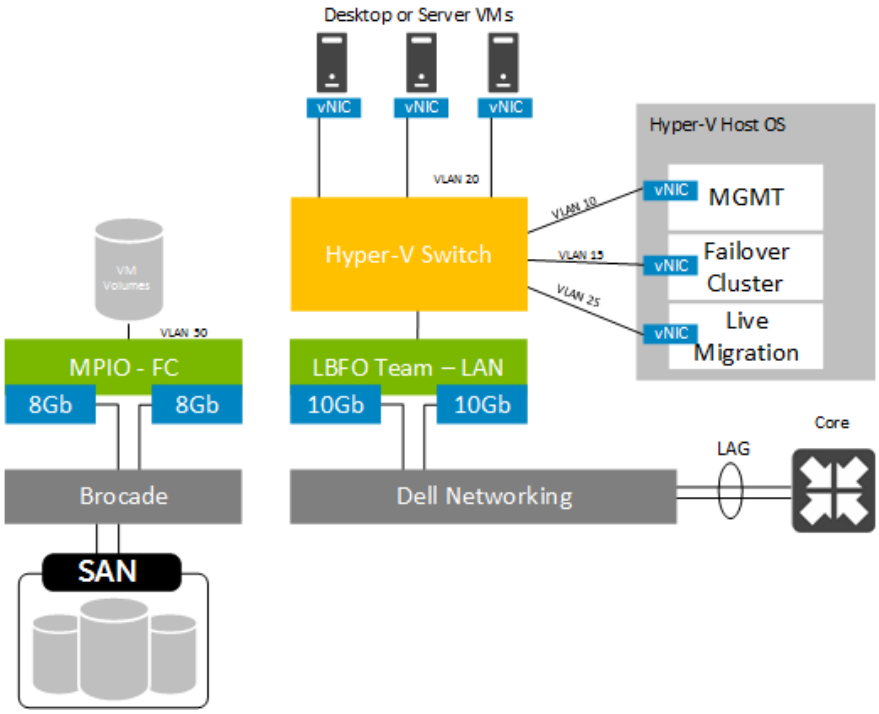

5.4 Scaling Guidance

- 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	Desktops per instance	Additional virtual RDSH servers added to the farm	Additional physical servers to host virtual RDSH servers.
Database Services	Concurrent connections, responsiveness of reads/ writes	Migrate databases to a dedicated SQL server and increase the number of management nodes	Additional RAM and CPU for the management nodes
File Services	Concurrent connections, responsiveness of reads/ writes	Split user profiles and home directories between multiple file servers in the cluster. File services can also be migrated to the optional NAS device to provide high availability.	Additional RAM and CPU for the management nodes
Monitoring Services	Managed agents/units (dependent on SQL performance as well)	Add additional monitoring servers and migrate databases to a dedicated SQL server	Additional RAM and CPU for the management nodes
Secure Gateway	Concurrent connections, Responsiveness of Network	Additional hosts, added to Farm	Additional virtual machine resources (RAM and CPU)

