



Dell Wyse Datacenter for Citrix Reference Architecture

Dell Wyse Datacenter for Citrix

Dell Wyse Technical Marketing

May 2014

Revisions

Date	Description
April 2014	Initial release
May 2014	Final release

THIS DOCUMENT IS FOR INFORMATIONAL PURPOSES ONLY, AND MAY CONTAIN TYPOGRAPHICAL ERRORS AND TECHNICAL INACCURACIES. THE CONTENT IS PROVIDED AS IS, HARDWARE SELECTIONS CONTAINED WITHIN ARE FROM THE BASIS OF BEST WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND.

Copyright © 2014 Dell Inc. All rights reserved. Reproduction of this material in any manner whatsoever without the express written permission of Dell Inc. is strictly forbidden. For more information, contact Dell.

Dell, the Dell logo, and the Dell badge are trademarks of Dell Inc. Microsoft and Windows are registered trademarks of Microsoft Corporation in the United States and/or other countries. VMware is a registered trademark of VMware, Inc. Citrix and XenDesktop are registered trademarks of Citrix Systems, Inc. Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Dell Inc. disclaims any proprietary interest in trademarks and trade names other than its own.

Contents

1 Introduction.....	5
1.1 Purpose.....	5
1.2 Scope	5
1.3 What's New in This Release	5
1.4 External Publications & Industry Vertical Information:	6
2 Solution Architecture Overview	7
2.1 Introduction	7
2.1.1 Physical Architecture Overview	8
2.1.2 Dell Wyse Datacenter – Solution Layers.....	9
2.2 Local Tier 1.....	10
2.2.1 Pilot Offerings	10
2.2.2 Local Tier 1 for Rack Servers (iSCSI).....	11
2.2.3 Local Tier 1 for Blade Servers	13
2.3 Shared Tier 1 Rack.....	15
2.3.1 Shared Tier 1 – Rack – 500 Users (iSCSI – EQL)	15
2.3.2 Shared Tier 1 – Rack (iSCSI – EQL)	16
2.3.3 Shared Tier 1 – Rack – 1000 Users (FC – CML)	17
2.4 Shared Tier 1 Blade.....	19
2.4.1 Shared Tier 1 – Blade – 500 users (iSCSI – EQL).....	19
2.4.2 Shared Tier 1 – Blade (iSCSI – EQL)	20
2.4.3 Shared Tier 1 – Blade (FC – CML)	23
3 Hardware Components.....	27
3.1 Network.....	27
3.1.1 Force10 S55 (ToR Switch)	27
3.1.2 Force10 S60 (1Gb ToR Switch).....	28
3.1.3 Force10 S4810 (10Gb ToR Switch)	30
3.1.4 Brocade 6510 (FC ToR Switch)	31
3.1.5 PowerEdge M I/O Aggregator (10Gb Blade Interconnect).....	32
3.1.6 PowerConnect M6348 (1Gb Blade Interconnect)	33
3.1.7 Brocade M5424 (FC Blade Interconnect)	33
3.2 Servers	35

3.2.1 PowerEdge R720	35
3.2.2 PowerEdge M620	35
3.3 Storage	36
3.3.1 EqualLogic Tier 1 Storage (iSCSI)	36
3.3.2 EqualLogic Tier 2 Storage (iSCSI)	36
3.3.3 EqualLogic Configuration	41
3.3.4 Compellent Storage (FC)	41
3.3.5 NAS	44
3.4 Dell Wyse Cloud Clients	45
3.4.1 ThinOS – T10D	45
3.4.2 ThinOS – D10D	46
3.4.3 Windows Embedded 7 – Z90Q7	46
3.4.4 Windows Embedded 8 – Z90Q8	46
3.4.5 Suse Linux – Z50D	47
3.4.6 Dell Wyse Zero – Xenith 2	47
3.4.7 Dell Wyse Zero – Xenith Pro 2	47
3.4.8 Dell Wyse Cloud Connect	48
3.4.9 Dell Venue 11 Pro	48
3.4.10 Dell Chromebook 11	48
4 Software Components	49
4.1 Citrix XenDesktop	49
4.1.1 Machine Creation Services (MCS)	50
4.1.2 Provisioning Services (PVS)	51
4.1.3 Personal vDisk	52
4.1.4 Citrix Profile Manager	53
4.1.5 XenApp	53
4.2 VDI Hypervisor Platforms	56
4.2.1 VMware vSphere 5	56
4.2.2 Microsoft Windows Server 2012 Hyper-V	56
4.3 Citrix NetScaler	57
4.4 Citrix CloudBridge	59
5 Solution Architecture for XenDesktop 7	60
5.1 Compute Server Infrastructure	60
5.1.1 Local Tier 1 Rack	60
5.1.2 Local Tier 1 Blade	61
5.1.3 Shared Tier 1 Rack	61

5.1.4 Shared Tier 1 Blade	64
5.1.5 XenApp Compute Host	66
5.2 Management Server Infrastructure.....	67
5.2.1 SQL Databases	67
5.2.2 DNS.....	68
5.3 Storage Architecture Overview.....	68
5.3.1 Local Tier 1 Storage	68
5.3.2 Shared Tier 1 Storage.....	69
5.3.3 Shared Tier 2 Storage	70
5.3.4 Storage Networking – EqualLogic iSCSI	70
5.3.5 Storage Networking – Compellent Fiber Channel (FC)	71
5.4 Virtual Networking	72
5.4.1 Local Tier 1 – Rack – iSCSI	72
5.4.2 Local Tier 1 – Blade.....	76
5.4.3 Shared Tier 1 – Rack – iSCSI	78
5.4.4 Shared Tier 1 – Rack – FC.....	81
5.4.5 Shared Tier 1 – Blade – iSCSI	84
5.4.6 Shared Tier 1 – Blade – FC	87
5.5 Scaling Guidance.....	89
5.5.1 Windows 8.1 – vSphere.....	91
5.5.2 Windows 8.1 – Hyper-V	91
5.6 Solution High Availability.....	92
5.6.1 Compute Layer HA (Local Tier 1).....	93
5.6.2 vSphere HA (Shared Tier 1)	94
5.6.3 Hyper-V HA (Shared Tier 1)	95
5.6.4 Management Server High Availability.....	95
5.6.5 XenApp Server High Availability	96
5.6.6 Provisioning Services High Availability	96
5.6.7 Windows File Services High Availability	97
5.6.8 SQL Server High Availability	98
5.7 Dell Wyse Datacenter for XenDesktop Communication Flow	99
6 Customer Provided Stack Components	100
6.1 Customer Provided Storage Requirements.....	100
6.2 Customer Provided Switching Requirements	101
7 Solution Performance and Testing.....	102
7.1 Load Generation and Monitoring	102

7.1.1 Login VSI 4 – Login Consultants	102
7.1.2 Liquidware Labs Stratusphere UX	102
7.1.3 EqualLogic SAN HQ	103
7.1.4 VMware vCenter	103
7.1.5 Microsoft Perfmon	103
7.2 Performance Analysis Methodology	104
7.2.1 Resource Utilization	104
7.2.2 EUE (Tools Info)	106
7.2.3 EUE (Real User Info)	106
7.2.4 Dell Wyse Datacenter Workloads and Profiles	106
7.2.5 Dell Wyse Datacenter Profiles	107
7.2.6 Dell Wyse Datacenter Workloads	107
7.2.7 Workloads Running on Shared Graphics Profile	109
7.2.8 Workloads Running on Pass-through Graphics Profile	109
7.3 Testing and Validation	110
7.3.1 Testing Process	110
7.4 XenDesktop Test Results	111
7.4.1 PVS Write Cache Comparisons (for XenDesktop)	113
7.4.2 Recommended Densities	116
7.4.3 VMware vSphere 5.5 Test Charts	116
7.4.4 Microsoft Windows 2012 R2 Hyper-V Test Charts	123
7.4.5 Local Tier 1 for Blades Test Results	128
7.4.6 XenDesktop with Personal vDisk Enabled	139
7.4.7 EqualLogic PS6210XS Hybrid Array Testing Summary	140
7.4.8 Dell Compellent Storage Testing Summary	148
7.5 XenApp Test Results	153
7.5.1 Configuration	153
7.5.2 Test Results Summary	154
7.5.3 VMware vSphere 5.5 Test Charts	159
7.5.4 Microsoft Windows 2012 R2 Hyper-V Test Charts	164
7.5.5 Microsoft Windows 2012 R2 Physical Server	171
7.5.6 XenApp Scaling Guidance	172
Acknowledgements	174
About the Authors	175

1 Introduction

1.1 Purpose

This document describes:

1. Dell Wyse Datacenter for Citrix XenDesktop Reference Architecture scaling from 50 to 50K+ VDI users.
2. Solution options encompass a combination of solution models including local disks, iSCSI or Fiber Channel based storage options.

This document addresses the architecture design, configuration and implementation considerations for the key components of the architecture required to deliver virtual desktops via XenDesktop 7.1 on Windows Server Hyper-V 2012 R2 or VMware vSphere 5.

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 design 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 solution scaling and component selection guidance.

1.3 What's New in This Release

- Citrix XenDesktop 7.1 support (including Windows 8.1, Server 2012 R2)
- Citrix XenApp sizing and guidance updates
- Expanded PVS write cache comparison and guidance
- Support for new EqualLogic 6210 series
- Dell Wyse Cloud Client updates
- Test results updated for XenDesktop 7.1 using Login VSI 4
- High performance SAN-Less offering for blade servers: [LINK](#)
- High IOPS persistent desktops using Atlantis ILIO: [LINK](#)
- Citrix vGPU for graphics acceleration offering: [LINK](#)

1.4 External Publications & Industry Vertical Information:

- Dell Wyse Datacenter for Citrix and Dell PowerEdge VRTX: [LINK](#)
- Dell Wyse Datacenter for Citrix and Microsoft Lync 2013: [LINK](#)
- Dell Wyse Datacenter for Mobile Clinical Computing: [LINK](#)
- Dell Wyse Datacenter for Government: [LINK](#)
- Dell Wyse Datacenter for Education: [LINK](#)
- Validated 3000 User PVS deployment using Dell Compellent All Flash Array: [LINK](#)

2 Solution Architecture Overview

2.1 Introduction


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
- Cloud Client Layer


These components have been integrated and tested to provide the optimal balance of high performance and lowest cost per user. Additionally, the Dell Wyse Datacenter Solution includes an approved extended list of optional components in the same categories. These components give IT departments the flexibility to custom tailor the solution for environments with unique VDI features, scale or performance needs. The Dell Wyse Datacenter stack is designed to be a cost effective starting point for IT departments looking to migrate to a fully virtualized desktop environment slowly. This approach allows you to grow the investment and commitment as needed or as your IT staff becomes more comfortable with VDI technologies.

2.1.1 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 profile/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.




2.1.2 Dell Wyse Datacenter – Solution Layers


Only a single high performance Force10 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 LAN and iSCSI traffic be separated into discrete switching Fabrics. Additional switches are added and stacked as required to provide High Availability for the Network layer.




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 pictured below).



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 VMs necessary to support the VDI infrastructure.



The storage layer consists of options provided by EqualLogic (EQL) for iSCSI and Compellent arrays for Fiber Channel to suit your Tier 1 and Tier 2 scaling and capacity needs.

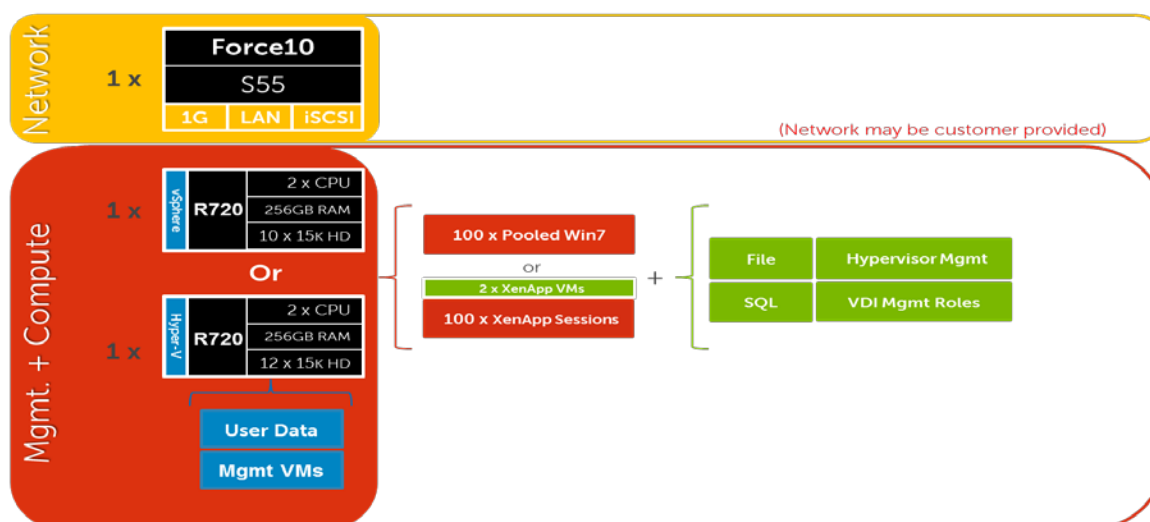


2.2 Local Tier 1

2.2.1 Pilot Offerings

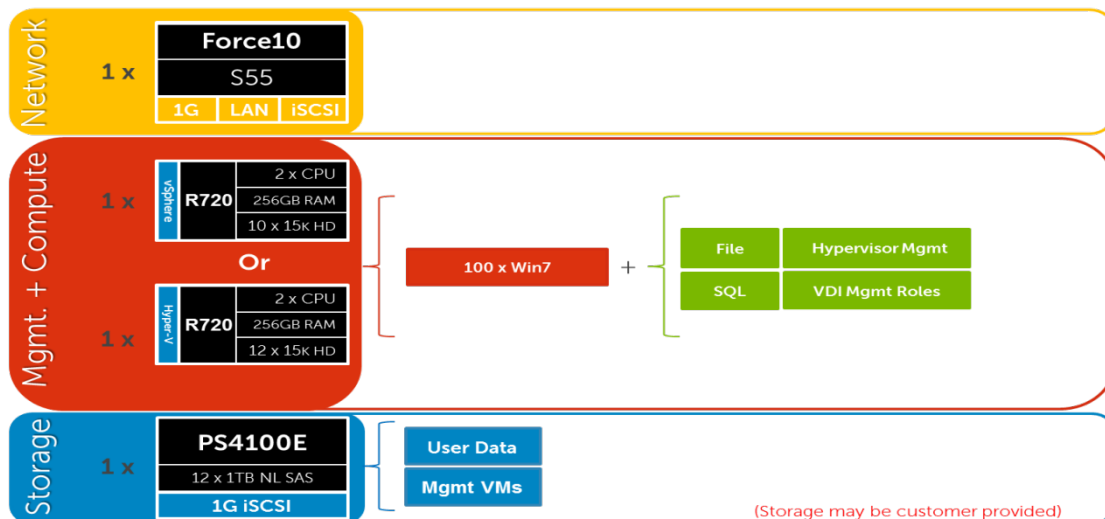
2.2.1.1 100 User Combined Pilot

For a very small deployment or pilot effort to familiarize oneself with the solution architecture, we offer a 100 user combined pilot solution. This architecture is non-distributed with all VDI, Management, and storage functions on a single host running either vSphere or Hyper-V. If additional scaling is desired, you can grow into a larger distributed architecture seamlessly with no loss on initial investment. Our recommended delivery mechanism for this architecture is Citrix Machine Creation Services (MCS).



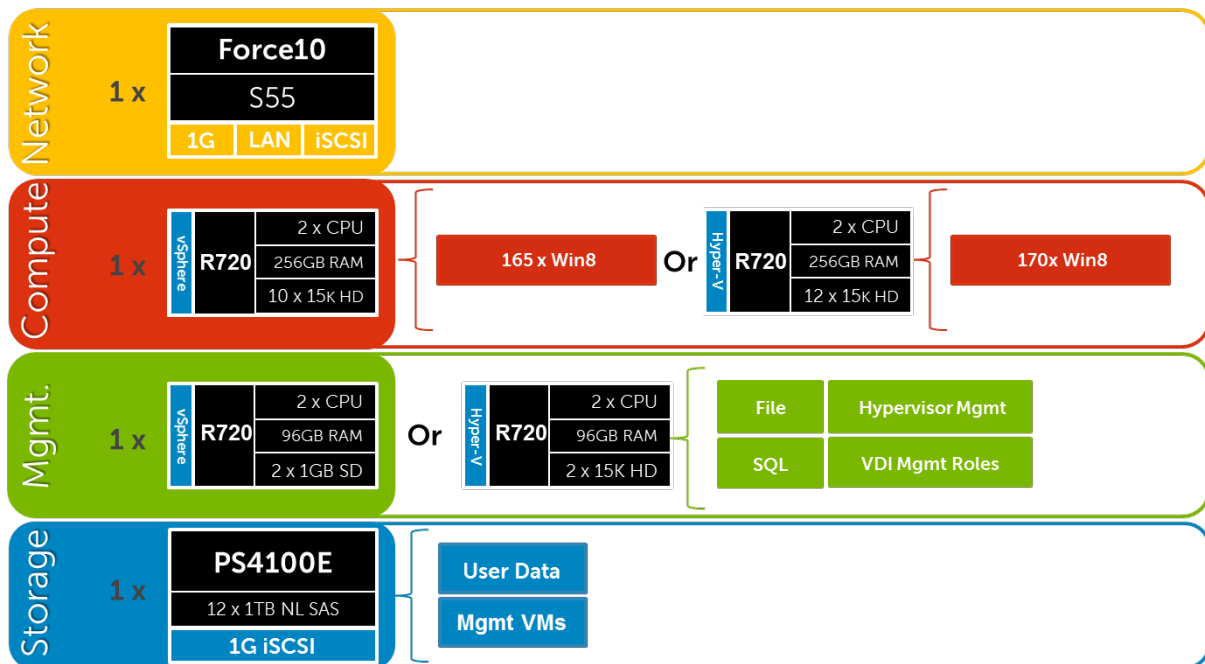
2.2.1.2 100 User Scale Ready Pilot

In addition to the 100 user combined offering we also offer a scale ready version that includes Tier 2 storage. The basic architecture is the same but customers looking to scale out quickly benefit by building out into Tier 2 initially.



2.2.2 Local Tier 1 for Rack Servers (iSCSI)


The Local Tier 1 solution model provides a scalable rack-based configuration that hosts user VDI sessions on local disk in the compute layer. vSphere or Hyper-V based solutions are available and scale based on the chosen hypervisor.



2.2.2.1 Local Tier 1 – Network Architecture (iSCSI)

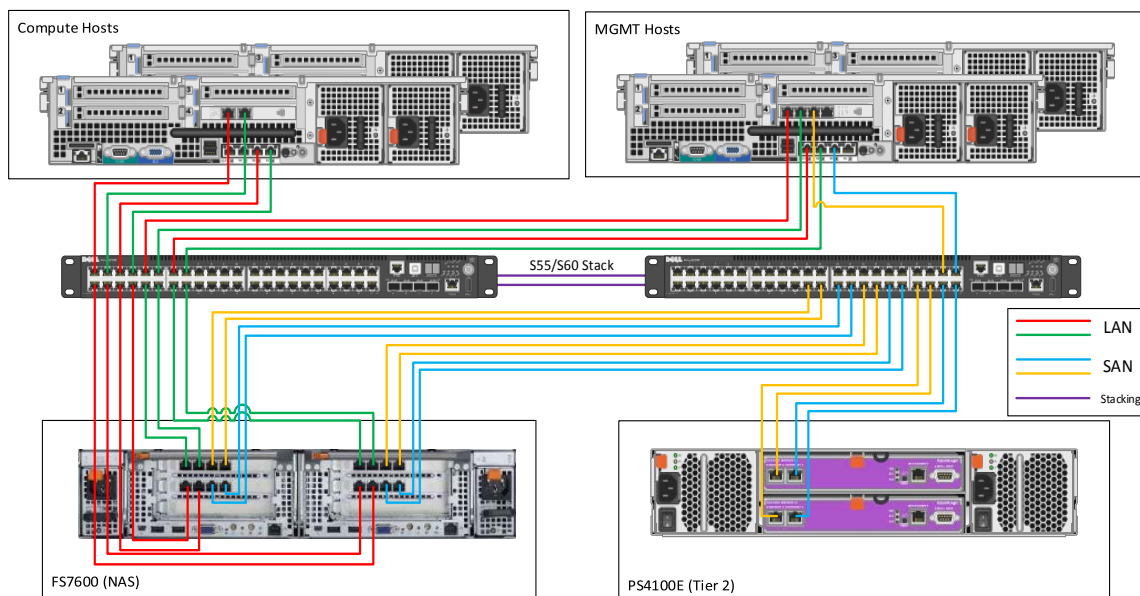
In the Local Tier 1 architecture, a single Force10 switch is shared among all network connections for both management and compute, up to 1000 users. Over 1000 users Dell Wyse Solutions Engineering recommends separating the network Fabrics to isolate iSCSI and LAN traffic as well as making each switch stack redundant. Only the management servers connect to iSCSI storage in this model. All Top of Rack (ToR) traffic has been designed to be 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.



2.2.2.2 Local Tier 1 – Cabling (Rack – HA)

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

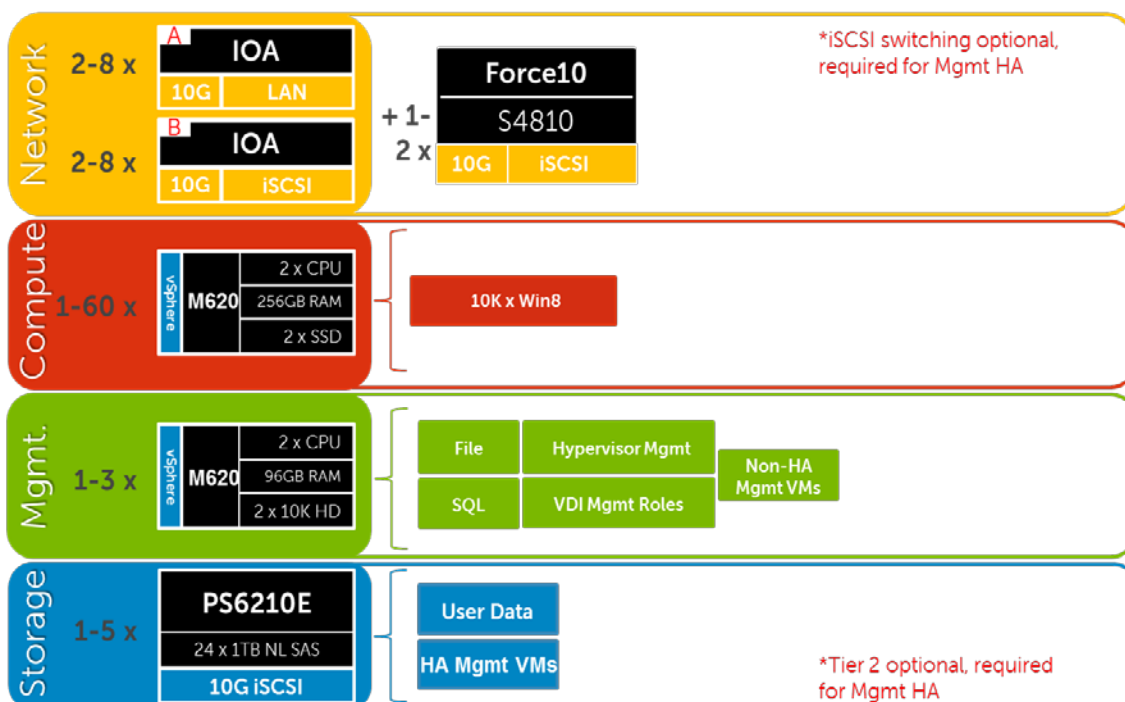


2.2.2.3 Local Tier 1 – Rack Scaling Guidance (iSCSI)

Local Tier 1 HW Scaling (iSCSI)				
User Scale	ToR LAN	ToR 1Gb iSCSI	EQL T2	EQL NAS
0-1000	S55		4100E	
0-1000 (HA)	S55	S55	4100E	FS7600
0-3000	S55	S55	6100E	FS7600
3000-6000	S55	S55	6500E	FS7600
6000+ users	S60	S60	6500E	FS7600


2.2.3 Local Tier 1 for Blade Servers

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




2.2.3.1 Local Tier 1 – Network Architecture

In the Local Tier 1 architecture for blades, there is no need for a ToR switch unless 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.



2.2.3.2 Local Tier 1 – Cabling (Blades - 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.



2.2.3.3 Local Tier 1 – Blade Scaling Guidance

If limited rack depth is an issue at higher user scales, the 6210E are used in lieu of the 6510E.

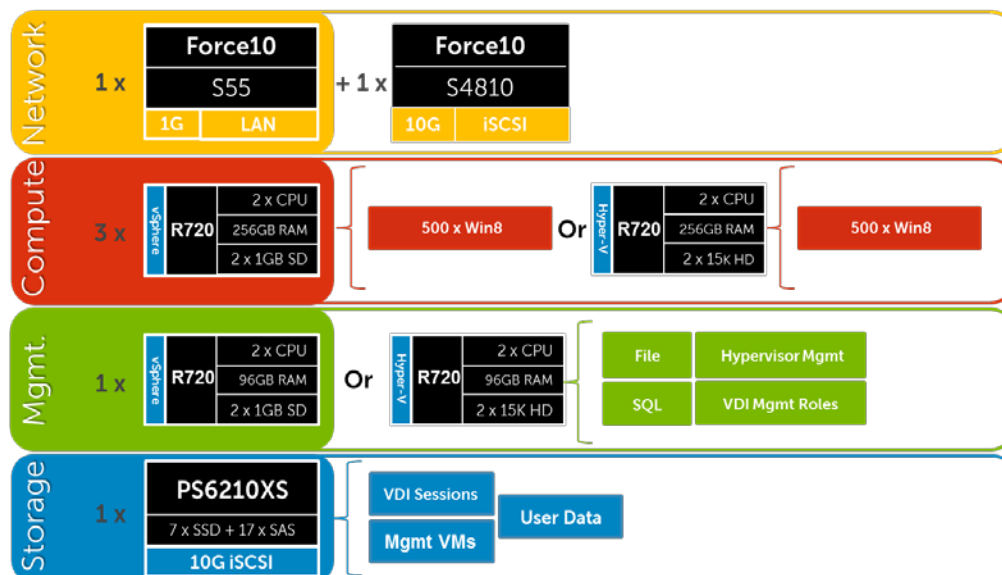
Local Tier 1 HW scaling (Blade)				
User Scale	Blade LAN (A Fabric)	Blade iSCSI (B Fabric)	ToR 10Gb iSCSI	EQL T2
0-10000 (no HA)	IOA	-	-	-
0-1000 (HA)	IOA	IOA	S4810	4110E
0-6000 (HA)	IOA	IOA	S4810	6210E
6000+ (HA)	IOA	IOA	S4810	6510E

2.3 Shared Tier 1 Rack

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

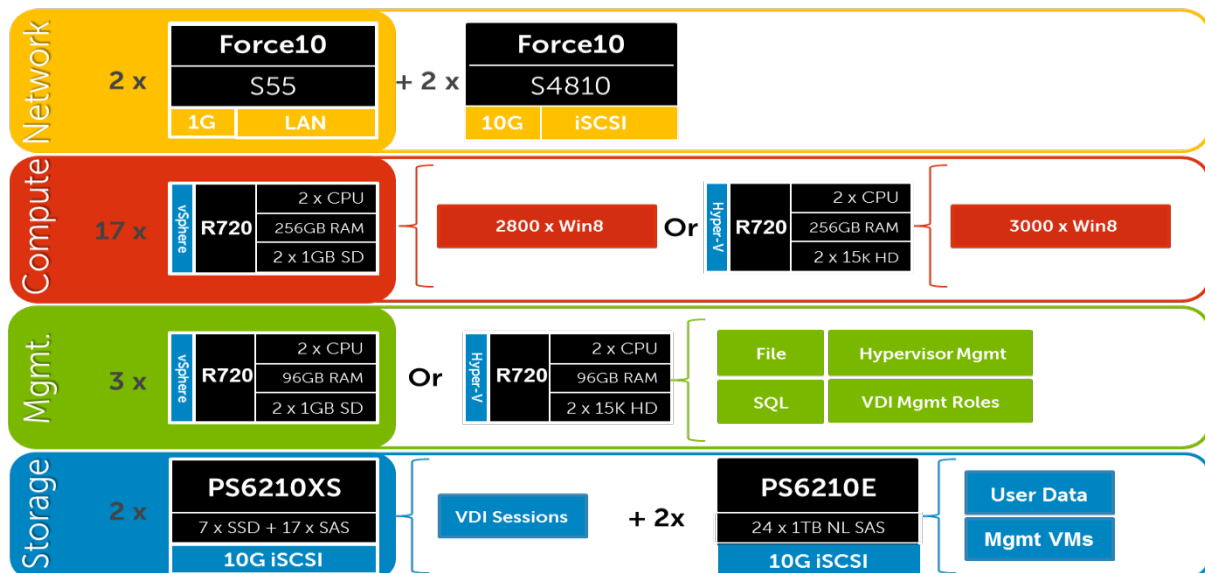
2.3.1 Shared Tier 1 – Rack – 500 Users (iSCSI – EQL)

For POCs or small deployments, Tier1 and Tier2 are combined on a single 6210XS storage array. Above 500 users, a separate array is used for Tier 2.



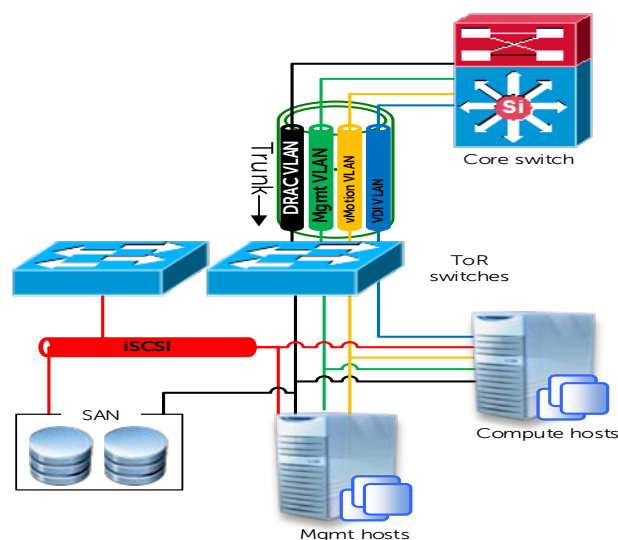
2.3.2 Shared Tier 1 – Rack (iSCSI – EQL)

For over 500 users on EqualLogic, the storage layers are separated into discrete arrays. The figure below depicts a 3000 user build with separate network Fabrics for LAN and iSCSI traffic. Additional 6210XS arrays are added for Tier 1 as the user count scales, just as the Tier 2 array models change also based on scale. The 4110E, 6210E, and 6510E are 10Gb Tier 2 array options. NAS is recommended above 1000 users to provide HA for file services.




2.3.2.1 Shared Tier 1 Rack – Network Architecture (iSCSI)

In the Shared Tier 1 architecture for rack servers, both management and compute servers connect to shared storage. 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.



2.3.2.2 Shared Tier 1 Cabling – (Rack – EQL)

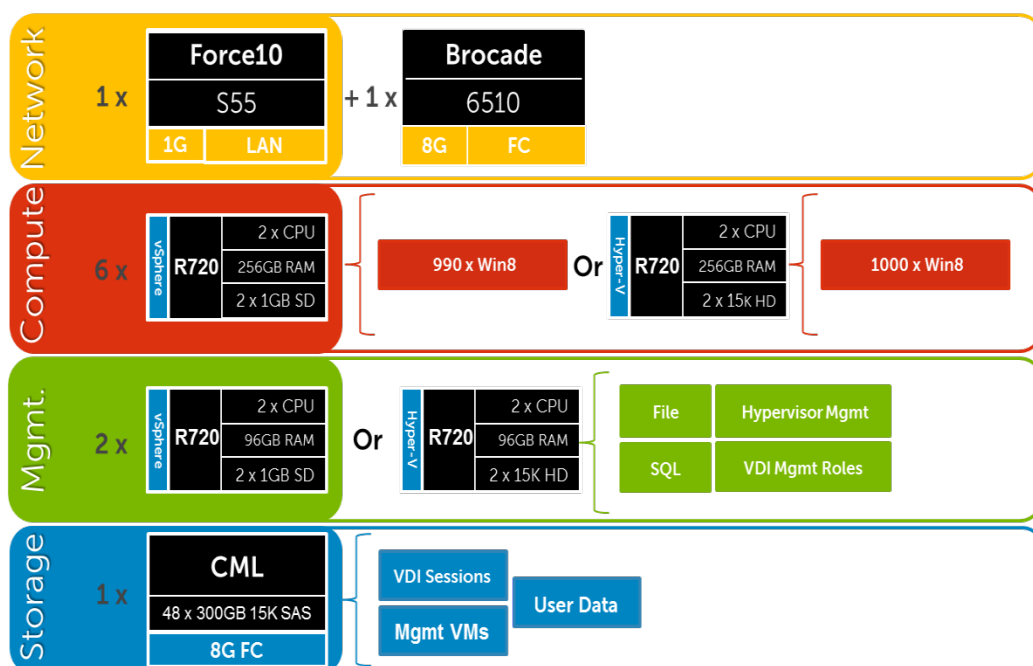


2.3.2.3 Shared Tier 1 Rack Scaling Guidance (iSCSI)

Shared Tier 1 HW scaling (Rack – iSCSI)					
User Scale	ToR LAN	ToR 10Gb iSCSI	EQL T1	EQL T2	EQL NAS
0-500	S55	S4810	6210XS	-	-
500-1000	S55	S4810	6210XS	4110E	-
0-1000 (HA)	S55	S4810	6210XS	4110E	NX3300
0-3000	S55	S4810	6210XS	6210E	NX3300
3000-6000	S55	S4810	6210XS	6510E	NX3300
6000+	S60	S4810	6210XS	6510E	NX3300


2.3.3 Shared Tier 1 – Rack – 1000 Users (FC – CML)

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 1000 user array. For 2000 users, 96 total disks in 4 shelves are required. Please see section 3.3.3 for more information.




2.3.3.1 Shared Tier 1 Rack – Network Architecture (FC)

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.



2.3.3.2 Shared Tier 1 Cabling (Rack – CML)



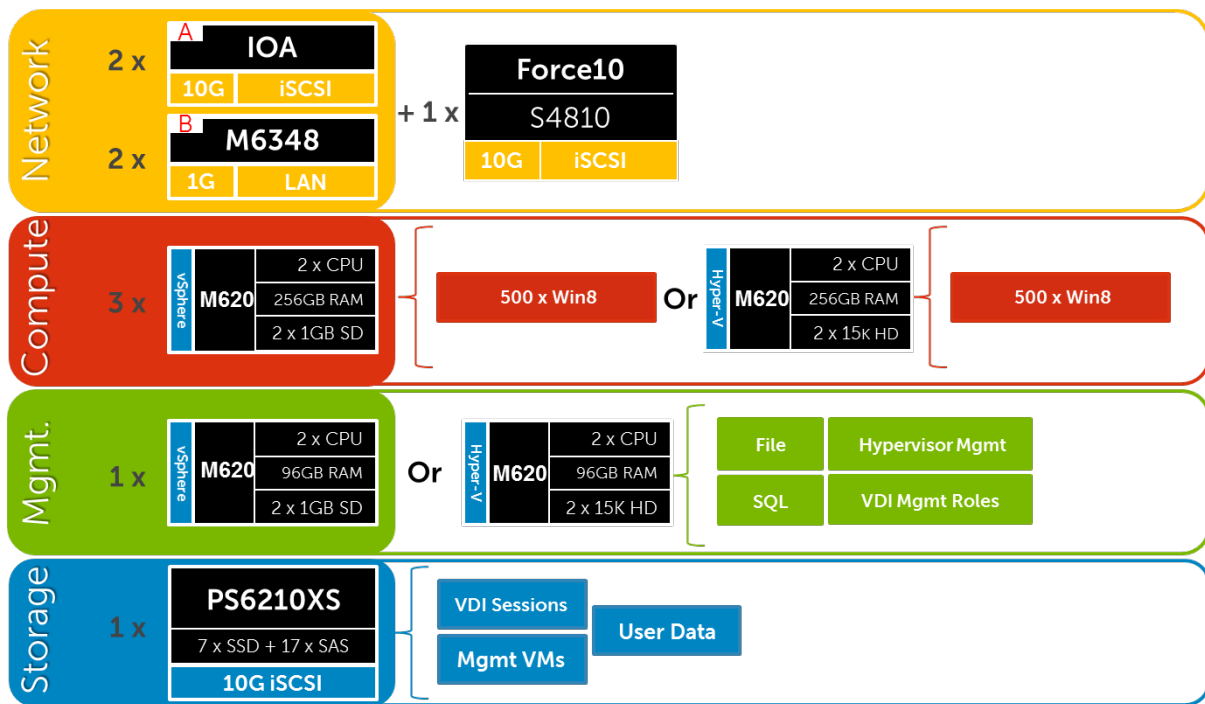
2.3.3.3 Shared Tier 1 Rack Scaling Guidance (FC)

Shared Tier 1 HW scaling (Rack - FC)					
User Scale	LAN Network	FC Network	CML T1	CML T2	CML NAS
0-1000	S55	6510	15K SAS	-	-
0-1000 (HA)	S55	6510	15K SAS	NL SAS	FS8600
1000-6000	S55	6510	15K SAS	NL SAS	FS8600
6000+	S60	6510	15K SAS	NL SAS	FS8600

2.4 Shared Tier 1 Blade

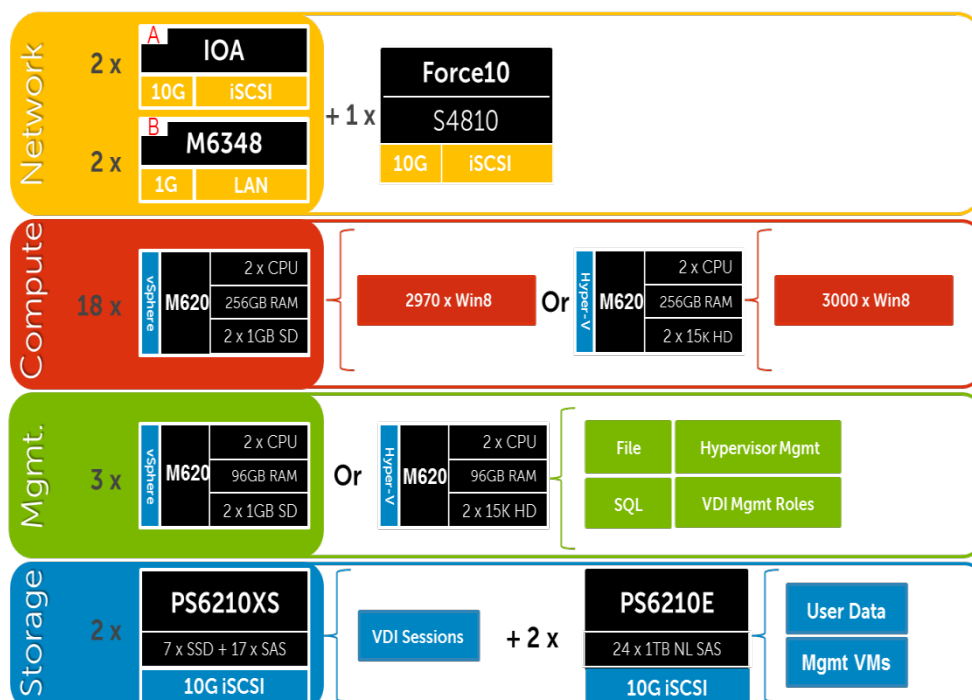
2.4.1 Shared Tier 1 – Blade – 500 users (iSCSI – EQL)

As is the case in the Shared Tier 1 model using rack servers, blades can also be used in a 500 user bundle by combining Tier 1 and Tier 2 on a single 6210XS array. Above 500 users, separate Tier 1 and Tier 2 storage into discrete arrays.




2.4.2 Shared Tier 1 – Blade (iSCSI – EQL)

Above 1000 users the Storage Tiers need to be separated to maximize the performance of the 6210XS for VDI sessions. At this scale we also separate LAN from iSCSI switching. Load balancing and NAS are added optionally for HA. The figure below depicts a 3000 user solution:




2.4.2.1 Shared Tier 1 Blade – Network Architecture (iSCSI)

In the Shared Tier 1 architecture for blades, only iSCSI is switched through a ToR switch. There is no need to switch LAN ToR since the M6348 in the chassis supports LAN to the blades and are uplinked to the core or distribution layers directly. The M6348 has 16 external ports per switch that are optionally used for DRAC/ IPMI traffic. For greater redundancy, a ToR switch is used to support iDRAC/IPMI used outside of the chassis. Both Management and Compute servers connect to all VLANs in this model. The following diagram illustrates the chassis switch uplink ports to ToR/ core switch connections, vSwitch assignments, as well as logical VLAN flow in relation to the core switch.



2.4.2.2 Shared Tier 1 Blade Cabling (Blade – EQL)

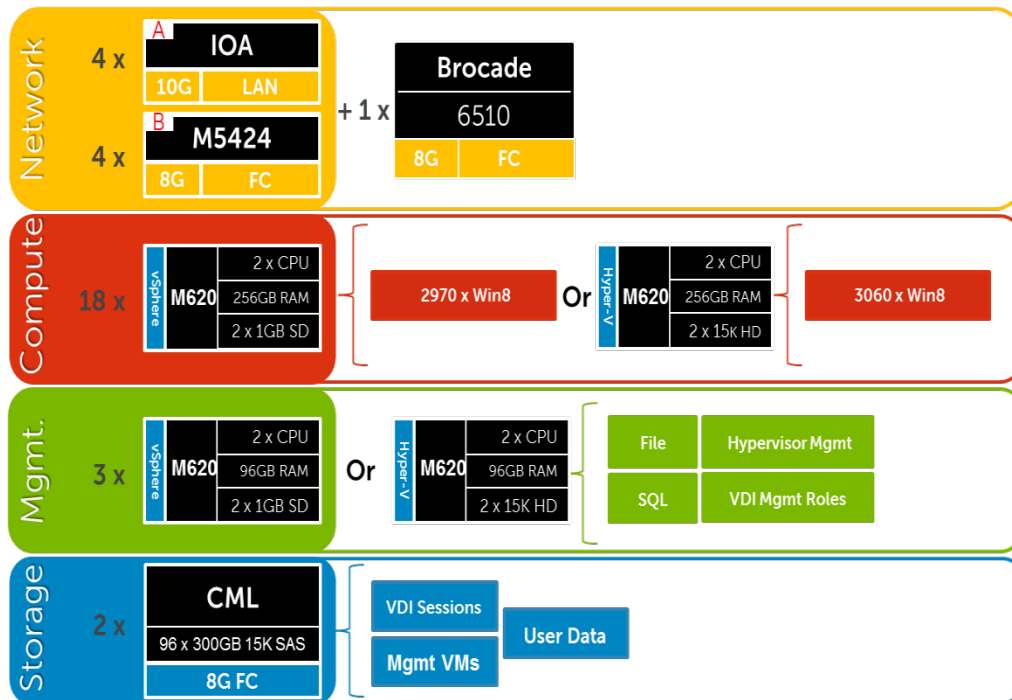


2.4.2.3 Shared Tier 1 Blade Scaling Guidance (iSCSI)


Shared Tier 1 HW scaling (Blade - iSCSI)						
User Scale	Blade LAN (B Fabric)	Blade iSCSI (A Fabric)	ToR 10Gb iSCSI	EQL T1	EQL T2	EQL NAS
0-500	M6348	IOA	S4810	6210XS	-	-
500-1000	M6348	IOA	S4810	6210XS	4110E	-
0-1000 (HA)	M6348	IOA	S4810	6210XS	4110E	NX3300
0-3000	M6348	IOA	S4810	6210XS	6210E	NX3300
3000-6000	M6348	IOA	S4810	6210XS	6510E	NX3300
6000+	M6348	IOA	S4810	6210XS	6510E	NX3300

2.4.3 Shared Tier 1 – Blade (FC – CML)


Fiber channel is also an option in the Shared Tier 1 model for blades. There are a few key differences using FC with blades instead of iSCSI: Blade chassis interconnects, FC HBAs in the servers, and FC IO cards in the Compellent arrays. ToR FC switching is optional if a suitable FC infrastructure is already in place. The image below depicts a 4000 user stack.



2.4.3.1 Shared Tier 1 Blade – Network Architecture (FC)



2.4.3.2 Shared Tier 1 Cabling (Blade – CML)



2.4.3.3 Shared Tier 1 Blade Scaling Guidance (FC)

Shared Tier 1 HW scaling (Blade - FC)						
User Scale	Blade LAN (A Fabric)	Blade FC (B Fabric)	ToR FC	CML T1	CML T2	CML NAS
0-500	IOA	5424	6510	15K SAS	-	-
500-1000	IOA	5424	6510	15K SAS	-	-
0-1000 (HA)	IOA	5424	6510	15K SAS	NL SAS	FS8600
1000-6000	IOA	5424	6510	15K SAS	NL SAS	FS8600
6000+	IOA	5424	6510	15K SAS	NL SAS	FS8600

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 Force10 S55 (ToR Switch)

The Dell Force10 S-Series S55 1/10Gb ToR (Top-of-Rack) switch is optimized for lowering operational costs while increasing scalability and improving manageability at the network edge. Optimized for high-performance data center applications, the S55 is recommended for Dell Wyse Datacenter deployments of 6000 users or less and leverages a non-blocking architecture that delivers line-rate, low-latency L2 and L3 switching to eliminate network bottlenecks. The high-density S55 design provides 48Gb access ports with up to four modular 10Gb uplinks in just 1-RU to conserve valuable rack space. The S55 incorporates multiple architectural features that optimize data center network efficiency and reliability, including IO panel to PSU airflow or PSU to IO panel airflow for hot/cold aisle environments, and redundant, hot-swappable power supplies and fans. A "scale-as-you-grow" ToR solution that is simple to deploy and manage, up to 8 S55 switches are stacked to create a single logical switch by utilizing Dell Force10's stacking technology and high-speed stacking modules.

Model	Features	Options	Uses
Force10 S55	44 x BaseT (10/100/1000) + 4 x SFP	Redundant PSUs	ToR switch for LAN and iSCSI in Local Tier 1 solution
		4 x 1Gb SFP ports that support copper or fiber	
		12Gb or 24Gb stacking (up to 8 switches)	
		2 x modular slots for 10Gb uplinks or stacking modules	




Guidance:

- 10Gb uplinks to a core or distribution switch are the preferred design choice using the rear 10Gb uplink modules. If 10Gb to a core or distribution switch is unavailable the front 4 x 1Gb SFP ports are used.
- The front 4 SFP ports can support copper cabling and are upgraded to optical if a longer run is needed.