5.5 Solution High Availability

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

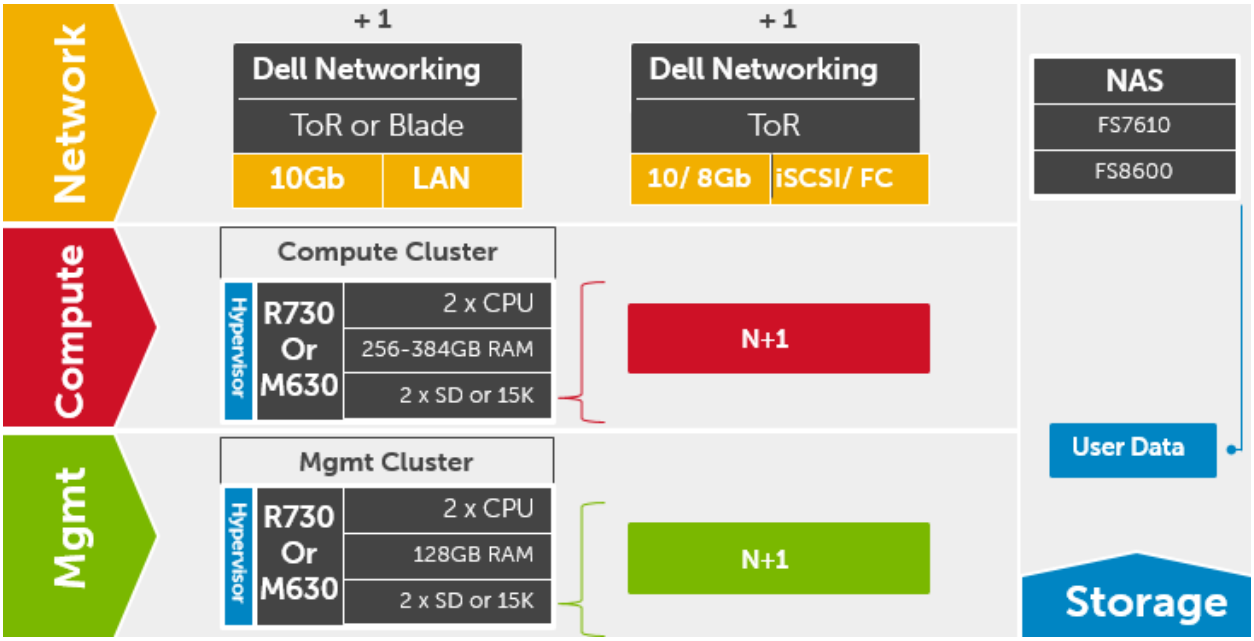

The HA options provide redundancy for all critical components in the stack while improving the performance and efficiency of the solution as a whole.

- Additional switches added to the existing thereby equally spreading each host's network connections across multiple switches.
- Additional ESXi or Hyper-V hosts added in the compute or mgmt layers to provide N+1 protection. Failover clusters are utilized with Shared Tier 1 storage for desktop HA.
- A number of enhancements occur at the Management layer, the first of which is the addition of another host. The Management hosts are configured in a failover cluster to allow live migration and protection of management VMs. All applicable management server roles can be duplicated in the cluster and utilize native load balancing functionality if available.
- SQL Server databases also are protected through the addition and configuration of an "AlwaysOn" Failover Cluster Instance or Availability Group.

- For the Storage Spaces solution, an additional SOFS host and JBOD are added to provide a "2x2" configuration (2 x SOFS hosts + 2 x JBODs).

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.