For more information on the S55 switch and Dell Force10 networking, please visit: [LINK](#)

3.1.1.1 Force10 S55 Stacking

The Top of Rack switches in the Network layer are optionally stacked with additional switches, if greater port count or redundancy is desired. Each switch will need a stacking module plugged into a rear bay and connected with a stacking cable. The best practice for switch stacks greater than 2 is to cable in a ring configuration with the last switch in the stack cabled back to the first. Uplinks need to be configured on all switches in the stack back to the core to provide redundancy and failure protection.




Please reference the following Force10 whitepaper for specifics on stacking best practices and configuration: [LINK](#)

3.1.2 Force10 S60 (1Gb ToR Switch)

The Dell Force10 S-Series S60 is a high-performance 1/10Gb access switch optimized for lowering operational costs at the network edge and is recommended for Dell Wyse Datacenter deployments over 6000 users. The S60 answers the key challenges related to network congestion in data center ToR (Top-of-Rack) and service provider aggregation deployments. As the use of bursty applications and services continue to increase, huge spikes in network traffic that can cause network congestion and packet loss, also become more common. The S60 is equipped with the industry's largest

packet buffer (1.25 GB), enabling it to deliver lower application latency and maintain predictable network performance even when faced with significant spikes in network traffic. Providing 48 line-rate Gb ports and up to four optional 10Gb uplinks in just 1-RU, the S60 conserves valuable rack space. Further, the S60 design delivers unmatched configuration flexibility, high reliability, and power and cooling efficiency to reduce costs.

Model	Features	Options	Uses
Force10 S60	44 x BaseT (10/100/1000) + 4 x SFP High performance High Scalability	Redundant PSUs	Higher scale ToR switch for LAN in Local + Shared Tier 1 and iSCSI in Local Tier 1 solution
		4 x 1Gb SFP ports the support copper or fiber	
		12Gb or 24Gb stacking (up to 12 switches)	
		2 x modular slots for 10Gb uplinks or stacking modules	



Guidance:

- 10Gb uplinks to a core or distribution switch are the preferred design choice using the rear 10Gb uplink modules. If 10Gb to a core or distribution switch is unavailable the front 4 x 1Gb SFP ports are used.
- The front 4 SFP ports can support copper cabling and are upgraded to optical if a longer run is needed.
- The S60 is appropriate for use in solutions scaling higher than 6000 users.

For more information on the S60 switch and Dell Force10 networking, please visit: [LINK](#)

3.1.2.1 S60 Stacking

The S60 switches are optionally stacked with 2 or more switches, if greater port count or redundancy is desired. Each switch will need a stacking module plugged into a rear bay and


connected with a stacking cable. The best practice for switch stacks greater than 2 is to cable in a ring configuration with the last switch in the stack cabled back to the first. Uplinks need to be configured on all switches in the stack back to the core to provide redundancy and failure protection.



3.1.3 Force10 S4810 (10Gb ToR Switch)

The Dell Force10 S-Series S4810 is an ultra-low latency 10/40Gb Top-of-Rack (ToR) switch purpose-built for applications in high-performance data center and computing environments. Leveraging a non-blocking, cut-through switching architecture, the S4810 delivers line-rate L2 and L3 forwarding capacity with ultra-low latency to maximize network performance. The compact S4810 design provides industry-leading density of 48 dual-speed 1/10Gb (SFP+) ports as well as four 40Gb QSFP+ uplinks to conserve valuable rack space and simplify the migration to 40Gb in the data center core (Each 40Gb QSFP+ uplink can support four 10Gb ports with a breakout cable). Priority-based Flow Control (PFC), Data Center Bridge Exchange (DCBX), Enhance Transmission Selection (ETS), coupled with ultra-low latency and line rate throughput, make the S4810 ideally suited for iSCSI storage, FCoE Transit & DCB environments.

Model	Features	Options	Uses
Force10 S4810	48 x SFP+ (1Gb/10Gb) + 4 x QSFP+ (40Gb) Redundant Power Supplies	Single-mode/ multimode optics, TwinAx, QSFP+ breakout cables	ToR switch for iSCSI in Shared Tier 1 solution
		Stack up to 6 switches or 2 using VLT, using SFP or QSFP ports	




Guidance:

- The 40Gb QSFP+ ports are split into 4 x 10Gb ports using breakout cables for stand-alone units, if necessary. This is not supported in stacked configurations.
- 10Gb or 40Gb uplinks to a core or distribution switch is the preferred design choice.
- The front 4 SFP ports can support copper cabling and are upgraded to optical if a longer run is needed.
- The S60 is appropriate for use in solutions scaling higher than 6000 users.

For more information on the S4810 switch and Dell Force10 networking, please visit: [LINK](#)

3.1.3.1 S4810 Stacking


The S4810 switches are optionally stacked up to 6 switches or configured to use Virtual Link Trunking (VLT) up to 2 switches. Stacking is supported on either SFP or QSFP ports as long as that port is configured for stacking. The best practice for switch stacks greater than 2 is to cable in a ring configuration with the last switch in the stack cabled back to the first. Uplinks need to be configured on all switches in the stack back to the core to provide redundancy and failure protection.



3.1.4 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.



Guidance:

- The 6510 FC switches are licensed to limit the number of ports required for the deployment. If only 24 or fewer ports are required for a given implementation, then only those need to be licensed.
- Up to 239 Brocade switches are used in a single FC Fabric.

For more information on the Brocade 6510 switch, please visit: [LINK](#)

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

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	




Guidance:

- Onboard QSFP+ ports are used for stacking a pair of IOAs within a single chassis.
- An additional SFP+ or 10GBT FlexIO module is added to provide uplinks.

For more information on the Dell IOA switch, please visit: [LINK](#)

3.1.6 PowerConnect M6348 (1Gb Blade Interconnect)

Model	Features	Options	Uses
PowerConnect M6348	32 x internal (1Gb) + 16 x external Base-T + 2 x 10Gb SFP+ + 2 x 16Gb stacking/ CX4 ports	Stack up to 12switches	Blade switch for LAN traffic in Shared Tier 1 blade solution



Guidance:


- 10Gb uplinks to a core or distribution switch are the preferred design choice using TwinAx or optical cabling via the SFP+ ports.
- 16 x external 1Gb ports are used for Management ports, iDRACs and IPMI.
- Stack up to 12 switches using stacking ports.

3.1.7 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.




Guidance:

- 12 port model includes 2 x 8Gb transceivers, 24 port models include 4 or 8 transceivers.

Up to 239 Brocade switches are used in a single FC Fabric.

3.1.7.1 QLogic QME2572 Host Bus Adapter

The QLogic® QME2572 is a dual-channel 8Gb/s Fibre Channel host bus adapter (HBA) designed for use in PowerEdge™ M1000e blade servers. Doubling the throughput enables higher levels of server consolidation and reduces data-migration/backup windows. It also improves performance and ensures reduced response time for mission-critical and next generation killer applications. Optimized for virtualization, power, security and management, as well as reliability, availability and serviceability (RAS), the QME2572 delivers 200,000 I/Os per second (IOPS).



3.1.7.2 QLogic QLE2562 Host Bus Adapter

The QLE2562 is a PCI Express, dual port, Fibre Channel (FC) Host Bus Adapter (HBA). The QLE2562 is part of the QLE2500 HBA product family that offers next generation 8 Gb FC technology, meeting the business requirements of the enterprise data center. Features of this HBA includes throughput of 3200 MBps (full-duplex), 200,000 initiator and target I/Os per second (IOPS) per port, and StarPower™



technology-based dynamic and adaptive power management. Benefits include optimizations for virtualization, power, reliability, availability, and serviceability (RAS), and security.

3.2 Servers

3.2.1 PowerEdge R720

The rack server platform for the Dell Wyse Datacenter solution is the best-in-class Dell PowerEdge R720 (12G). This dual socket CPU platform runs the fastest Intel Xeon E5-2600 family of processors, can host up to 768GB RAM, and supports up to 16 2.5" SAS disks. The Dell PowerEdge R720 offers uncompromising performance and scalability in a 2U form factor. For more information please visit: [Link](#)



3.2.2 PowerEdge M620

The blade server platform recommendation for the Dell Wyse Datacenter solution is the PowerEdge M620. This half-height blade server is a feature-rich, dual-processor platform that offers a blend of density, performance, efficiency and scalability. The M620 offers remarkable computational density, scaling up to 24 cores, 2 socket Intel Xeon processors and 24 DIMMs (768GB RAM) of DDR3 memory in an extremely compact half-height blade form factor. For more information please visit: [Link](#)




3.3 Storage

3.3.1 EqualLogic Tier 1 Storage (iSCSI)

3.3.1.1 PS6210XS

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	Tier 1 array for Shared Tier 1 solution model requiring greater per user capacity. (10Gb – iSCSI)




3.3.2 EqualLogic Tier 2 Storage (iSCSI)

The following arrays are used for management VM storage and user data, depending on the scale of the deployment. Please refer to the hardware tables in section 2 or the "Uses" column of each array below. For more information on Dell EqualLogic offerings, please visit: [Link](#)


3.3.2.1 PS4100E

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 (1Gb – iSCSI)
		24TB – 12 x 2TB HDs	
		36TB – 12 x 3TB HDs	




3.3.2.2 PS4110E

Model	Features	Options	Uses
EqualLogic PS4110E	12 drive bays (NL-SAS/ 7200k RPM), dual HA controllers, Snaps/Clones, Async replication, SAN HQ, 10Gb	12TB – 12 x 1TB HDs	Tier 2 array for 1000 users or less in Shared Tier 1 solution model (10Gb – iSCSI)
		24TB – 12 x 2TB HDs	
		36TB – 12 x 3TB HDs	




3.3.2.3 PS6100E

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	




3.3.2.4 PS6210E

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	




3.3.2.5 PS6500E

Model	Features	Options	Uses
EqualLogic PS6500E	48 drive SATA/ NL SAS array, dual HA controllers, Snaps/Clones, Async replication, SAN HQ, 1Gb	48TB – 48 x 1TB SATA	Tier 2 array for Local Tier 1 solution model up to 2500 users each (1Gb – iSCSI)
		96TB – 48 x 2TB SATA	
		144TB – 48 x 3TB NL SAS	




3.3.2.6 PS6510E

Model	Features	Options	Uses
EqualLogic PS6510E	48 drive SATA/ NL SAS array, dual HA controllers, Snaps/Clones, Async replication, SAN HQ, 10Gb	48TB – 48 x 1TB SATA	Tier 2 array for Shared Tier 1 solution model up to 2500 users each (10Gb – iSCSI)
		96TB – 48 x 2TB SATA	
		144TB – 48 x 3TB NL SAS	



3.3.3 EqualLogic Configuration


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



3.3.4 Compellent Storage (FC)

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:

- Fast Track – Dynamic placement of most frequently accessed data blocks on the faster outer tracks of each spinning disk. Lesser active data blocks remain on the inner tracks. Fast track is well-complimented when used in conjunction with Thin Provisioning.
- Data Instant Replay – Provides continuous data protection using snapshots called Replays. Once the base of a volume has been captured, only incremental changes are then captured going forward. This allows for a high number of Replays to be scheduled over short intervals, if desired, to provide maximum protection.



3.3.4.1 Compellent Tier 1

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

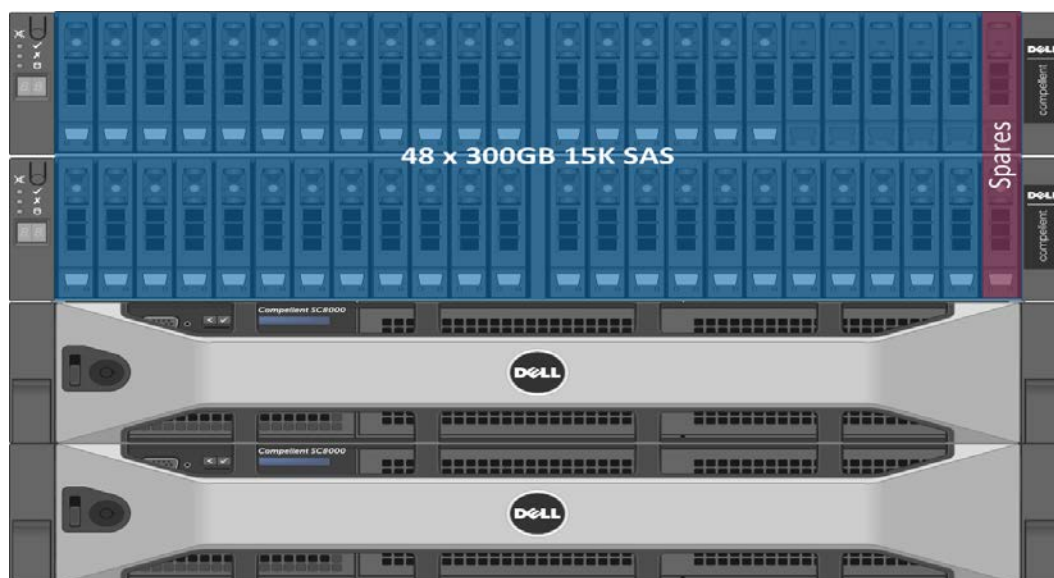
specific workload characteristics in section 7. SSDs are added for use in scenarios where boot storms or provisioning speeds are an issue.

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)	6.3

Tier 1 Scaling Guidance:

Users	Controller Pairs	Disk Shelves	15K SAS Disks	RAW Capacity	Use
500	1	1	22	7TB	T1 + T2
1000	1	2	48	15TB	T1 + T2
2000	1	4	96	29TB	T1 + T2

Example of a 1000 user Tier 1 array:



3.3.4.2 Compellent Tier 2


Compellent Tier 2 storage is completely 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 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)

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

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

Example of a 1000 user Tier 2 array:



For more information on Dell Compellent, please visit: [Link](#)

3.3.5 NAS

3.3.5.1 FS7600

Model	Features	Scaling	Uses
EqualLogic FS7600	Dual active-active controllers, 24GB cache per controller (cache mirroring), SMB & NFS support, AD-integration. Up to 2 FS7600 systems in a NAS cluster (4 controllers). 1Gb iSCSI via 16 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.

3.3.5.2 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).




3.3.5.3 PowerVault NX3300 NAS

Model	Features	Options	Uses
PowerVault NX3300	Cluster-ready NAS built on Microsoft® Windows® Storage Server 2008 R2 Enterprise Edition,	1 or 2 CPUs, 1Gb and 10Gb NICs (configurable).	Scale out NAS for Shared Tier 1 on EqualLogic or Compellent, to provide file share HA (iSCSI).



3.4 Dell Wyse Cloud Clients



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

3.4.1 ThinOS – T10D

T10D sets the standard for thin clients. Providing an exceptional user experience, the T10D 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. Boot up in just seconds and log in securely to almost any network. The T10D delivers a superior Citrix VDI user experience, along with usability and management features found in premium thin clients. The T10D delivers outstanding performance based on its dual core system-on-a-chip (SoC) design, and a built-in media processor delivers smooth multimedia, bi-directional audio and Flash playback. Flexible mounting options let you position the T10D vertically or



horizontally on your desk, on the wall or behind your display. Using about 7-watts of power in full operation, the T10D creates very little heat for a greener, more comfortable working environment. [Link](#)

3.4.2 ThinOS – D10D



Designed for knowledge workers and power users, the new Dell Wyse D10D 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 D10D combines Dell Wyse ThinOS with a dual-core AMD 1.4 GHz processor and a revolutionary unified graphics engine for an outstanding user experience. The D10D 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 D10D offers a unique combination of performance, efficiency, and affordability. For more information, please visit: [Link](#)

3.4.3 Windows Embedded 7 – Z90Q7

The Dell Wyse Z90Q7 is a super high-performance Windows Embedded Standard 7 thin client for virtual desktop environments. Featuring a quad-core AMD processor, and an integrated graphics engine that significantly boost performance; the Z90Q7 achieves exceptional speed and power for the most demanding VDI and embedded Windows applications, rotational 3D graphics, 3D simulation and modeling, unified communications, and multi-screen HD multimedia. Take a unit from box to productivity in minutes. Just select the desired configuration and the Z90Q7 does the rest automatically—no need to reboot. Scale to tens of thousands of endpoints with Dell Wyse WDM software or leverage your existing Microsoft System Center Configuration Manager platform. The Z90Q7 is the thin client for power users who need workstation-class performance on their desktop or within a desktop virtualization environment (x86 or x64). For more information, please visit: [Link](#)

3.4.4 Windows Embedded 8 – Z90Q8

Dell Wyse Z90Q8 is a super high-performance Windows Embedded 8 Standard thin client for virtual desktop environments. Featuring a quad-core AMD processor, the Z90Q8 offers a vibrant Windows 8 experience and achieves exceptional speed and power for the most demanding embedded Windows applications, rich 3D graphics and HD multimedia. And you can scale to tens of thousands of Z90Q8 endpoints with Dell Wyse Device Manager (WDM) software, or leverage your existing Microsoft System Center Configuration Manager platform. With single-touch or multi-touch capable displays, the Z90Q8 adds the ease of an intuitive touch user experience. The Z90Q8 is an ideal thin client for offering a high-performance Windows 8 experience with the most demanding mix of virtual desktop or cloud applications (x86 or x64). For more information please visit: [Link](#)

3.4.5 Suse Linux – Z50D

Designed for power users, the new Dell Wyse Z50D is the highest performing thin client on the market. Highly secure and ultra-powerful, the Z50D combines Dell Wyse-enhanced SUSE Linux Enterprise with a dual-core AMD 1.65 GHz processor and a revolutionary unified engine for an unprecedented user experience. The Z50D eliminates performance constraints for high-end, processing-intensive applications like computer-aided design, multimedia, HD video and 3D modeling. Scalable enterprise-wide management provides simple deployment, patching and updates. Take a unit from box to productivity in minutes with auto configuration. Delivering unmatched processing speed and power, security and display performance, it's no wonder no other thin client can compare. For more information, please visit: [Link](#)



3.4.6 Dell Wyse Zero – Xenith 2



Establishing a new price/performance standard for zero clients for Citrix, the new Dell Wyse Xenith 2 provides an exceptional user experience at a highly affordable price for Citrix XenDesktop and XenApp environments. With zero attack surface, the ultra-secure Xenith 2 offers network-borne viruses and malware zero target for attacks. Xenith 2 boots up in just seconds and delivers exceptional performance for Citrix XenDesktop and XenApp users while offering usability and management features found in premium Dell Wyse cloud client devices. Xenith 2 delivers outstanding performance based on its system-on-chip (SoC) design optimized with its Dell Wyse zero architecture and a built-in media processor delivers smooth multimedia, bi-directional audio and Flash playback. Flexible mounting options let you position Xenith 2 vertically or horizontally. Using about 7 Watts of power in full operation, the Xenith 2 creates very little heat for a greener working environment. For more information, please visit: [Link](#)

3.4.7 Dell Wyse Zero – Xenith Pro 2

Dell Wyse Xenith Pro 2 is the next-generation zero client for Citrix HDX and Citrix XenDesktop, delivering ultimate performance, security and simplicity. With a powerful dual core AMD G-series CPU, Xenith Pro 2 is faster than competing devices. This additional computing horsepower allows dazzling HD multimedia delivery without overtaxing your server or network. Scalable enterprise-wide management provides simple deployment, patching and updates—your Citrix XenDesktop server configures it out-of-the-box to your preferences for plug-and-play speed and ease of use. Virus and malware immune, the Xenith Pro 2 draws under 9 watts of power in full operation—that's less than any PC on the planet. For more information please visit: [Link](#)



3.4.8 Dell Wyse Cloud Connect



Designed to promote bring-your-own-device (BYOD) environments, Dell Wyse Cloud Connect allows you to securely access and share work and personal files, presentations, applications and other content from your business or your home. Managed through Dell Wyse Cloud Client Manager software-as-a-service (SaaS), IT managers can ensure that each Cloud Connect device is used by the appropriate person with the right permissions and access to the appropriate apps and content based on role, department and location. Slightly larger than a USB memory stick, Cloud Connect is an ultra-compact multimedia-capable device. Simply plug it into any available

Mobile High-Definition Link (MHL) / HDMI port on a TV or monitor, attach a Bluetooth keyboard and mouse, and you're off and running. Easy to slip into your pocket or bag, it enables an HD-quality window to the cloud, great for meetings and presentations while on business travel, or for cruising the internet and checking email after a day of work. For more information, please visit: [Link](#)

3.4.9 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](#)

3.4.10 Dell Chromebook 11

The lightweight, easy-to-use Dell Chromebook 11 helps turn education into exploration - without the worries of safety or security. Priced to make 1:1 computing affordable today, Chromebook 11 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 4th Gen Intel® processor provides up to 10 hours of power. Encourage creativity with the Chromebook 11 and its multimedia features that include an 11.6" screen, stereo sound and webcam.

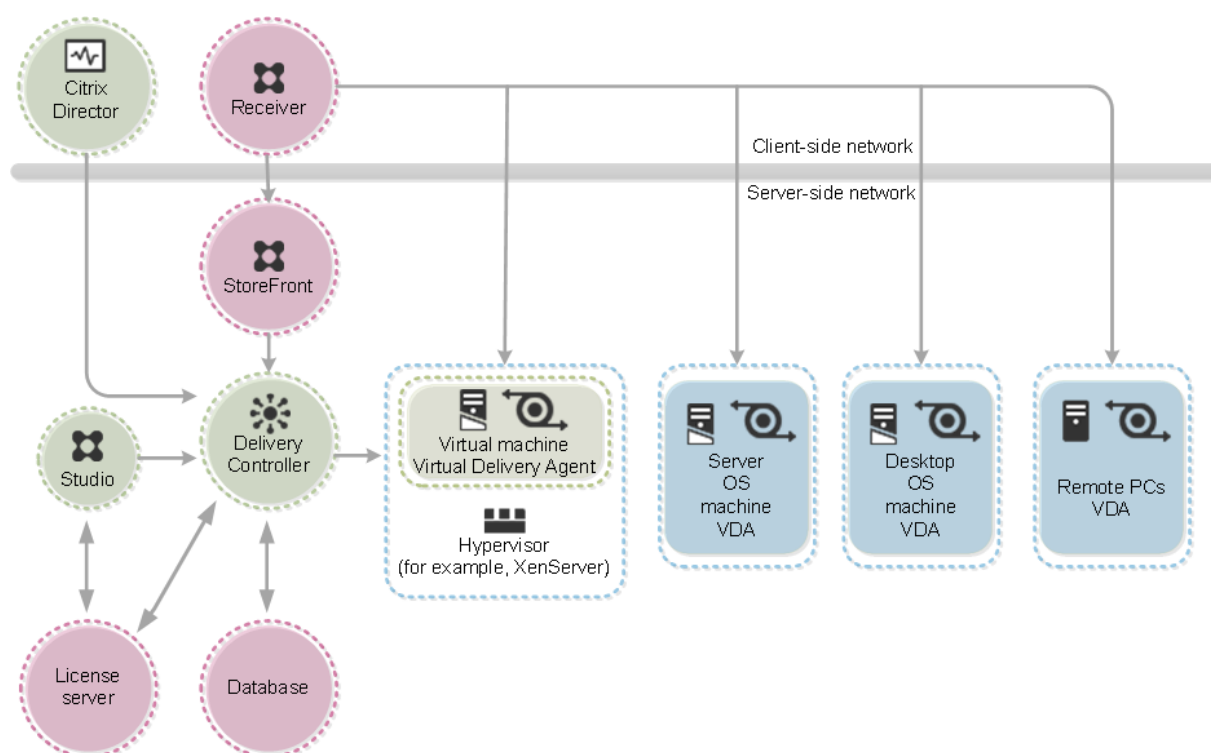


4 Software Components

4.1 Citrix XenDesktop

The solution is based on Citrix XenDesktop 7.1 which provides a complete end-to-end solution delivering Microsoft Windows virtual desktops or server-based hosted shared sessions to users on a wide variety of endpoint devices. Virtual desktops are dynamically assembled on demand, providing users with pristine, yet personalized, desktops each time they log on.

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



The core XenDesktop components include:

- **Studio**

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

- **Director**

Director is a web-based tool that enables IT support teams to monitor an environment, troubleshoot issues before they become system-critical, and perform support tasks for end users. You can also view and interact with a user's sessions using Microsoft Remote Assistance.

- **Receiver**

Installed on user devices, Citrix Receiver provides users with quick, secure, self-service access to documents, applications, and desktops from any of the user's devices including smartphones, tablets, and PCs. Receiver provides on-demand access to Windows, Web, and Software as a Service (SaaS) applications.

- **Delivery Controller (DC)**

Installed on servers in the data center, the controller authenticates users, manages the assembly of users' virtual desktop environments, and brokers connections between users and their virtual desktops.

- **StoreFront**

StoreFront authenticates users to sites hosting resources and manages stores of desktops and applications that user's access.

- **License Server**

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

- **Machine Creation Services (MCS)**

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

- **Provisioning Services (PVS)**

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

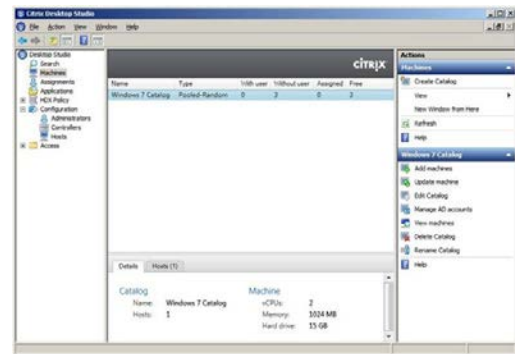
- **Virtual Delivery Agent (VDA)**

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

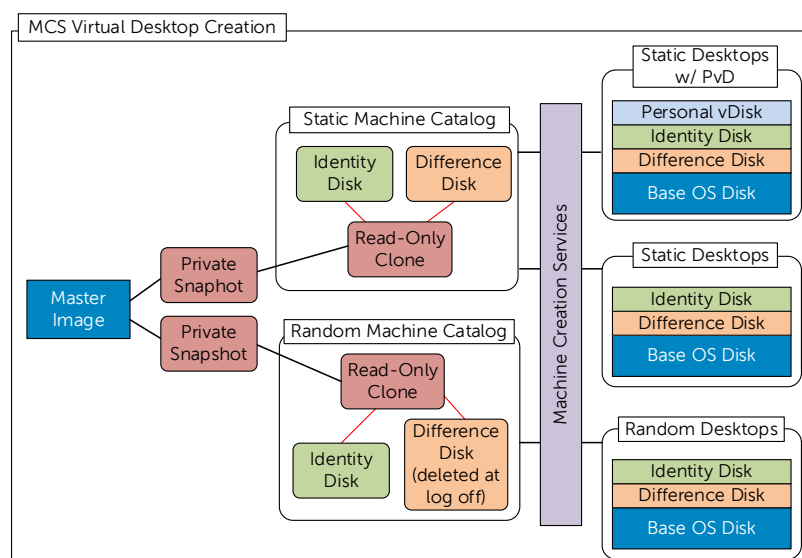
4.1.1 Machine Creation Services (MCS)

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

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



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



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

4.1.2 Provisioning Services (PVS)

PVS is an alternative method of image provisioning which uses streaming to share a single base vDisk image instead of copying images to VMs. PVS are used to deliver shared vDisk images to physical or virtual machines. Another potential use is the serial provisioning of XenApp to enable scale-out hosted shared desktop



infrastructure. Provisioning Services enables real-time streamed provisioning and re-provisioning which enable administrators to completely eliminate the need to manage and patch individual systems.

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

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

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

4.1.2.1 PVS Write Cache

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

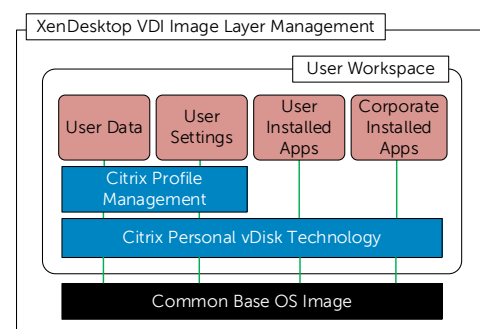
While there are 4 possible locations for storage of the write cache in PVS, the Dell Wyse Datacenter solution recommends placement of the PVS write cache in the target compute host's RAM with overflow enabled. We recommend using a cache size of 512MB for virtual desktops and 21GB for XenApp VMs delivered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

4.1.3 Personal vDisk

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

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

Personal vDisk provides the following benefits to XenDesktop:



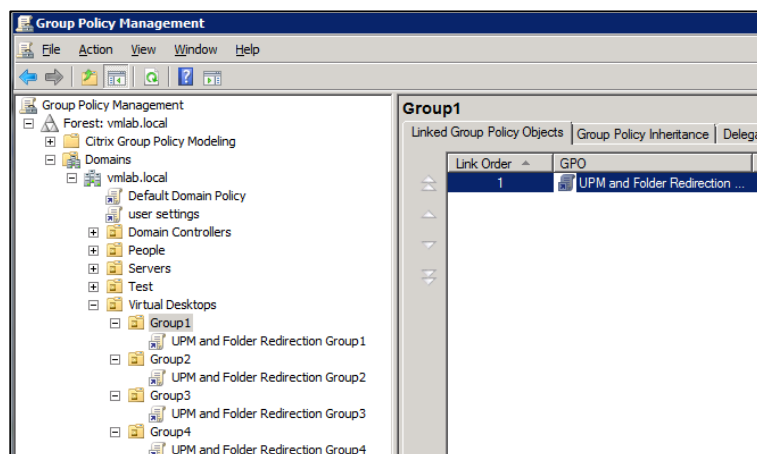
- Persistent personalization of user profiles, settings and data
- Enables deployment and management of user installed and entitlement based applications
- Fully compatible with Microsoft SCCM and App-V
- 100% persistence with VDI pooled Storage management
- Near Zero management overhead

4.1.4 Citrix Profile Manager

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

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

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



4.1.5 XenApp


Hosted desktops and applications, delivered via Microsoft XenApp, provide a sensible and low-cost entry into VDI by consolidating many basic task user sessions within a single Windows Server instance (formerly Terminal Services with XenApp). This technology can also be used to augment a XenDesktop deployment by delivering virtualized applications to user desktops via XenApp or App-V. XenDesktop 7 brings greater simplicity to the server hosted model simply by leveraging the Citrix Virtual Delivery Agent (VDA) installed within the Windows Server OS. The management infrastructure between desktop OS (desktop VMs) and server OS (XenApp) deployments is identical.

Hosted applications give users a customized application set with a “locally-installed” experience even though the applications are centrally installed and managed on XenApp servers. This can dramatically simplify the XenDesktop environment by leveraging a widely shared virtual desktop image, while at the same time extending the scalability of XenDesktop by alleviating resource load from the desktop compute servers. This two-Tiered approach to desktop and application delivery brings management simplification, a much quicker return on investment and the absolute best end-user experience.

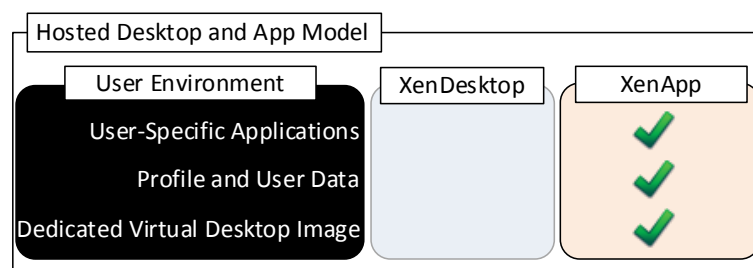
Benefits of hosted desktop sessions and applications:

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

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



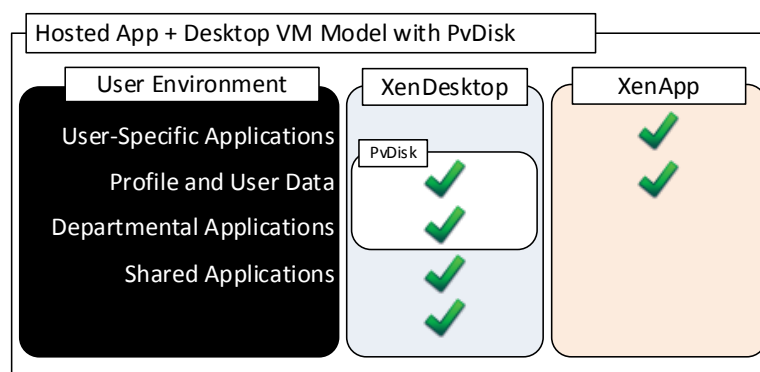
Alternatively, XenApp provides a platform for delivering a Windows server-based sessions to users who may not need a full desktop VM. Hosted desktops increase infrastructure resource utilization while reducing complexity as all applications and sessions are centrally managed.



4.1.5.1 XenDesktop with XenApp and Personal vDisk Integration

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

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



4.1.5.2 XenApp Integration into Dell Wyse Datacenter Architecture

The XenApp servers can exist as physical or virtualized instances of Windows Server 2012 R2. A minimum of one (1), up to a maximum of eight (8) virtual servers are installed per physical compute host. Since XenApp instances are easily added to an existing XenDesktop stack, the only additional components required are:

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

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

4.1.5.3 PVS Integration with XenApp

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

- PVS streaming is delivered via UDP, yet the application has built-in mechanisms to provide flow control, and retransmission as necessary.

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

4.2 VDI Hypervisor Platforms

4.2.1 VMware vSphere 5

VMware vSphere 5 is a virtualization platform used for building VDI and cloud infrastructures. vSphere 5 represents a migration from the ESX architecture to the ESXi architecture.

VMware vSphere 5 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 client and the vSphere web client.

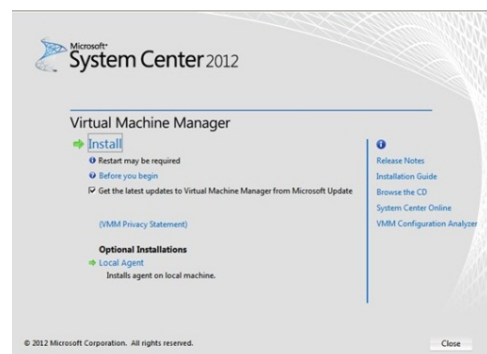


Throughout the Dell Wyse Datacenter solution, all VMware and Microsoft best practices and prerequisites for core services are adhered to (NTP, DNS, Active Directory, etc.). The vCenter 5 VM used in the solution is a single Windows Server 2008 R2 VM or vCenter 5 virtual appliance, residing on a host in the management Tier. SQL server is a core component of the Windows version of vCenter and is hosted on another VM also residing in the management Tier. It is recommended that all additional XenDesktop components be installed in a distributed architecture, one role per server VM.

4.2.2 Microsoft Windows Server 2012 Hyper-V

Windows Server 2012 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 processors 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. The primary management platform within a Hyper-V based XenDesktop virtualization environment is Microsoft Systems Center Virtual Machine Manager SP1 (SCVMM).

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

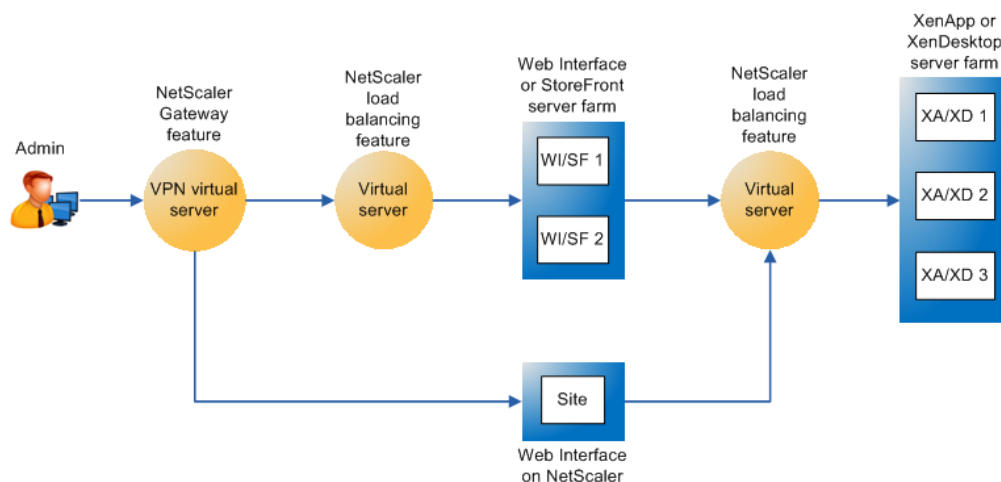
4.3 Citrix NetScaler

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

- Proven application acceleration such as [HTTP compression](#) and [caching](#)
- High application availability through advanced L4-7 [load balancer](#)
- Application security with an integrated [AppFirewall](#)
- Server offloading to significantly reduce costs and consolidate servers

Where Does a Citrix NetScaler Fit in the Network?

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



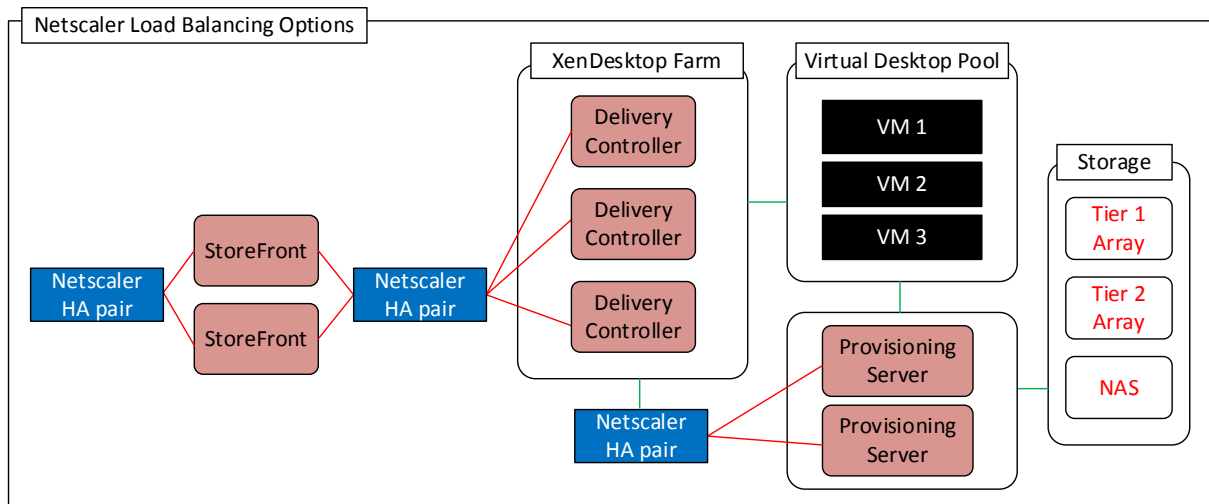
Global Server Load Balancing

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

XenDesktop HA with Netscaler White Paper: [Link](#)

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

- StoreFront Servers
- Licensing Server
- XenDesktop XML Service
- XenDesktop Desktop Director
- Provisioning Services TFTP Service
-



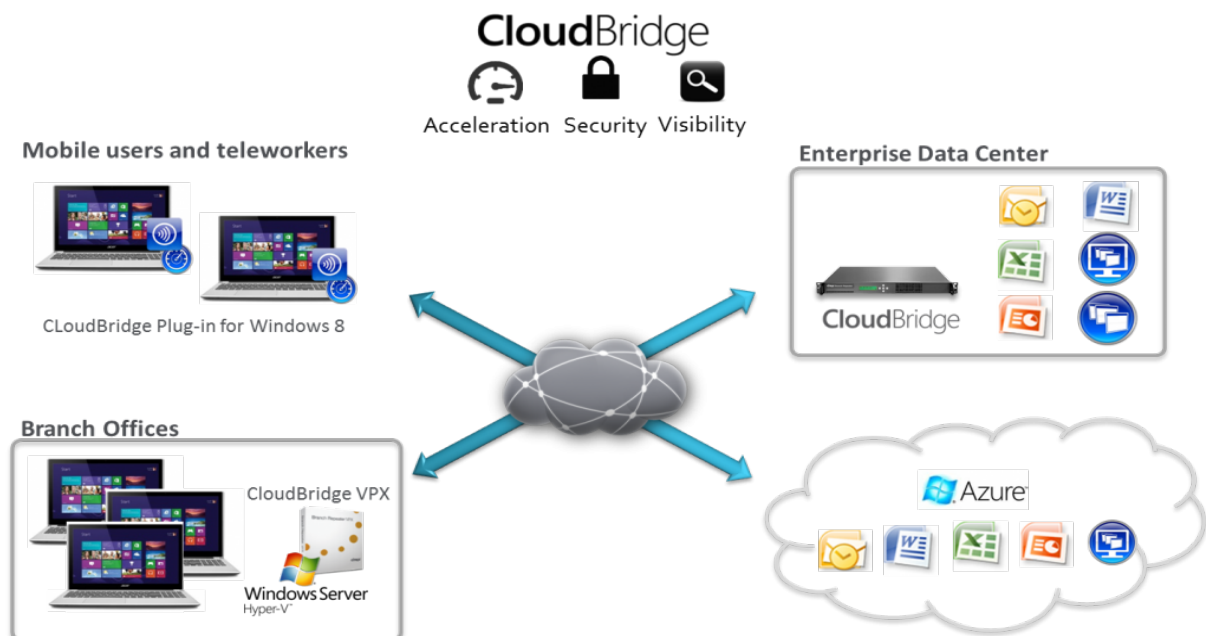
Citrix NetScaler is added to the Dell Wyse Datacenter mgmt. stack at any time and runs on the existing server infrastructure.

4.4 Citrix CloudBridge

Citrix CloudBridge provides a unified platform that connects and accelerates applications, and optimizes bandwidth utilization across public cloud and private networks. The only WAN optimization solution with integrated, secure, transparent cloud connectivity, CloudBridge allows enterprises to augment their data center with the infinite capacity and elastic efficiency provided by public cloud providers. CloudBridge delivers superior application performance and end-user experiences through a broad base of features, including:

- Market-leading enhancements for the Citrix XenDesktop user experience including HDX WAN optimization
- Secure, optimized networking between clouds
- Compression, de-duplication and protocol acceleration
- Acceleration of traditional enterprise applications
- Sophisticated traffic management controls and reporting
- Faster storage replication times and reduced bandwidth demands
- Integrated video delivery optimization to support increasing video delivery to branch offices
- Deliver a faster experience for all users

CloudBridge is ICA aware and enables IT organizations to accelerate, control and optimize all services – desktops, applications, multimedia and more – for corporate office, branch office and mobile users while dramatically reducing costs. With CloudBridge, branch office users experience a better desktop experience with faster printing, file downloads, video streaming and application start-up times.



For more information please visit: [Link](#)

5 Solution Architecture for XenDesktop 7

5.1 Compute Server Infrastructure

5.1.1 Local Tier 1 Rack

In the Local Tier 1 model, VDI sessions execute on 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. Because of this, the Compute and Management servers are configured with different add-on NICs to support their pertinent network Fabric connection requirements. Refer to section 2.4.3.2 for cabling implications. The Management server host has reduced RAM and CPU and does not require local disk to host the management VMs. Any of the compute options below can also be used for XenApp while optionally reducing the amount of RAM.

Local Tier 1 Compute Host – PowerEdge R720		
2 x Intel Xeon E5-2690v2 Processor (3Ghz)	OR	2 x Intel Xeon E5-2690v2 Processor (3Ghz)
256GB Memory (16 x 16GB DIMMs @ 1600Mhz)		256GB Memory (16 x 16GB DIMMs @ 1600Mhz)
VMware vSphere on internal 2GB Dual SD		Microsoft Hyper-V on 12 x 300GB 15K SAS disks
10 x 300GB SAS 6Gbps 15k Disks (VDI)		PERC H710 Integrated RAID Controller – RAID10
PERC H710 Integrated RAID Controller – RAID10		Broadcom 5720 1Gb QP NDC (LAN)
Broadcom 5720 1Gb QP NDC (LAN)		Broadcom 5720 1Gb DP NIC (LAN)
Broadcom 5720 1Gb DP NIC (LAN)		iDRAC7 Enterprise w/ vFlash, 8GB SD
iDRAC7 Enterprise w/ vFlash, 8GB SD		2 x 750W PSUs
2 x 750W PSUs		

Local Tier 1 Management Host – PowerEdge R720	
2 x Intel Xeon E5-2670v2 Processor (2.5Ghz)	2 x Intel Xeon E5-2670v2 Processor (2.5Ghz)
96GB Memory (6 x 16GB DIMMs @ 1600Mhz)	96GB Memory (6 x 16GB DIMMs @ 1600Mhz)
VMware vSphere on internal 2GB Dual SD	Microsoft Hyper-V on 2 x 300GB 15K SAS disks
Broadcom 5720 1Gb QP NDC (LAN/iSCSI)	Broadcom 5720 1Gb QP NDC (LAN/iSCSI)

Broadcom 5719 1Gb QP NIC (LAN/iSCSI)	Broadcom 5719 1Gb QP NIC (LAN/iSCSI)
iDRAC7 Enterprise w/ vFlash, 8GB SD	iDRAC7 Enterprise w/ vFlash, 8GB SD
2 x 750W PSUs	2 x 750W PSUs

5.1.2 Local Tier 1 Blade

In the Local Tier 1 model for blades, VDI sessions execute on local high-performance SSDs on each compute host. vSphere is the supported hypervisor in this solution due to its ability to run from integrated SD freeing the pair of SSDs for VDI execution only. In this model, shared storage is not required for Tier 2 unless management host-level HA is required. All management and desktop VMs is hosted locally on their respective blades. The recommended provisioning method is MCS with non-persistent desktops. Refer to Section 2.2.3 for solution cabling implications.

Local Tier 1 Compute Host – PowerEdge M620
2 x Intel Xeon E5-2690v2 Processor (3GHz)
256GB Memory (16 x 16GB DIMMs @ 1600Mhz)
VMware vSphere on 2 x 2GB internal SD
Broadcom 57810-k 10Gb DP KR NDC (LAN)
iDRAC7 Enterprise w/ vFlash, 8GB SD

Local Tier 1 Management Host – PowerEdge M620
2 x Intel Xeon E5-2670v2 Processor (2.5GHz)
96GB Memory (6 x 16GB DIMMs @ 1600Mhz)
VMware vSphere on 2 x 2GB internal SD
Broadcom 57810-k 10Gb DP KR NDC (iSCSI-HA)
1 x Intel i350 1Gb QP mezzanine (LAN)
iDRAC7 Enterprise w/ vFlash, 8GB SD

5.1.3 Shared Tier 1 Rack

In the Shared Tier 1 model, VDI 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 split QP NDC: 2 x 10Gb ports for iSCSI, 2 x 1Gb ports for LAN. Two additional DP NICs (2 x 1Gb + 2 x 10Gb) provide slot and connection-level redundancy for both network Fabrics. 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.