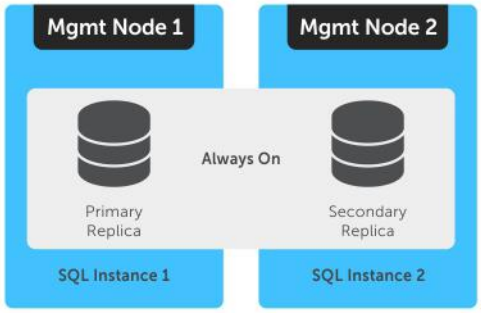

Please refer to these links for more information: LINK1 and LINK2

6 Dell Wyse Datacenter Architecture for vWorkspace

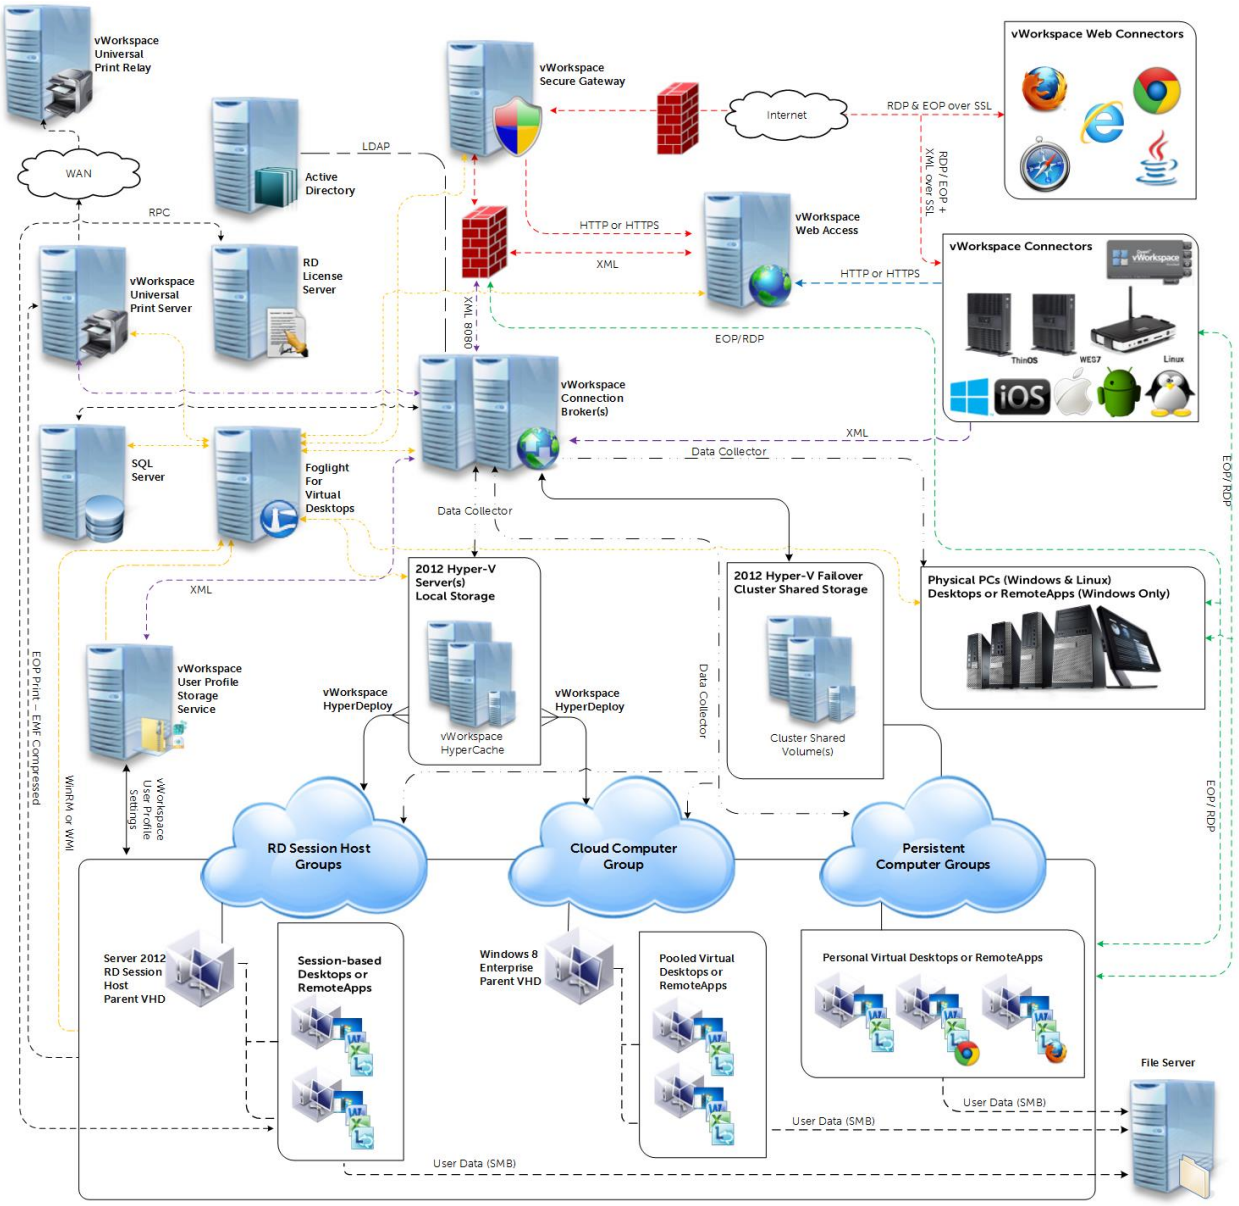

7 Solution Performance and Testing

At the time of publication here are the available density recommendations.

Hypervisor	Provisioning	Workload	Template OS	Config	User Density
Hyper-V	vWorkspace Pooled	Standard	Windows 8.1	Large Appliance	245
Hyper-V	vWorkspace Pooled	Enhanced	Windows 8.1	Large Appliance	205
Hyper-V	vWorkspace Pooled	Professional	Windows 8.1	Large Appliance	150
Hyper-V	vWorkspace Pooled	Enhanced	Server 2012 R2	Large Appliance	340
Hyper-V	vWorkspace Pooled	Standard	Windows 8.1	Small Appliance	110
Hyper-V	vWorkspace Pooled	Enhanced	Windows 8.1	Small Appliance	85
Hyper-V	vWorkspace Pooled	Professional	Windows 8.1	Small Appliance	70
Hyper-V	vWorkspace RDSH	Enhanced	Server 2012 R2	Small Appliance	185

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

Acknowledgements

Thanks to Manish Chacko of the Dell Wyse Technical Marketing team and Nicolas Cuendet of Dell Wyse Product Marketing for their support and assistance with Wyse Datacenter Appliances at Dell.

Thanks to Clayton Walther and Dilim Nwobu of the Dell Enterprise Solutions Group for their assistance and guidance on the Dell Storage with Microsoft Storage Spaces solutions.

About the Authors

Jerry Van Blaricom is a Systems Principal 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.

Peter Fine is the Sr. Principal Solutions Architect for enterprise VDI solutions at Dell. Peter has extensive experience and expertise on the broader Microsoft, Citrix and VMware solutions software stacks as well as in enterprise virtualization, storage, networking and enterprise data center design.

Rick Biedler is the Solutions Development Manager for VDI solutions at Dell, managing the development and delivery of Enterprise class Desktop virtualization solutions based on Dell Data center components and core virtualization platforms.