5.1.3.1 iSCSI

Shared Tier 1 Compute Host – PowerEdge R720		
2 x Intel Xeon E5-2690v2 Processor (3GHz)	OR	2 x Intel Xeon E5-2690v2 Processor (3GHz)
256GB Memory (16 x 16GB DIMMs @ 1600Mhz)		256GB Memory (16 x 16GB DIMMs @ 1600Mhz)
VMware vSphere on 2 x 2GB internal SD		Microsoft Hyper-V on 2 x 300GB 15K SAS disks
Broadcom 57800 2 x 10Gb SFP+ + 2 x 1Gb NDC		Broadcom 57800 2 x 10Gb SFP+ + 2 x 1Gb NDC
1 x Broadcom 5720 1Gb DP NIC (LAN)		1 x Broadcom 5720 1Gb DP NIC (LAN)
1 x Intel X520 2 x 10Gb SFP+ DP NIC (iSCSI)		1 x Intel X520 2 x 10Gb SFP+ DP NIC (iSCSI)
iDRAC7 Enterprise w/ vFlash, 8GB SD		iDRAC7 Enterprise w/ vFlash, 8GB SD
2 x 750W PSUs		2 x 750W PSUs

Shared Tier 1 Management Host – PowerEdge R720		
2 x Intel Xeon E5-2670v2 Processor (2.5GHz)	OR	2 x Intel Xeon E5-2670v2 Processor (2.5GHz)
96GB Memory (6 x 16GB DIMMs @ 1600Mhz)		96GB Memory (6 x 16GB DIMMs @ 1600Mhz)
VMware vSphere on 2 x 2GB internal SD		Microsoft Hyper-V on 2 x 300GB 15K SAS disks
Broadcom 57800 2 x 10Gb SFP+ + 2 x 1Gb NDC		Broadcom 57800 2 x 10Gb SFP+ + 2 x 1Gb NDC
1 x Broadcom 5720 1Gb DP NIC (LAN)		1 x Broadcom 5720 1Gb DP NIC (LAN)
1 x Intel X520 2 x 10Gb SFP+ DP NIC (iSCSI)		1 x Intel X520 2 x 10Gb SFP+ DP NIC (iSCSI)
iDRAC7 Enterprise w/ vFlash, 8GB SD		iDRAC7 Enterprise w/ vFlash, 8GB SD
2 x 750W PSUs		2 x 750W PSUs

5.1.3.2 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 R720	
2 x Intel Xeon E5-2690v2 Processor (3GHz)	2 x Intel Xeon E5-2690v2 Processor (3GHz)
256GB Memory (16 x 16GB DIMMs @ 1600Mhz)	256GB Memory (16 x 16GB DIMMs @ 1600Mhz)
VMware vSphere on 2 x 2GB internal SD	Microsoft Hyper-V on 2 x 300GB 15K SAS disks
Broadcom 5720 1Gb QP NDC (LAN)	Broadcom 5720 1Gb QP NDC (LAN)
Broadcom 5720 1Gb DP NIC (LAN)	Broadcom 5720 1Gb DP NIC (LAN)
2 x QLogic 2562 8Gb DP FC HBA	2 x QLogic 2562 8Gb DP FC HBA
iDRAC7 Enterprise w/ vFlash, 8GB SD	iDRAC7 Enterprise w/ vFlash, 8GB SD
2 x 750W PSUs	2 x 750W PSUs

Shared Tier 1 Management Host – PowerEdge R720	
2 x Intel Xeon E5-2670v2 Processor (2.5GHz)	2 x Intel Xeon E5-2670v2 Processor (2.5GHz)
96GB Memory (6 x 16GB DIMMs @ 1600Mhz)	96GB Memory (6 x 16GB DIMMs @ 1600Mhz)
VMware vSphere on 2 x 2GB internal SD	Microsoft Hyper-V on 2 x 300GB 15K SAS disks
Broadcom 5720 1Gb QP NDC (LAN)	Broadcom 5720 1Gb QP NDC (LAN)
Broadcom 5720 1Gb DP NIC (LAN)	Broadcom 5720 1Gb DP NIC (LAN)
2 x QLogic 2562 8Gb DP FC HBA	2 x QLogic 2562 8Gb DP FC HBA
iDRAC7 Enterprise w/ vFlash, 8GB SD	iDRAC7 Enterprise w/ vFlash, 8GB SD
2 x 750W PSUs	2 x 750W PSUs

In the above configurations, the R720-based Dell Wyse Datacenter Solution can support the following maximum user counts per server (PVS or MCS).

Local/ Shared Tier 1 Rack Densities		
Workload	vSphere	Hyper-V
Standard (Win8)	165	170
Enhanced (Win8)	110	120
Professional (Win8)	90	100

5.1.4 Shared Tier 1 Blade

The Dell M1000e Blade Chassis combined with the M620 blade server is the platform of choice for a high-density data center configuration. The M620 is a feature-rich, dual-processor, half-height blade server which offers a blend of density, performance, efficiency and scalability. The M620 offers remarkable computational density, scaling up to 16 cores, 2 socket Intel Xeon processors and 24 DIMMs (768GB RAM) of DDR3 memory in an extremely compact half-height blade form factor.

5.1.4.1 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 1Gb LAN.

Shared Tier 1 Compute Host – PowerEdge M620		
2 x Intel Xeon E5-2690v2 Processor (3GHz)	OR	2 x Intel Xeon E5-2690v2 Processor (3GHz)
256GB Memory (16 x 16GB DIMMs @ 1600Mhz)		256GB Memory (16 x 16GB DIMMs @ 1600Mhz)
VMware vSphere on 2 x 2GB internal SD		Microsoft Hyper-V on 2 x 300GB 15K SAS disks
Broadcom 57810 -k 10Gb DP KR NDC (iSCSI)		Broadcom 57810 -k 10Gb DP KR NDC (iSCSI)
1 x Intel i350 1Gb QP mezz (LAN)		1 x Intel i350 1Gb QP mezz (LAN)
iDRAC7 Enterprise w/ vFlash, 8GB SD		iDRAC7 Enterprise w/ vFlash, 8GB SD

Shared Tier 1 Management Host – PowerEdge M620	
2 x Intel Xeon E5-2670v2 Processor (2.5GHz)	2 x Intel Xeon E5-2670v2 Processor (2.5GHz)
96GB Memory (6 x 16GB DIMMs @ 1600Mhz)	96GB Memory (6 x 16GB DIMMs @ 1600Mhz)
VMware vSphere on 2 x 2GB internal SD	Microsoft Hyper-V on 2 x 300GB 15K SAS disks
Broadcom 57810-k 10Gb DP KR NDC (iSCSI)	Broadcom 57810-k 10Gb DP KR NDC (iSCSI)
1 x Intel i350 1Gb QP mezz (LAN)	1 x Intel i350 1Gb QP mezz (LAN)
iDRAC7 Enterprise w/ vFlash, 8GB SD	iDRAC7 Enterprise w/ vFlash, 8GB SD

5.1.4.2 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 M620	
2 x Intel Xeon E5-2690v2 Processor (3GHz)	2 x Intel Xeon E5-2690v2 Processor (3GHz)
256GB Memory (16 x 16GB DIMMs @ 1600Mhz)	256GB Memory (16 x 16GB DIMMs @ 1600Mhz)
VMware vSphere on 2 x 2GB internal SD	Microsoft Hyper-V on 2 x 300GB 15K SAS disks
Broadcom 57810-k 10Gb DP KR NDC (LAN)	Broadcom 57810-k 10Gb DP KR NDC (LAN)
1 x QLogic QME2572 8Gb FC mezz (FC)	1 x QLogic QME2572 8Gb FC mezz (FC)
iDRAC7 Enterprise w/ vFlash, 8GB SD	iDRAC7 Enterprise w/ vFlash, 8GB SD

Shared Tier 1 Management Host – PowerEdge M620	
2 x Intel Xeon E5-2670v2 Processor (2.5GHz)	2 x Intel Xeon E5-2670v2 Processor (2.5GHz)
96GB Memory (6 x 16GB DIMMs @ 1600Mhz)	96GB Memory (6 x 16GB DIMMs @ 1600Mhz)
VMware vSphere on 2 x 2GB internal SD	Microsoft Hyper-V on 2 x 300GB 15K SAS disks
Broadcom 57810-k 10Gb DP KR NDC (LAN)	Broadcom 57810-k 10Gb DP KR NDC (LAN)
1 x QLogic QME2572 8Gb FC mezz (FC)	1 x QLogic QME2572 8Gb FC mezz (FC)

iDRAC7 Enterprise w/ vFlash, 8GB SD

iDRAC7 Enterprise w/ vFlash, 8GB SD

In the above configuration, the M620-based Dell Wyse Datacenter Solutions can support the following maximum single server user densities (PVS or MCS):

Shared Tier 1 Blade Densities		
Workload	vSphere	Hyper-V
Basic (Win8)	165	170
Standard (Win8)	110	120
Premium (Win8)	90	100

5.1.5 XenApp Compute Host

The XenApp Virtual Server compute hosts share the same PowerEdge R720/ M620 hardware configuration as the XenDesktop compute hosts for both VMware vSphere and Microsoft Hyper-V with slightly reduced RAM. If PVS the cache in RAM feature is to be used, to achieve maximum density, we recommend using 256GB.

XenApp Virtual Server Compute Host – PowerEdge R720		
2 x Intel Xeon E5-2690v2 Processor (3GHz)	OR	2 x Intel Xeon E5-2690v2 Processor (3GHz)
128GB Memory (8 x 16GB DIMMs @ 1600Mhz)		128GB Memory (8 x 16GB DIMMs @ 1600Mhz)
VMware vSphere on 2 x 2GB internal SD		Microsoft Hyper-V on 12 x 300GB 15K SAS disks
10 x 300GB SAS 6Gbps 15k Disks (VDI)		PERC H710 Integrated RAID Controller – RAID10
PERC H710 Integrated RAID Controller – RAID10		Broadcom 5720 1Gb QP NDC (LAN)
Broadcom 5720 1Gb QP NDC (LAN)		Broadcom 5720 1Gb DP NIC (LAN)
Broadcom 5720 1Gb DP NIC (LAN)		iDRAC7 Enterprise w/ vFlash, 8GB SD
iDRAC7 Enterprise w/ vFlash, 8GB SD		2 x 750W PSUs
2 x 750W PSUs		

5.2 Management Server Infrastructure

The Management role requirements for the base solution are summarized below. Use data disks for role-specific application files such as data, logs and IIS web files in the Management volume.

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

Role	vCPU	RAM (GB)	NIC	OS + Data vDisk (GB)	Tier 2 Volume (GB)
DDC	4	8	1	40 + 5	-
Storefront	1	2	1	40 + 5	-
Citrix License Server	1	2	1	40 + 5	-
vCenter or SCVMM	2	8	1	40 + 5	50 (VMDK/ VHDX)
Provisioning Server	4	24	1	40 + 5	100 (VMDK/ VHDX)
SQL Server	2	8	1	40 + 5	200 (VMDK/ VHDX)
File Server	1	4	1	40 + 5	2048 (RDM/ PTD)
TOTALS	15	56	7	315	2398

5.2.1 SQL Databases

The Citrix and VMware databases are hosted by a single dedicated SQL 2008 R2 Server VM in the Management layer. Use caution during database setup to ensure that SQL data, logs, and TempDB are properly separated onto their respective volumes. Create all Databases that are required for:

- Citrix XenDesktop
- PVS
- vCenter or SCVMM

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

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

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

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

5.2.2 DNS




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

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

5.2.2.1 DNS for SQL

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

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

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

5.3 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.

5.3.1 Local Tier 1 Storage

Choosing the local Tier 1 storage option means that the virtualization 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 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.

5.3.2 Shared Tier 1 Storage

Choosing the shared Tier 1 option means that the virtualization compute hosts are deployed in a diskless mode and all leverage shared storage hosted on a high performance Dell storage array. In this model, shared storage is leveraged for Tier 1 used for VDI execution and write cache. Based on the heavy performance requirements of Tier 1 VDI execution, it is recommended to use separate arrays for Tier 1 and Tier 2. 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

5.3.3 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 and 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
SQL Witness	1	Tier 2	SQL (optional)	VMFS/ NTFS
Templates/ ISO	200	Tier 2	ISO storage (optional)	VMFS/ NTFS

5.3.4 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.

5.3.5 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

5.3.5.1 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.



5.4 Virtual Networking


5.4.1 Local Tier 1 – Rack – 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 sessions locally. Since the Management VMs is hosted on shared storage, they can take advantage of Live Migration. The following outlines the VLAN requirements for the Compute and Management hosts in this solution model:

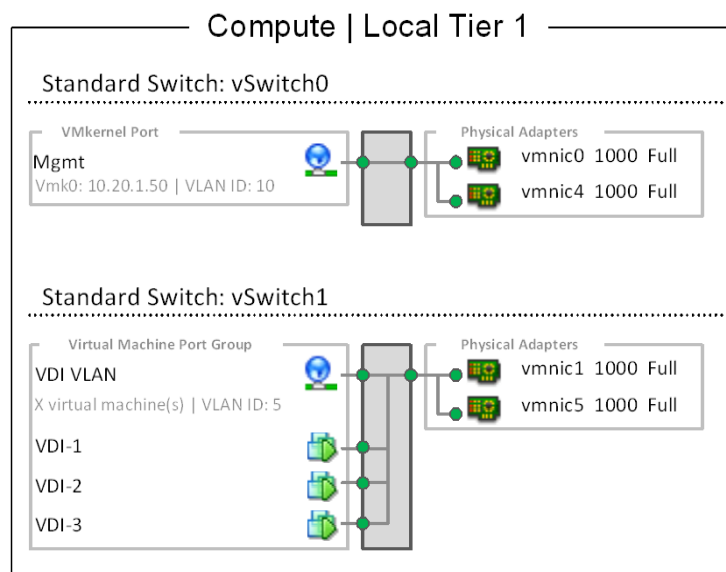
- Compute hosts (Local Tier 1)
 - Management VLAN: Configured for hypervisor infrastructure traffic – L3 routed via core switch
 - VDI VLAN: Configured for VDI session traffic – L3 routed via core switch
- Management hosts (Local Tier 1)
 - Management VLAN: Configured for hypervisor Management traffic – L3 routed via core switch
 - Live Migration VLAN: Configured for Live Migration traffic – L2 switched only, trunked from Core (HA only)
 - 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 iDRAC VLAN are configured for all hardware management traffic – L3 routed via core switch

Following best practices, LAN and block storage traffic is separated in solutions >1000 users. This traffic are combined within a single switch in smaller stacks to minimize buy-in costs, however, VLANs are required for each traffic type to enable traffic separation. Each Local Tier 1 Compute host will have a quad port NDC as well as a 1Gb dual port NIC. Configure the LAN traffic from the server to the ToR switch as a LAG.

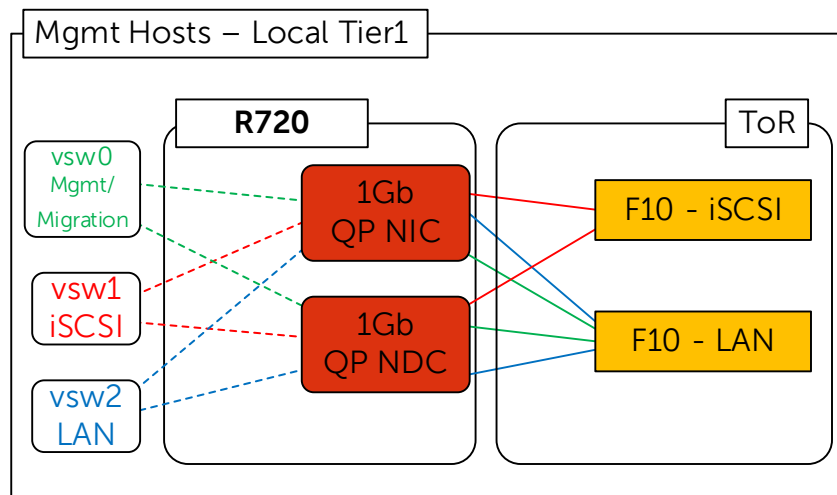
5.4.1.1 vSphere



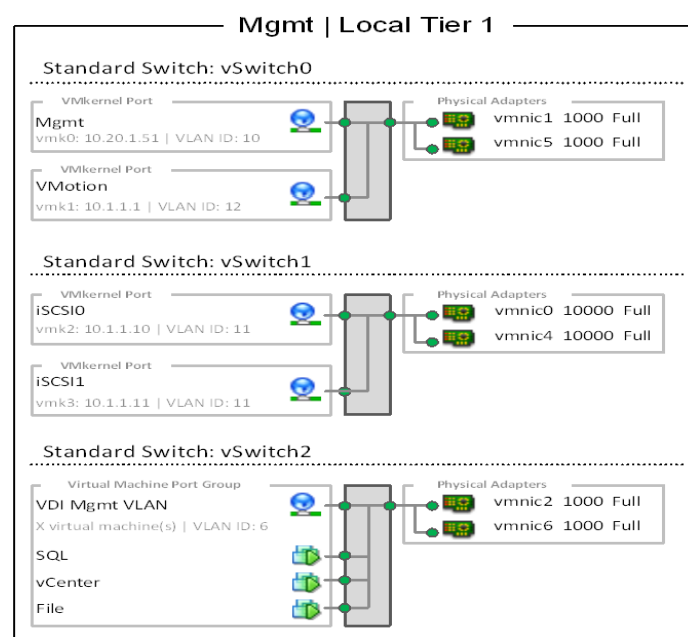
The Compute host will require two vSwitches, one for VDI LAN traffic, and another for the ESXi Management. Configure both vSwitches so that each is physically connected to both the onboard NIC as well as the add-on NIC. Set all NICs and switch ports to auto negotiate. vSphere vSwitch configuration details are shown in the figure below.



The Management hosts have a slightly different configuration since they also access iSCSI storage. The add-on NIC for the Management hosts is a 1Gb quad port NIC. Three ports of both the NDC and add-on NIC are used for the required connections. iSCSI traffic is isolated onto its own vSwitch with redundant ports and connections from all three vSwitches. Connections pass through both the NDC and add-on NIC per the diagram below. LAN traffic from the server is configured to the ToR switch as a LAG.

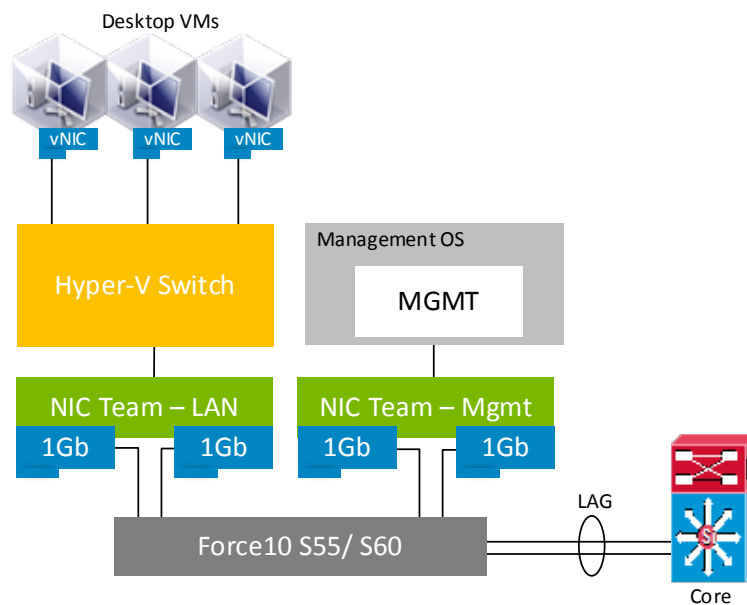


As shown in the diagram below, vSwitch0 carries traffic for both Management and Live Migration which is VLAN-tagged so that either NIC can serve traffic for either VLAN. The Management VLAN is L3 routable while the Live Migration VLAN is L2 non-routable.

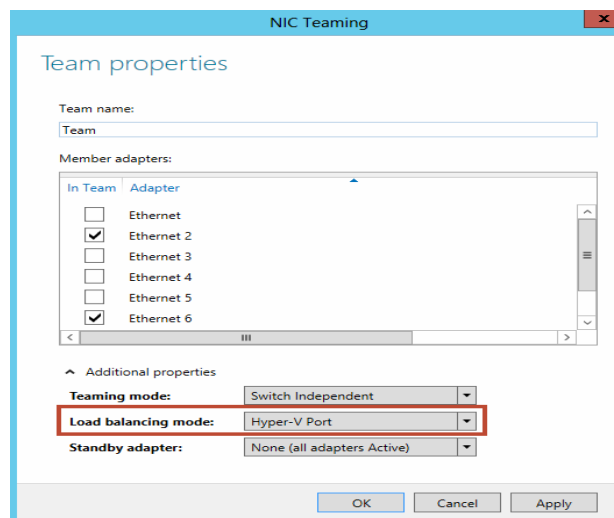


5.4.1.2 Hyper-V

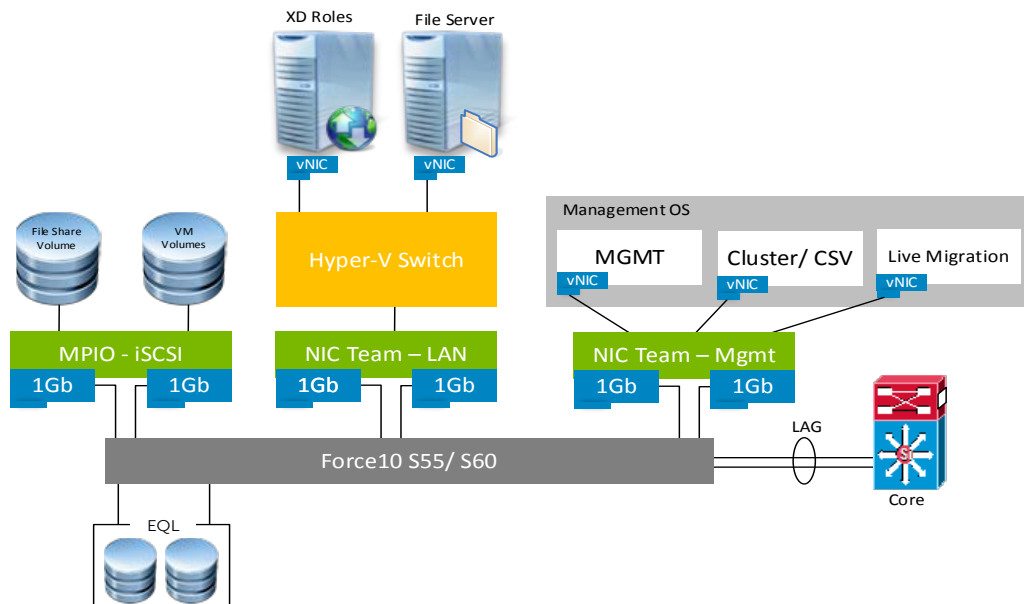
The Hyper-V configuration, while identical in core requirements and hardware, is executed differently due to how Hyper-V and Windows Server 2013 implement networking and virtual switches. As shown in the diagram below, native Windows Server 2012 NIC Teaming is utilized to load balance and provide resiliency for network connections. For the compute host in this scenario, two NIC teams are configured, one for the Hyper-V switch and one for the management OS. The desktop VMs connect to the Hyper-V switch, while the management OS connects directly to the NIC team dedicated to the Mgmt network. The desktop VMs connect to the Hyper-V switch, while the management OS connects directly to the NIC team dedicated to the Mgmt network.



The NIC team for the Hyper-V switch is configured as switch independent, and the Hyper-V port for the load balancing mode with all adapters set to active. This team is used exclusively by Hyper-V.



The management hosts in this scenario connect to shared storage and as a result have a slightly different configuration to support that requirement. Server 2012 supports native MPIO functions but we recommend using the Dell EqualLogic HIT Kit for MPIO to the iSCSI arrays. The rest of the configuration is basically the same as the compute host with two NIC teams, one for Hyper-V and the other for the management OS. If the management host is clustered or utilizes live migration, additional vNICs (Team Interfaces) with dedicated VLANs is configured on the Mgmt NIC team to support these functions.



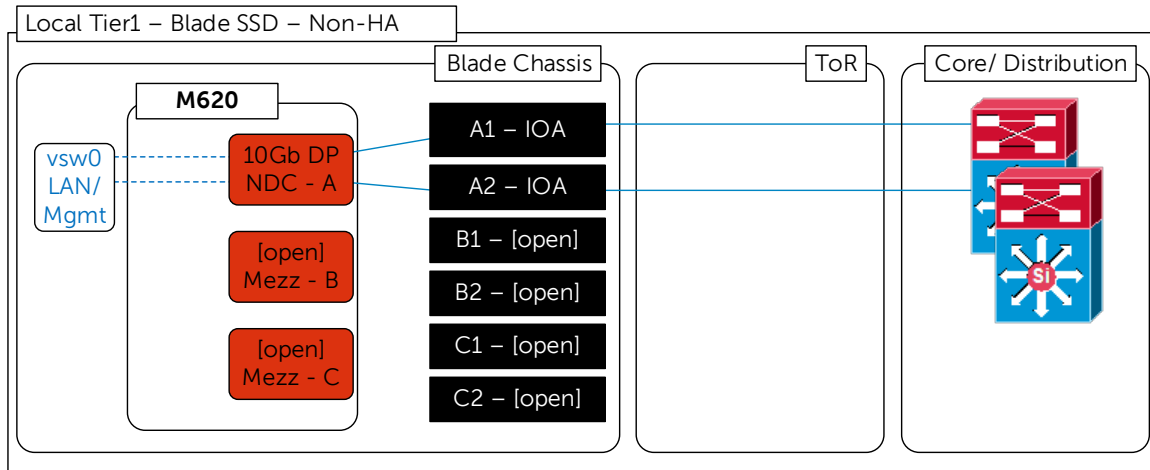
5.4.2 Local Tier 1 – Blade

The network configuration in this model will vary slightly between the Compute and Management hosts. The Management VMs are hosted on local storage and optionally on shared Tier 2 if Live Migration is a requirement. The following outlines the VLAN requirements for the Compute and Management hosts in this solution model:

- 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
 - VDI Management VLAN: Configured for VDI infrastructure traffic – L3 routed via core switch
 - Live Migration VLAN: Configured for Live Migration traffic – L2 switched only, trunked from Core (**HA**)
 - iSCSI VLAN: Configured for iSCSI traffic – L2 switched only via ToR switch (**HA**)
- An optional iDRAC VLAN are configured for all hardware management traffic – L3 routed via core switch

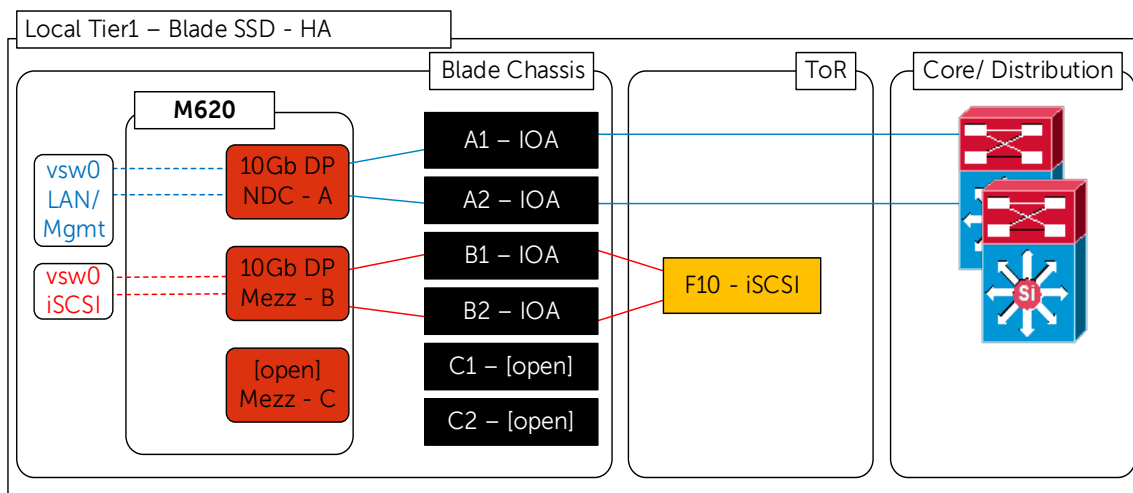
5.4.2.1 vSphere – Non HA

Each Local Tier 1 Compute and Management blade host have a 10Gb dual port LOM in the A Fabric that will flow through 2 x IOA blade interconnects. The B and C Fabrics are left open for optional expansion. Connections should pass through the blade mezzanines and interconnects per the diagram below. Configure the LAN traffic from the chassis interconnects to the ToR switch as a LAG if possible.



5.4.2.2 vSphere – HA

Following best practices, iSCSI and LAN traffic is physically separated into discrete Fabrics for the HA design. If mgmt. HA is required, the B Fabric is populated with a 10Gb DP NIC connecting through an additional pair of IOA blade interconnects.



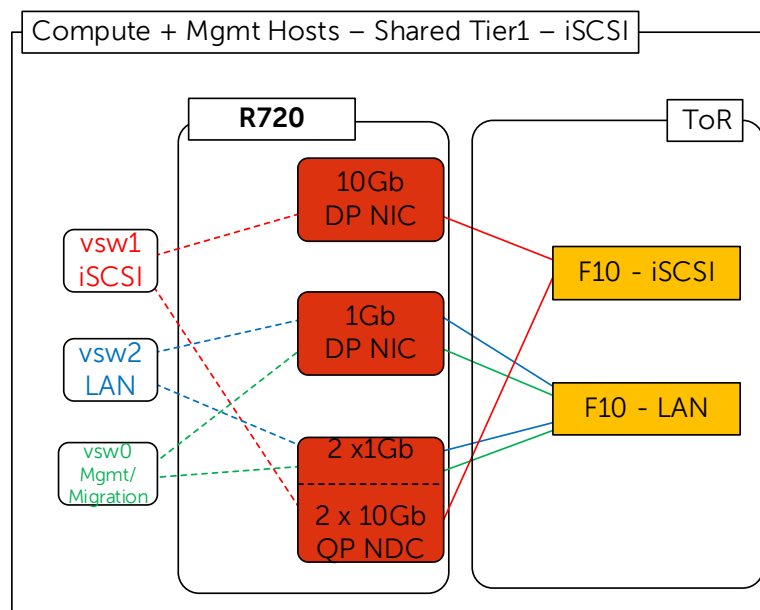
5.4.3 Shared Tier 1 – Rack – iSCSI

The network configuration in this model is identical between the Compute and Management hosts. Both need access to iSCSI storage since they are hosting VDI sessions from shared storage and both can leverage Live Migration as a result as well. The following outlines the VLAN requirements for the Compute and Management hosts in this solution model:

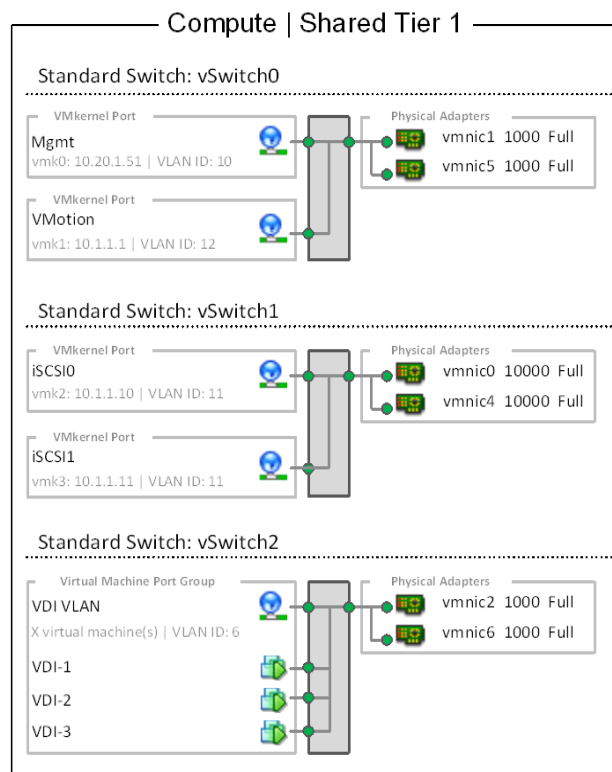
- Compute hosts (Shared Tier 1)
 - Management VLAN: Configured for hypervisor Management traffic – L3 routed via core switch
 - Live Migration VLAN: Configured for Live Migration traffic – L2 switched only, trunked from Core
 - iSCSI VLAN: Configured for iSCSI traffic – L2 switched only via ToR switch
 - 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
 - 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 optional iDRAC VLAN are 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 will have a quad port NDC (2 x 1Gb + 2 x 10Gb SFP+), a 10Gb dual port NIC, as well as a 1Gb dual port NIC. Isolate iSCSI onto its own vSwitch with redundant ports. Connections from all 3 vSwitches should pass through both the NDC and add-on NICs per the diagram below. Configure the LAN traffic from the server to the ToR switch as a LAG.

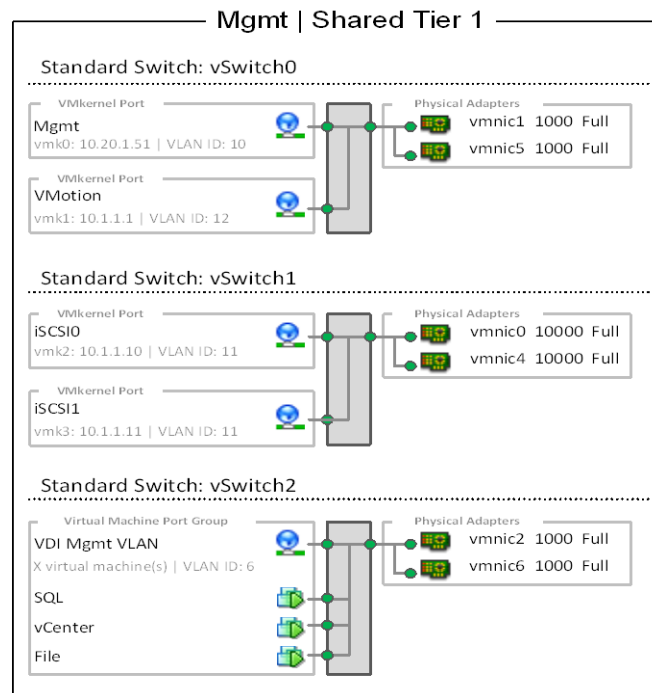
5.4.3.1 vSphere



vSwitch0 carries traffic for both Management and Live Migration which is VLAN-tagged so that either NIC can serve traffic for either VLAN. The Management VLAN is L3 routable while the Live Migration VLAN is L2 non-routable.

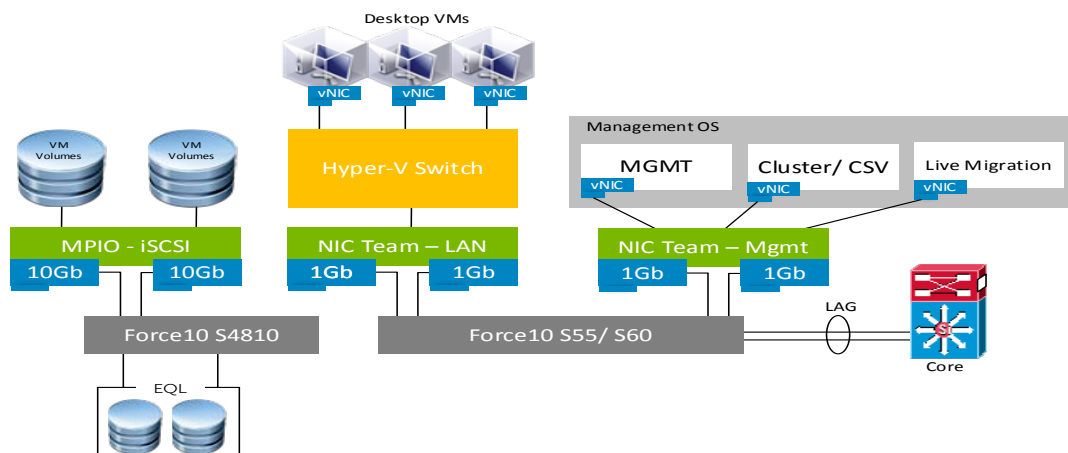


The Management server is configured identically except for the VDI Management VLAN which is fully routed and separated from the VDI VLAN used on the Compute host. Care is taken to ensure that all vSwitches are assigned redundant NICs that are NOT from the same PCIe device.

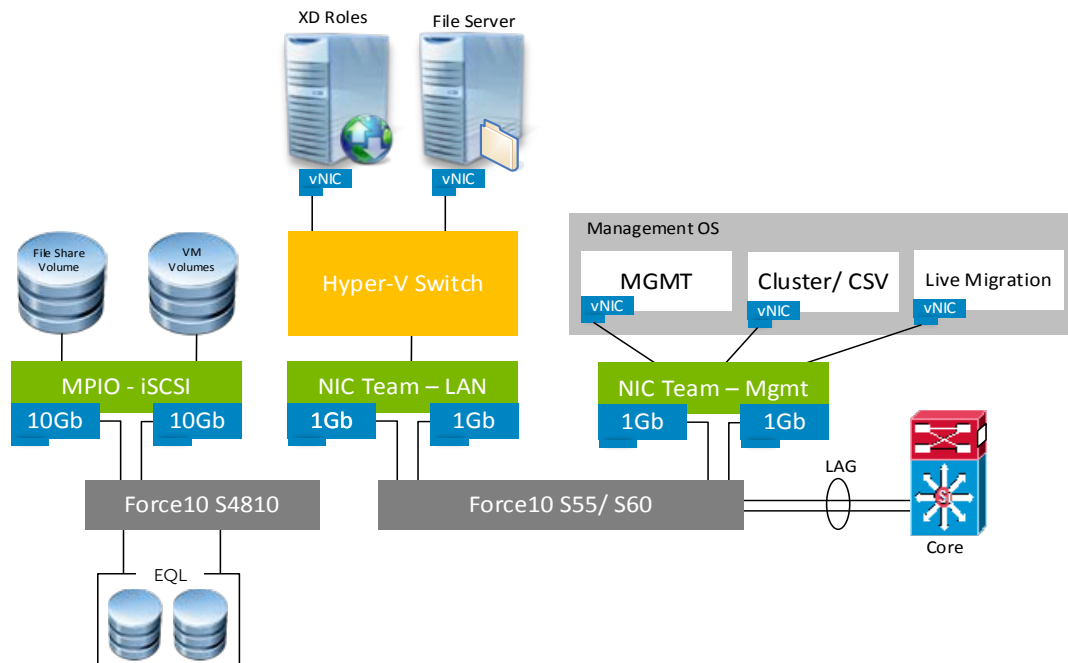


5.4.3.2 Hyper-V

The Hyper-V configuration, while identical in core requirements and hardware to the vSphere configuration, is executed a bit differently. Native Windows Server 2012 NIC Teaming is utilized to load balance and provide resiliency for network connections. The compute and management hosts are configured identically in this scenario with both connecting to shared storage. Two NIC teams is configured, one for the Hyper-V switch and one for the management OS, with Dell MPIO used to connect to shared storage. The desktop VMs will connect to the Hyper-V switch, while the management OS will connect directly to the NIC team dedicated to the Mgmt network.



Management hosts are configured in the same manner.



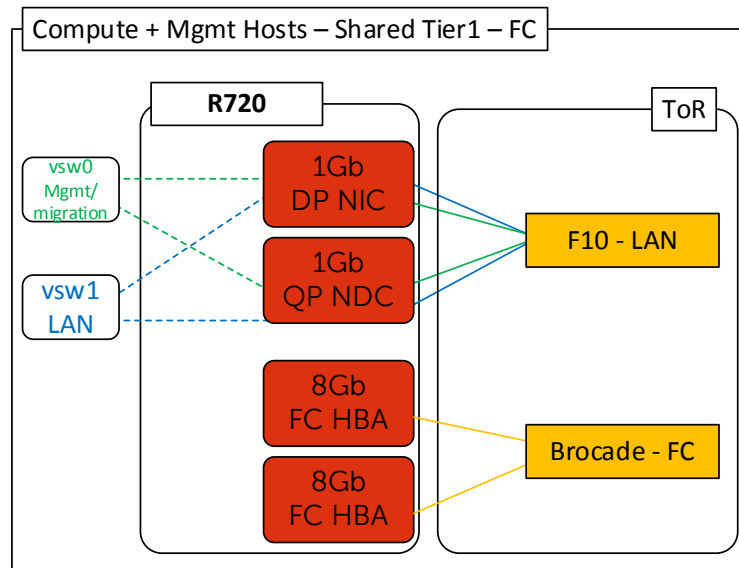
5.4.4 Shared Tier 1 – Rack – FC

Using Fiber Channel based storage eliminates the need to build iSCSI into the network stack but requires additional Fabrics to be built out. The network configuration in this model is identical between the Compute and Management hosts. Both need access to FC storage since they are hosting VDI sessions from shared storage and both can leverage Live Migration as a result. The following outlines the VLAN requirements for the Compute and Management hosts in this solution model:

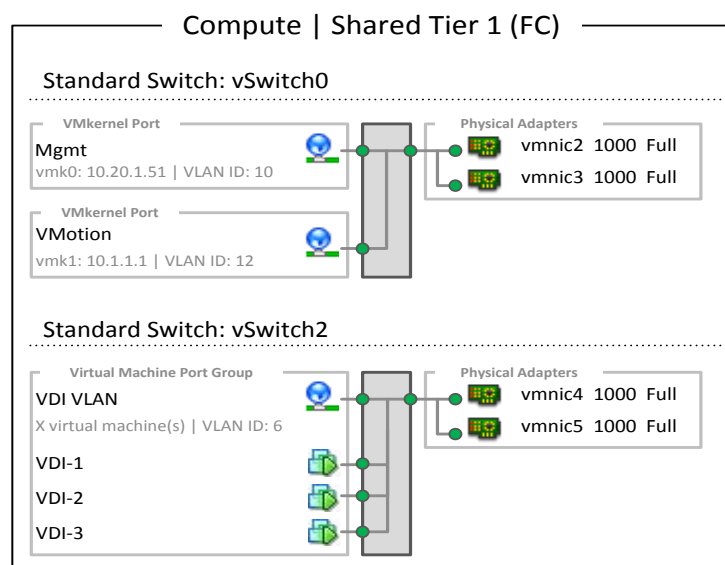
- Compute hosts (Shared Tier 1)
 - Management VLAN: Configured for hypervisor Management traffic – L3 routed via core switch
 - Live Migration VLAN: Configured for Live Migration traffic – L2 switched only, trunked from Core
 - VDI VLAN: Configured for VDI session traffic – L3 routed via core switch
- 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
 - VDI Management VLAN: Configured for VDI infrastructure traffic – L3 routed via core switch
- An optional iDRAC VLAN are 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 1Gb), a 1Gb dual port NIC, as well as 2 x 8Gb dual port FC HBAs. Connections from both vSwitches pass through both the NDC and add-on NICs per the diagram below. LAN traffic from the server to the ToR switch is configured as a LAG.

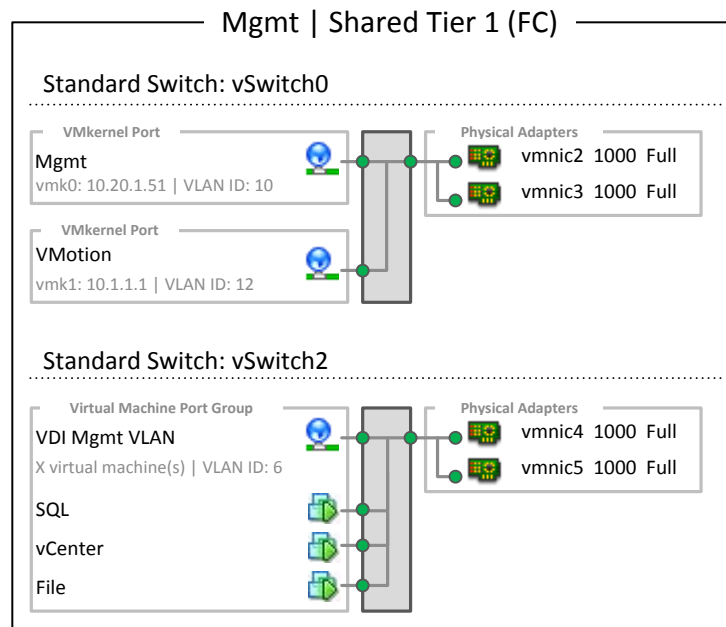
5.4.4.1 vSphere



vSwitch0 carries traffic for both Management and Live Migration which is VLAN-tagged so that either NIC can serve traffic for either VLAN. The Management VLAN is L3 routable while the Live Migration VLAN is L2 non-routable.

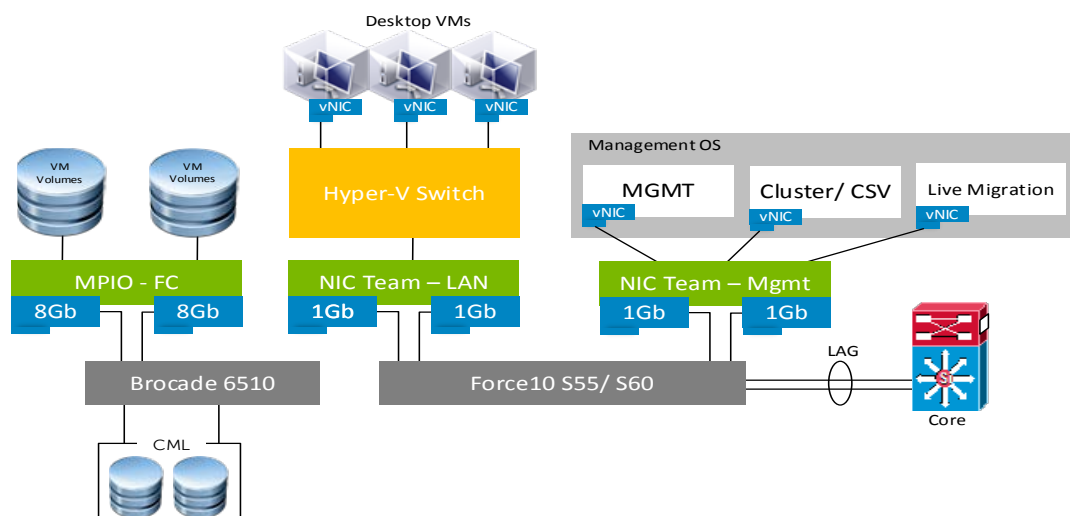


The Management server is configured identically except for the VDI Management VLAN which is fully routed but is separated from the VDI VLAN used on the Compute host.

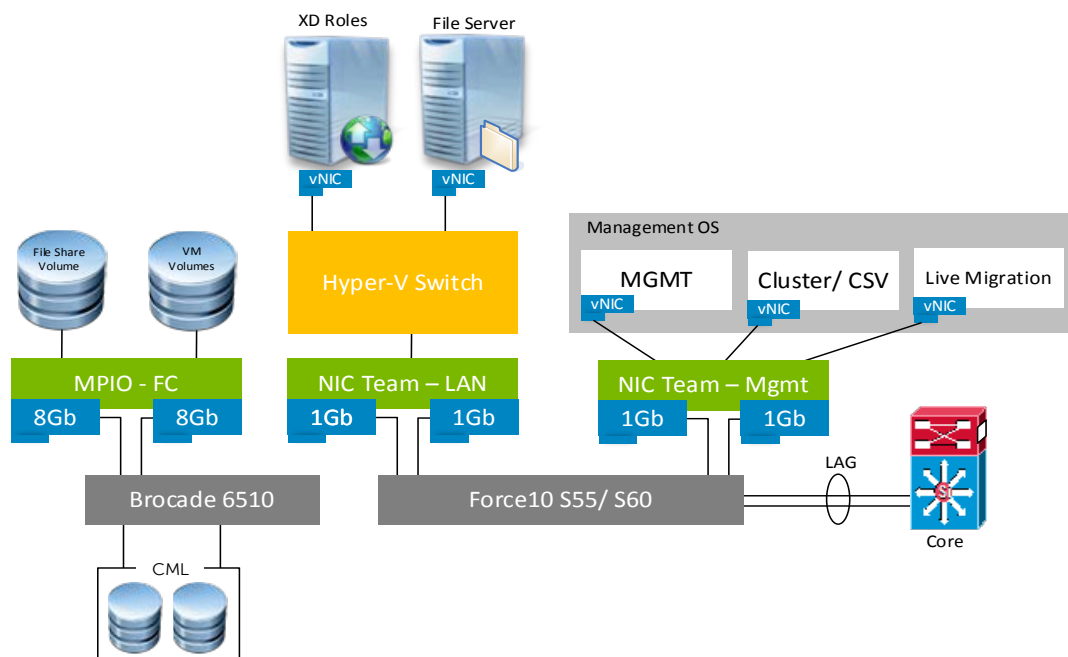


5.4.4.2 Hyper-V

The Hyper-V configuration, while identical in core requirements and hardware to the vSphere configuration, is executed a bit differently. Native Windows Server 2012 NIC Teaming is utilized to load balance and provide resiliency for network connections. The compute and management hosts are configured identically in this scenario with both connecting to shared storage. Two NIC teams is configured, one for the Hyper-V switch and one for the management OS, with Dell MPIO used to connect to shared storage. The desktop VMs will connect to the Hyper-V switch, while the management OS will connect directly to the NIC team dedicated to the Mgmt network.



Management hosts are configured in the same manner.



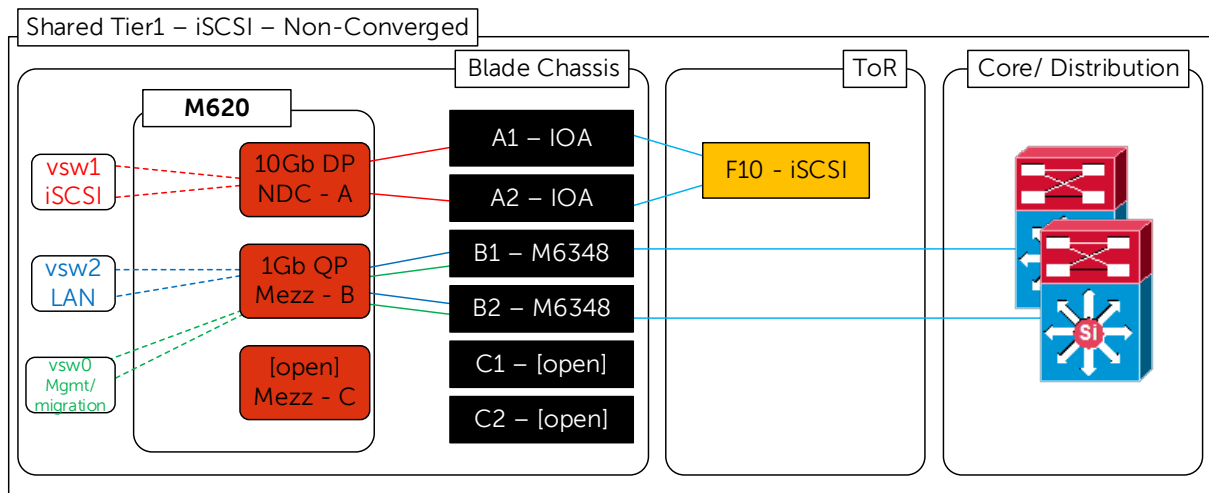
5.4.5 Shared Tier 1 – Blade – iSCSI

The network configuration in this model is identical between the Compute and Management hosts. The following outlines the VLAN requirements for the Compute and Management hosts in this solution model:

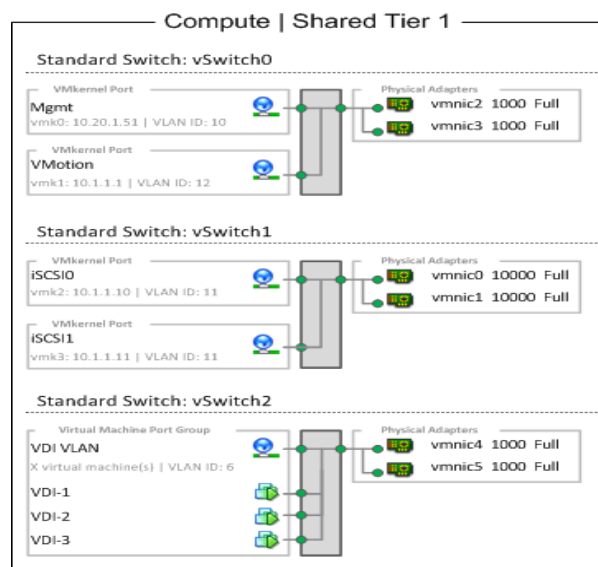
- Compute hosts (Shared Tier 1)
 - Management VLAN: Configured for hypervisor Management traffic – L3 routed via core switch
 - Live Migration VLAN: Configured for Live Migration traffic – L2 switched only, trunked from Core
 - iSCSI VLAN: Configured for iSCSI traffic – L2 switched only via ToR switch
 - 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
 - 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 optional iDRAC VLAN are 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 blade host will have a 10Gb dual port LOM in the A Fabric and a 1Gb quad port NIC in the B Fabric. 10Gb iSCSI traffic will flow through A Fabric using 2 x IOA blade interconnects. 1Gb LAN traffic will flow through the B Fabric using 2 x M6348 blade interconnects. The C Fabric is left open for future expansion. Connections from 10Gb and 1Gb traffic vSwitches should pass through the blade mezzanines and interconnects per the diagram below. LAN traffic from the server to the ToR switch is configured as a LAG.

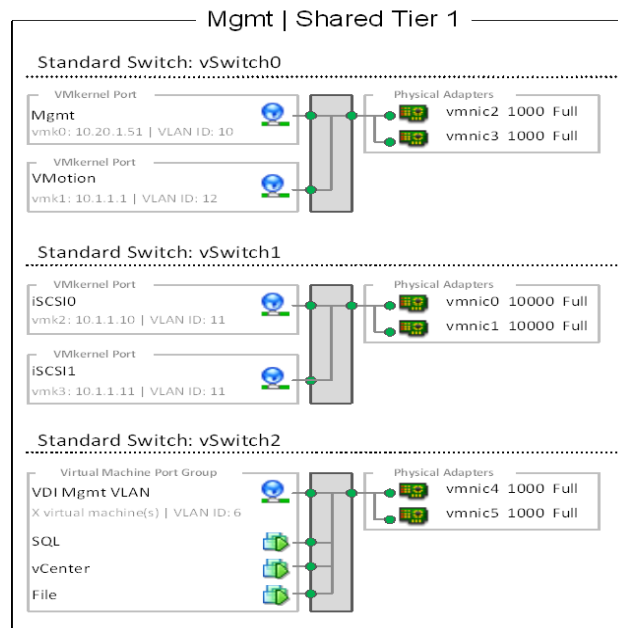
5.4.5.1 vSphere



vSwitch0 carries traffic for both Management and Live Migration which is VLAN-tagged so that either NIC can serve traffic for either VLAN. The Management VLAN is L3 routable while the Live Migration VLAN is L2 non-routable.

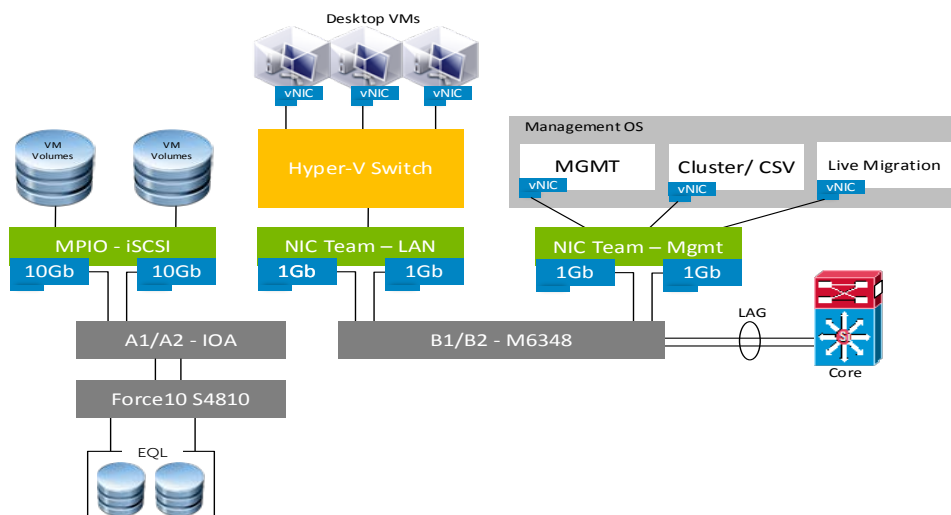


The Management server is configured identically except for the VDI Management VLAN which is fully routed but is separated from the VDI VLAN used on the Compute host.

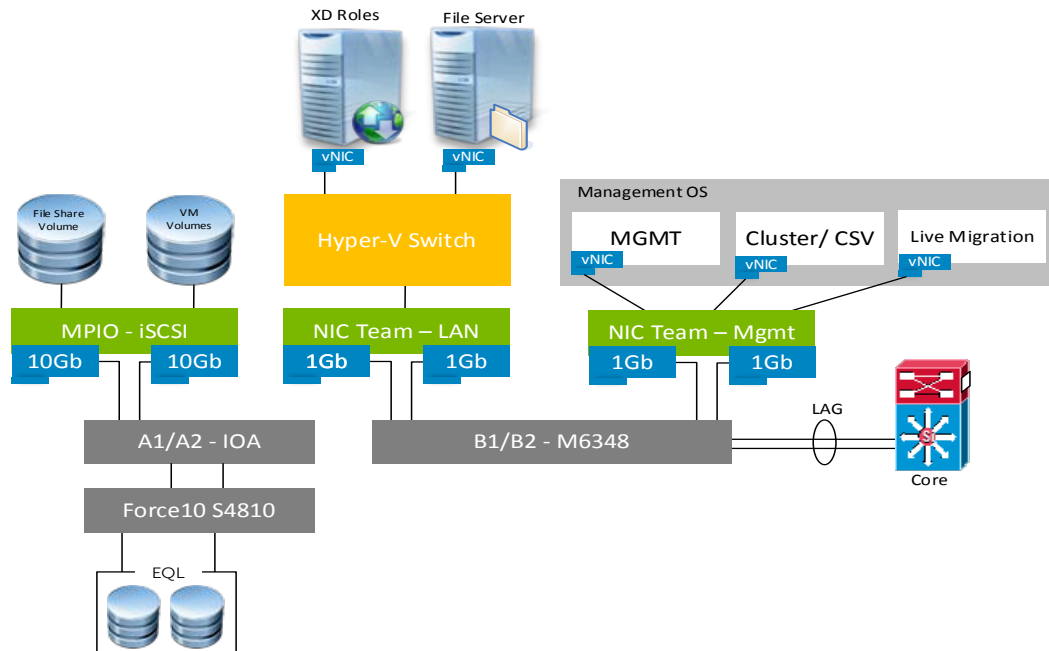


5.4.5.2 Hyper-V

The Hyper-V configuration, while identical in core requirements and hardware to the vSphere configuration, is executed a bit differently. Native Windows Server 2012 NIC Teaming is utilized to load balance and provide resiliency for network connections. The compute and management hosts are configured identically in this scenario with both connecting to shared storage. Two NIC teams is configured, one for the Hyper-V switch and one for the management OS, with Dell MPIO used to connect to shared storage. The desktop VMs will connect to the Hyper-V switch, while the management OS will connect directly to the NIC team dedicated to the Mgmt network.



Management hosts are configured in the same manner.



5.4.6 Shared Tier 1 – Blade – FC

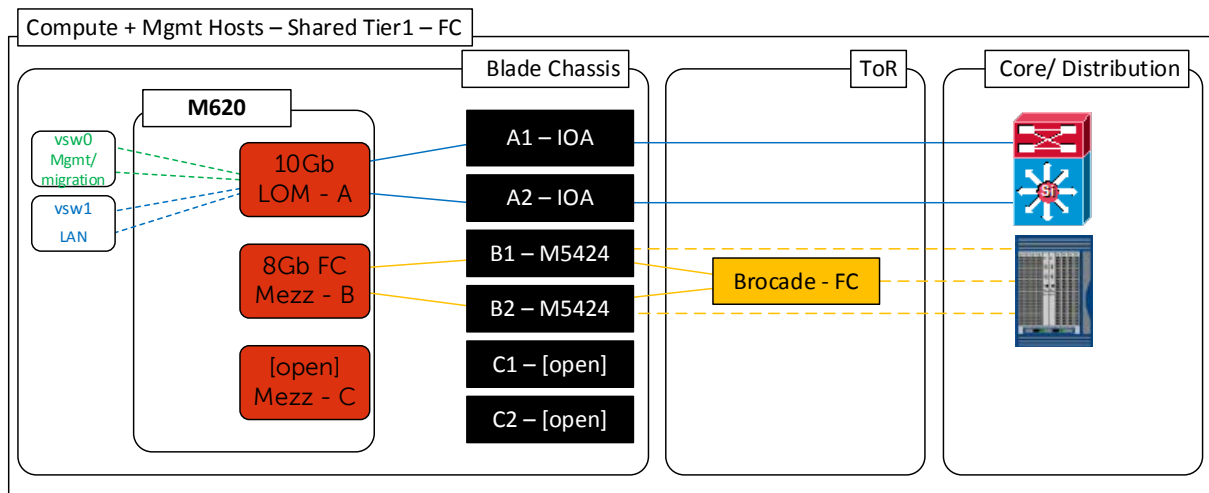
Using Fiber Channel based storage eliminates the need to build iSCSI into the network stack but requires additional Fabrics to be built out. The network configuration in this model is identical between the Compute and Management hosts. The following outlines the VLAN requirements for the Compute and 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
 - 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
 - VDI Management VLAN: Configured for VDI infrastructure traffic – L3 routed via core switch
- An optional DRAC VLAN are 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 blade will have a 10Gb dual port LOM in the A Fabric and an 8Gb dual port HBA in the B Fabric. All LAN and Mgmt traffic will flow through the A Fabric using 2 x IOA blade interconnects partitioned to the connecting blades. 8Gb FC traffic will flow through the B Fabric

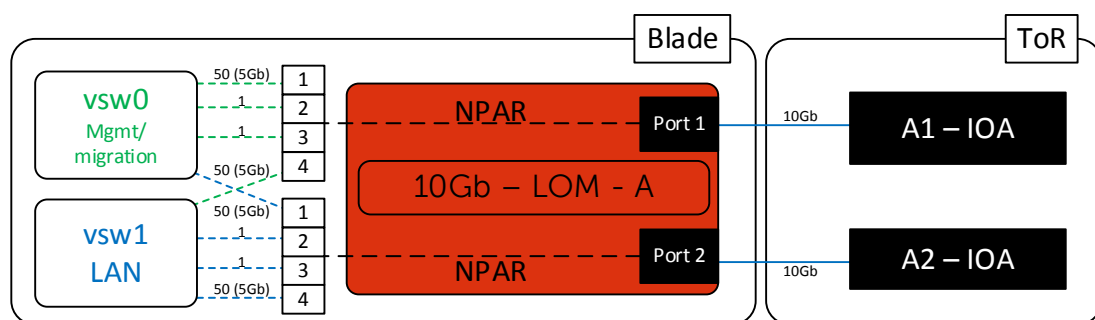
using 2 x M5424 blade interconnects. The C Fabric is left open for future expansion. Connections from the vSwitches and storage Fabrics should pass through the blade mezzanines and interconnects per the diagram below. Configure the LAN traffic from the server to the ToR switch as a LAG.

5.4.6.1 vSphere



5.4.6.1.1 Shared Tier 1 – Blade – Network Partitioning

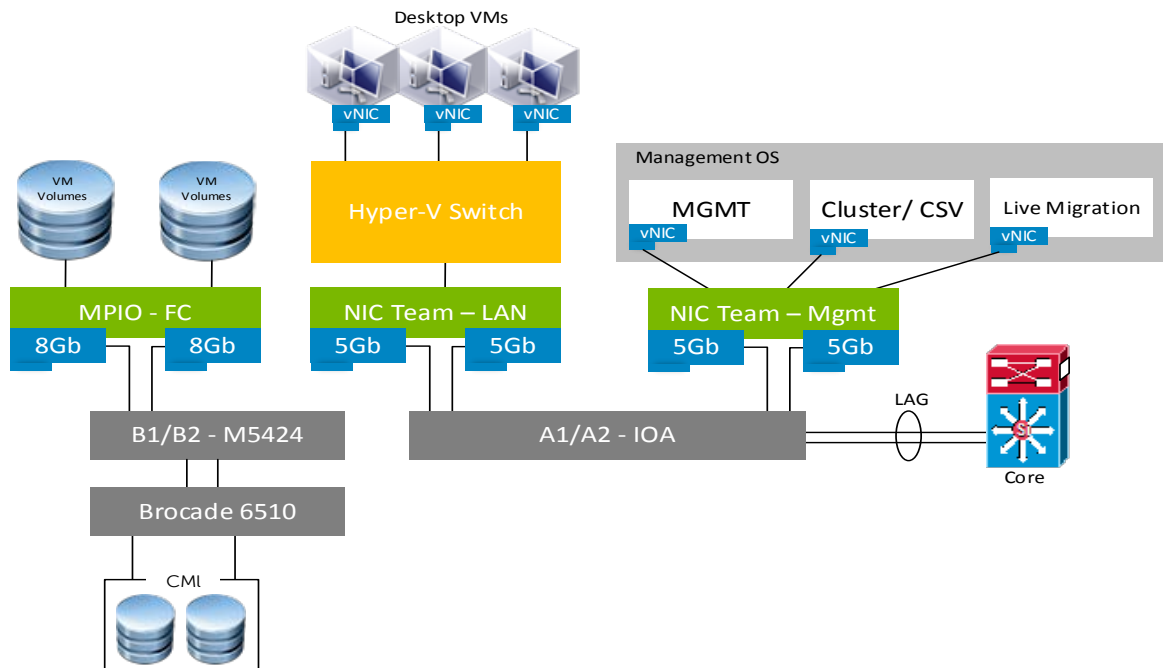
Network partitioning (NPAR) takes place within the UEFI of the 10GB LOMs of each blade in the A Fabric. Partitioning allows a 10Gb NIC to be split into multiple logical NICs that are assigned differing amounts of bandwidth. 4 partitions are defined per NIC with the amounts specified below. We only require 2 partitions per NIC port so the unused partitions receive a bandwidth of 1. Partitions are oversubscribed, but not the reverse. We are partitioning out a total of 4 x 5Gb NICs with the remaining 4 unused. Use care to ensure that each vSwitch receives a NIC from each physical port for redundancy.



5.4.6.2 Hyper-V

The Hyper-V configuration, while identical in core requirements and hardware to the vSphere configuration, is executed a bit differently. Native Windows Server 2012 NIC Teaming is utilized to load balance and provide resiliency for network connections. The compute and management hosts

are configured identically in this scenario with both connecting to shared storage. Two NIC teams is configured, one for the Hyper-V switch and one for the management OS, with Dell MPIO used to connect to shared storage. The desktop VMs will connect to the Hyper-V switch, while the management OS will connect directly to the NIC team dedicated to the Mgmt network. The LAN connections are partitioned using NPAR to provide 2 x 5Gb vNICs to each NIC team. Please see section 4.7.5.1.1 above for more information.



5.5 Scaling Guidance

Each component of the solution architecture scales independently according to the desired number of supported users. PVS scales differently from MCS, as does the hypervisor in use. Using the new Intel Ivy Bridge CPUs, rack and blade servers now scale equally from a compute perspective.

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

Component	Metric	Horizontal scalability	Vertical scalability
Virtual Desktop Host/Compute Servers	VMs per physical host	Additional hosts and clusters added as necessary	Additional RAM or CPU compute power
Provisioning Servers	Desktops per instance	Additional servers added to the Provisioning Server farm	Additional network and I/O capacity added to the servers
Desktop Delivery Servers	Desktops per instance (dependent on SQL performance as well)	Additional servers added to the XenDesktop Site	Additional virtual machine resources (RAM and CPU)
XenApp Servers	Desktops per instance	Additional virtual servers added to the XenDesktop Site	Additional physical servers to host virtual XenApp servers.
Storefront Servers	Logons/ minute	Additional servers added to the Storefront environment	Additional virtual machine resources (RAM and CPU)
Database Services	Concurrent connections, responsiveness of reads/ writes	Migrate databases to a dedicated SQL server and increase the number of management nodes	Additional RAM and CPU for the management nodes
File Services	Concurrent connections, responsiveness of reads/ writes	Split user profiles and home directories between multiple file servers in the cluster. File services can also be migrated to the optional NAS device to provide high availability.	Additional RAM and CPU for the management nodes

The tables below indicate the server platform, desktop OS, hypervisor and delivery mechanism.

5.5.1 Windows 8.1 – vSphere

Rack or Blade, Win8.1, vSphere, PVS										
Standard User Count	Enhanced User Count	Professional User Count	Physical Mgmt Servers	Physical Host Servers	M1000e Blade Chassis	Virtual PVS Servers	Virtual Delivery Controllers	Virtual License Servers	Virtual StoreFront Servers	Virtual vCenter Server
165	110	90	1	1	1	1	1	1	1	1
1000	770	630	1	7	1	1	1	1	1	1
2000	1430	1170	2	13	1	2	1	1	1	1
3000	2090	1710	3	19	2	3	1	1	1	1
4000	2750	2250	3	25	2	4	1	1	1	1
5000	3410	2790	4	31	3	5	1	1	1	1
6000	4070	3330	5	37	3	5	1	1	1	1
7000	4730	3870	5	43	3	6	1	1	1	1
8000	5390	4410	6	49	4	7	1	1	1	1
9000	6050	4950	7	55	4	8	1	1	1	1
10000	6710	5490	7	61	5	9	1	1	1	1

Rack or Blade, Win8.1, vSphere, MCS									
Standard User Count	Enhanced User Count	Professional User Count	Physical Mgmt Servers	Physical Host Servers	M1000e Blade Chassis	Virtual Delivery Controllers	Virtual License Servers	Virtual StoreFront Servers	Virtual vCenter Server
165	110	90	1	1	1	1	1	1	1
1000	770	630	2	7	1	1	1	1	1
2000	1430	1170	2	13	1	1	1	1	1
3000	2090	1710	2	19	2	1	1	1	1
4000	2750	2250	2	25	2	1	1	1	1
5000	3410	2790	3	31	3	1	1	1	1

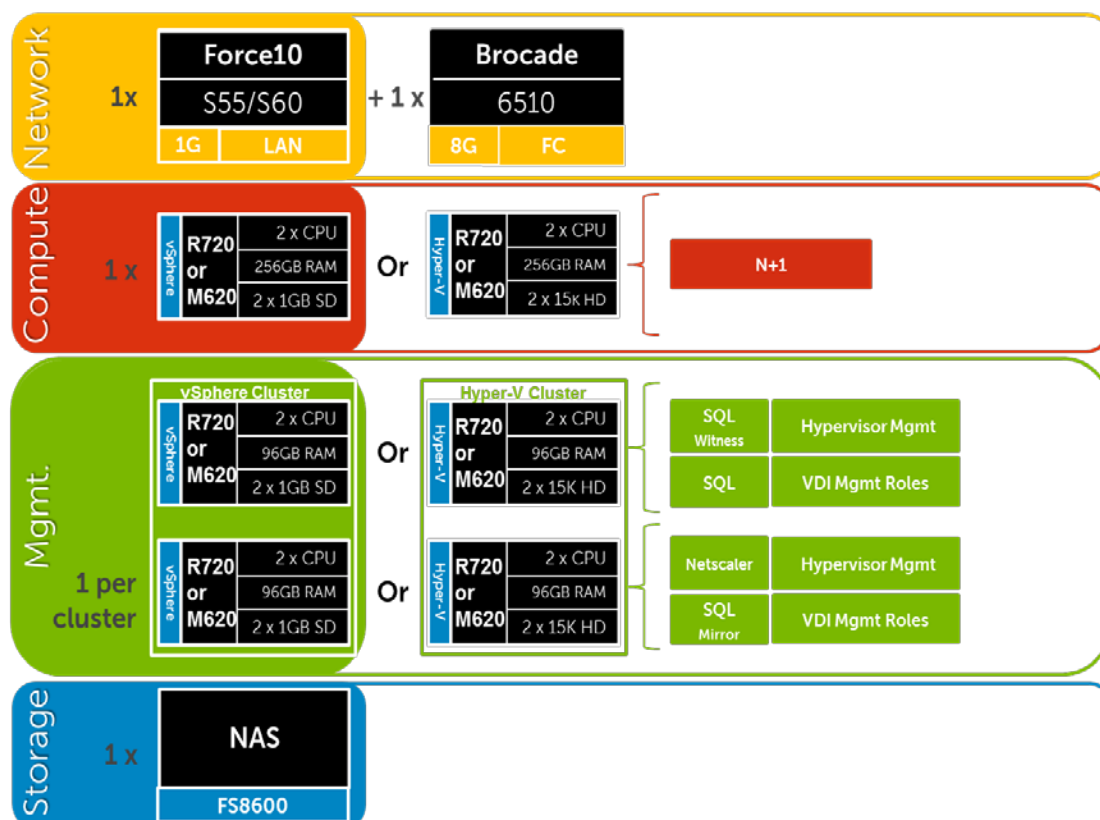
5.5.2 Windows 8.1 – Hyper-V

Rack or Blade, Win8.1, Hyper-V, PVS										
Standard User Count	Enhanced User Count	Professional User Count	Physical Mgmt Servers	Physical Host Servers	M1000e Blade Chassis	Virtual PVS Servers	Virtual Delivery Controllers	Virtual License Servers	Virtual StoreFront Servers	Virtual SCVMM Server
170	120	100	1	1	1	1	1	1	1	1
1000	720	600	1	6	1	1	1	1	1	1
2000	1440	1200	2	12	1	2	1	1	1	1
3000	2160	1800	3	18	2	3	1	1	1	1
4000	2880	2400	3	24	2	4	1	1	1	1
5000	3600	3000	4	30	3	5	1	1	1	1
6000	4320	3600	5	36	3	5	1	1	1	1
7000	5040	4200	5	42	3	6	1	1	1	1
8000	5760	4800	6	48	4	7	1	1	1	1
9000	6360	5300	7	53	4	8	1	1	1	1
10000	7080	5900	7	59	5	9	1	1	1	1

Rack or Blade, Win8.1, Hyper-V, MCS									
Standard User Count	Enhanced User Count	Professional User Count	Physical Mgmt Servers	Physical Host Servers	M1000e Blade Chassis	Virtual Delivery Controllers	Virtual License Servers	Virtual StoreFront Servers	Virtual SCVMM Server
170	120	100	1	1	1	1	1	1	1
1000	720	600	2	6	1	1	1	1	1
2000	1440	1200	2	12	1	1	1	1	1
3000	2160	1800	2	18	2	1	1	1	1
4000	2880	2400	2	24	2	1	1	1	1
5000	3600	3000	3	30	3	1	1	1	1

5.6 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 for LAN, iSCSI, or FC 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 the management layer, SQL is mirrored or clustered, NetScaler is leveraged for load balancing, and a NAS device are used to host file shares. Storage protocol switch stacks and NAS selection will vary based on chosen solution architecture. The following diagram depicts the additional components required to make a standard solution HA.

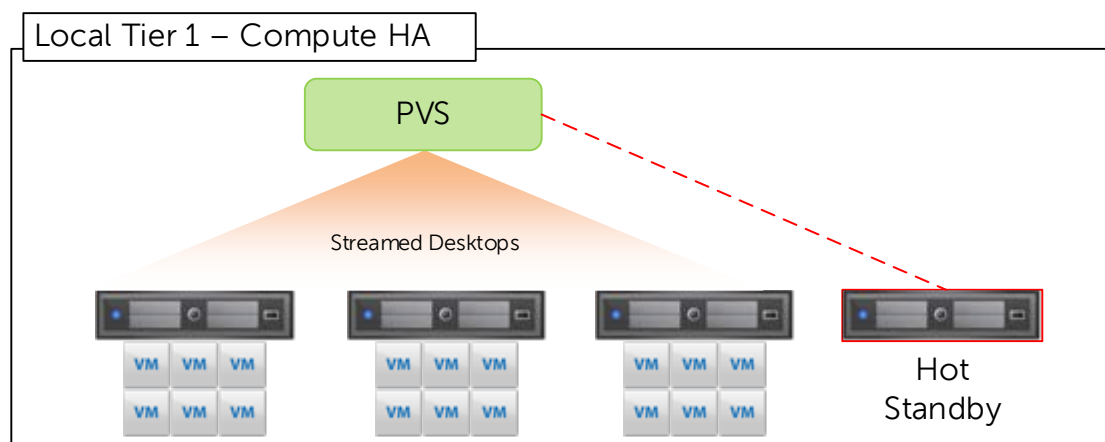


The HA options provide redundancy for all critical components in the stack while improving the performance and efficiency of the solution as a whole.

- An additional switch is added at the network Tier which is configured with the original as a stack and thereby equally spreading each host's network connections across both switches.
- At the compute Tier an additional ESXi host is added to provide N+1 protection provided by Citrix PVS. In a rack based solution with local Tier 1 storage, there is no vSphere HA cluster in the compute Tier as VMs run on local disks.
- A number of enhancements occur at the Management Tier, the first of which is the addition of another host. The Management hosts are configured in an HA cluster. All applicable Citrix server roles can then be duplicated on the new host where connections to each are load balanced via the addition of a virtual NetScaler appliance. SQL Server databases also receive greater protection through the addition and configuration of a 3-way SQL mirror with a witness.

5.6.1 Compute Layer HA (Local Tier 1)

The optional HA bundle adds an additional host in the Compute and Management layers to provide redundancy and additional processing power to spread out the load. The Compute layer in this model does not leverage shared storage so hypervisor HA does not provide a benefit here. To protect the Compute layer, with the additional server added, PVS will ensure that if a single host fails, its load is spun up on the hot standby.



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

vSphere HA Admission control are configured one of three ways to protect the cluster. This will vary largely by customer preference but the most manageable and predictable options are percentage reservations or a specified hot standby. Reserving by percentage will reduce the overall per host density capabilities but will make some use of all hardware in the cluster. Additions and subtractions

of hosts will require the cluster to be manually rebalanced. Specifying a failover host, on the other hand, will ensure maximum per host density numbers but will result in hardware sitting idle.

Admission Control
The vSphere HA Admission control policy determines the amount of cluster capacity that is reserved for VM failovers. Reserving more failover capacity allows more failures to be tolerated but reduces the number of VMs that can be run.

☒ Enable: Disallow VM power on operations that violate availability constraints
☐ Disable: Allow VM power on operations that violate availability constraints

Admission Control Policy
Specify the type of policy that admission control should enforce.


☐ Host failures the cluster tolerates: 1

☒ Percentage of cluster resources reserved as failover spare capacity: 25 % CPU, 25 % Memory

☐ Specify failover hosts: 0 hosts specified. Click to edit.

5.6.2 vSphere HA (Shared Tier 1)


As shown below, both Compute and Management hosts are identically configured, within their respective Tiers, and leverage shared storage thus making full use of vSphere HA. The Compute hosts are configured in an HA cluster following the boundaries of vCenter with respect to limits imposed by VMware (3000 VMs per vCenter). This will result in multiple HA clusters managed by multiple vCenter servers.



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

5.6.3 Hyper-V HA (Shared Tier 1)

The computer layer hosts are provided using typical N+1 fashion as shown in the diagram below. An additional host is added to the pool and is configured to absorb additional capacity or as a standby node should another fail.




5.6.4 Management Server High Availability

The applicable core Citrix roles are load balanced via DNS namespaces by default. In environments requiring HA, NetScaler VPX is introduced to manage load-balancing efforts. XenDesktop, PVS, and vCenter configurations are stored in SQL which is protected via the SQL mirror.

Citrix license server is protected from host hardware failure and Citrix licensing grace period by default. If customer desires, the license server are optionally protected further via the form of a cold stand-by VM residing on an opposing management host. A vSphere scheduled task is used, for example, to clone the VM to keep the stand-by VM current. Note – In the HA option, there is no file server VM, its duties have been replaced by introducing a NAS head.

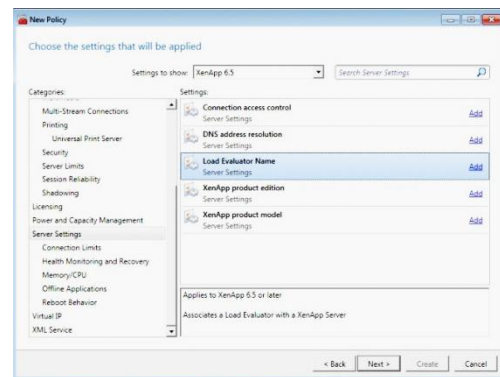
The following protects each of the critical infrastructure components in the solution:

- The Management hosts are configured in a Hyper-V cluster (Node and Disk Majority).
- The storage volume that hosts the Management VMs is upgraded to a Cluster Shared Volume (CSV) so all hosts in a cluster can read and write to the same volumes.
- SQL Server mirroring is configured with a witness to further protect SQL.



5.6.5 XenApp Server High Availability

The high availability configuration of XenApp virtual servers is introduced to avoid application or desktop unavailability due to a physical server failure. Enabling this configuration typically involves adding an additional one or more physical XenApp server hosts and then the appropriate number of virtual XenApp servers to attain the desired level of resiliency and high availability. A single server is configured with the addition of two (2) to eight (8) additional virtual servers per physical host. Once joined to the farm, the additional XenApp virtual servers will leverage native XenApp high availability via Load Evaluators and failover policies.



5.6.6 Provisioning Services High Availability

The following components of the Provisioning Server Hierarchy are delivered in a highly available configuration:

- Provisioning Server
- Database Storage Device (SAN)
- vDisk Shared Storage System (SAN).

With the HA feature, a target device can connect to its vDisk through any Provisioning Server listed in the boot file. The HA feature is enabled on the Properties' Options tab for the vDisk.

The target device attempts to connect to one Provisioning Server at a time, in the order listed. If the attempt fails, it tries the next Provisioning Server in the list. If the attempt with the last server in the list fails, it repeats the process. From the Console Tool menu, you use the Configure Bootstrap option to add a Provisioning Server to the list.

The connecting Provisioning Server does not necessarily become the Provisioning Server that accesses the vDisk on behalf of the target device. The connecting Provisioning Server chooses the Provisioning Server that accesses the vDisk according to the Boot Behavior property of the connecting target device.

The HA feature requires that each Provisioning server has access to a shared copy of the database, an identical copy of the vDisk and the ability to write to the Write Cache. Within the environment the database is hosted on the clustered SQL server not one of the Provisioning servers and as such the SQL platform will ensure the database remains available. The vDisks and Write Cache are configured in two ways:

- A shared storage system that ensures the availability of the Provisioning Server database, and vDisks, assuming that the Write Cache is hosted on the Provisioning server
- Each Provisioning server has access to an identical copy of the vDisk via the same local path, and the Write cache located on the target device

Within the environment the second option is utilized, this is commonly referred to as a distributed model because the vDisks are located on each Provisioning server not on a shared storage platform.

Each Provisioning Server is responsible for the Write cache of each desktop that it is hosting. This Write cache is readable by the resilient Provisioning Server and as it is held on the target device (in this case virtual machine). In the event of a Provisioning Server failure, all desktops that were hosted will transfer to an alternate provisioning server for that site and users is unaware of the failure.

Streaming Services run under a user account with Service account credentials. The Service account credential (user account name and password) is a domain account that is configured on each Provisioning Server, in order to access the Streaming Service and the shared database.

5.6.7 Windows File Services High Availability


High availability for file services are provided by the Dell FS7600, FS8600 or PowerVault NX3300 clustered NAS devices. To ensure proper redundancy, the NAS cabling is distributed between ToR switches.

Unlike the FS8600, the FS7600 and NX3300 do not support for 802.1q (VLAN tagging) so the connecting switch ports are configured with native VLANs, both iSCSI and LAN/ VDI traffic ports. Best practice dictates that all ports are connected on both controller nodes. The backend ports are used for iSCSI traffic to the storage array as well as for internal NAS functionality (cache mirroring and cluster heart beat). Front-end ports are configured using Adaptive Load Balancing or a LAG (LACP).

The Dell Wyse Solutions Engineering recommendation is to configure the original file server VM to use RDMs or PTDs to access the storage LUNs; therefore migration to the NAS is simplified by changing the presentation of these LUNs from the file server VM to the NAS.

5.6.8 SQL Server High Availability

HA for SQL is provided via a 3-server synchronous mirror configuration that includes a witness (High safety with automatic failover). 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 active SQL server. Place the principal VM that will host the primary copy of the data on the first Management host. Place the mirror and witness VMs on the second or later Management hosts. Mirror all critical databases to provide HA protection.

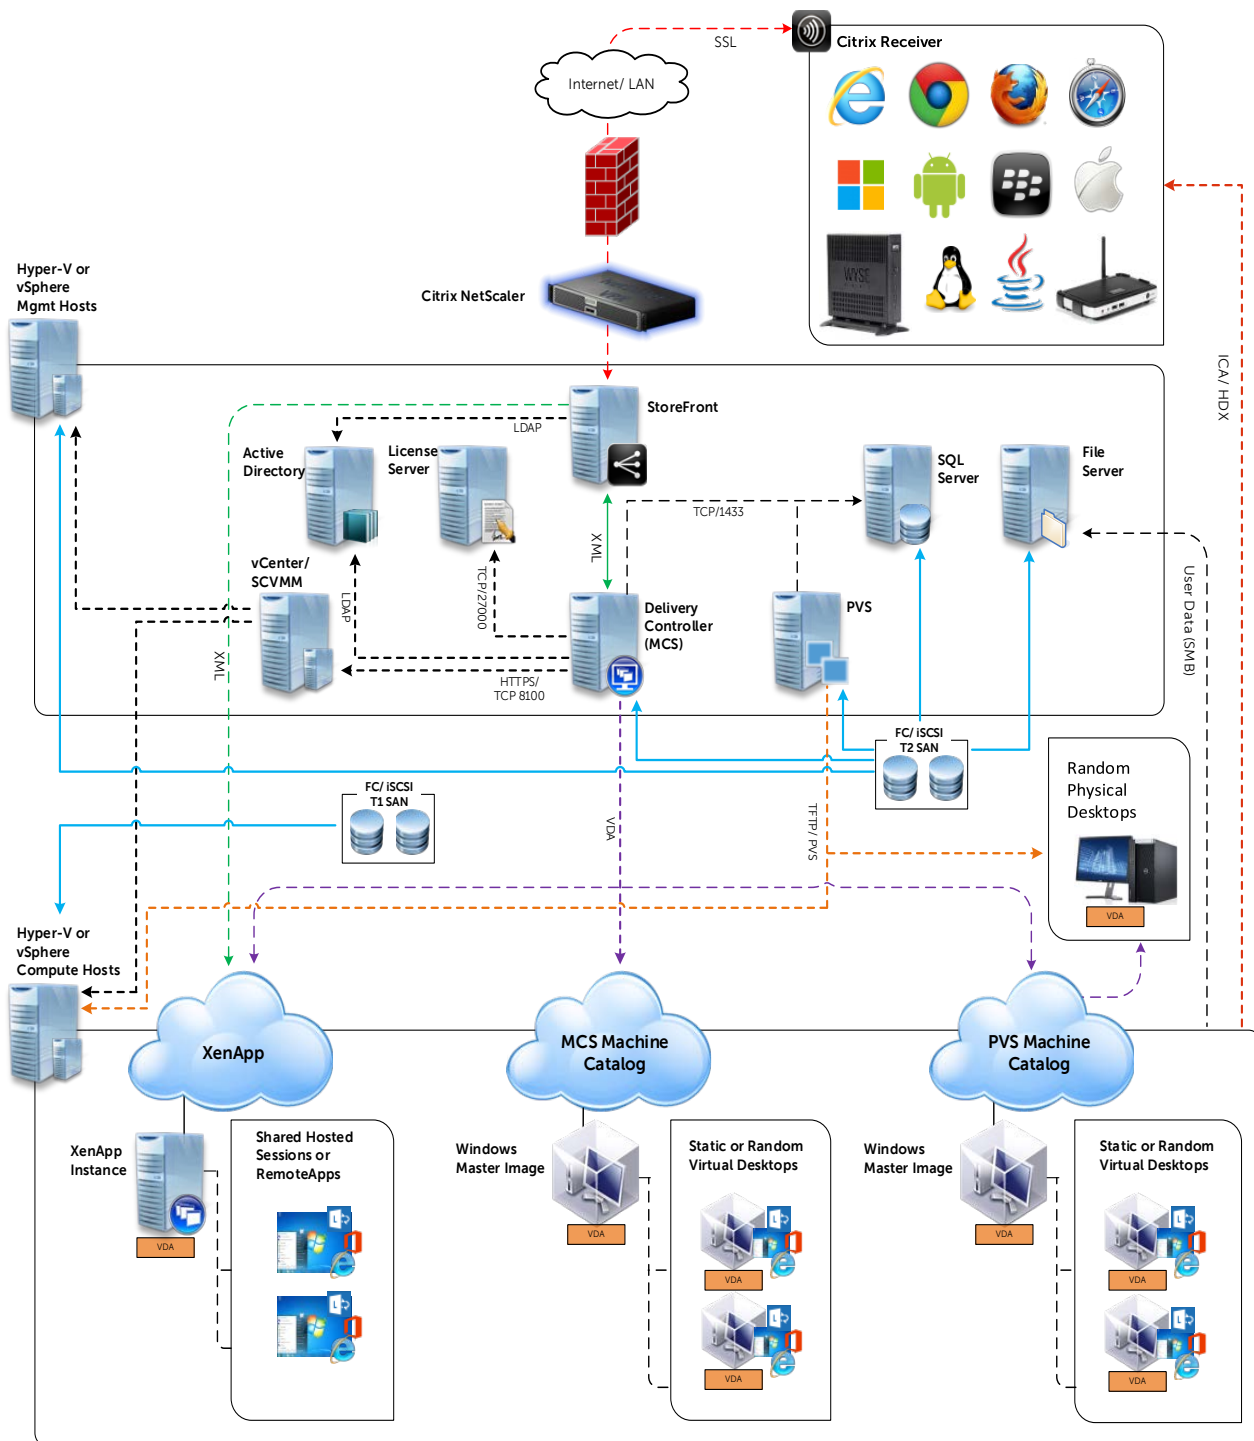


Please refer to this Citrix document for more information: [LINK](#)

The following article details the step-by-step mirror configuration: [LINK](#)

Additional resources are found in TechNet: [LINK1](#) and [LINK2](#)

5.7 Dell Wyse Datacenter for XenDesktop Communication Flow



6 Customer Provided Stack Components

6.1 Customer Provided Storage Requirements

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

Feature	Minimum Requirement	Notes
Total Tier 2 Storage	User count and workload dependent	1Gb/ 10Gb iSCSI or FC storage required on NL SAS disks minimally.
Tier 1 IOPS Requirement	(Total Users) x workload IOPS	6-30 IOPS per user may be required depending on workload. T1 storage is capable of providing user IOPS requirement concurrently to all hosted users.
Tier 2 IOPS Requirement	(Total Users) x 1 IOPS	File share usage and size of deployment may shift this requirement.
Data Networking	1Gb Ethernet for LAN/ T2 SCSI 10Gb Ethernet for T1 iSCSI 8Gb FC	Data networking traffic is isolated on dedicated NICs and HBAs in each applicable host.
RAID Support	10, 6	RAID 10 is leveraged for T1 local storage and if required for shared T2. RAID 6 is used in shared T1.

6.2 Customer Provided Switching Requirements

In the event that a customer wishes to provide his or her own rack network switching solution for a Dell Wyse Datacenter solution, the following minimum hardware requirements must be met:

Feature	Minimum Requirement	Notes
Switching Capacity	Line rate switch	1Gb or 10Gb switching pertinent to solution being implements. 1Gb switching for iSCSI is only suitable for T2. T1 iSCSI requires 10Gb.
10Gbps Ports	Uplink to Core Shared T1 iSCSI/ Converged	Dell Wyse Datacenter leverages 10Gb for Shared Tier 1 solution models. 10Gb uplinks to the core network are also recommended.
1Gbps Ports	5x per Management server per Compute Server per Storage Array	5x 6x Dell Wyse Datacenter leverages 1Gbps network connectivity for LAN traffic and T2 storage in Local Tier 1 solution models.
VLAN Support	IEEE 802.1Q tagging and port-based VLAN support.	
Stacking Capability	Yes	The ability to stack switches into a consolidated management framework is preferred to minimize disruption and planning when uplinking to core networks.

7 Solution Performance and Testing

7.1 Load Generation and Monitoring

7.1.1 Login VSI 4 – Login Consultants

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

7.1.2 Liquidware Labs Stratusphere UX

Stratusphere UX was used during each test run to gather data relating to User Experience and desktop performance. Data was gathered at the Host and Virtual Machine layers and reported back to a central server (Stratusphere Hub). The hub was then used to create a series of "Comma Separated Values" (.csv) reports which have then been used to generate graphs and summary tables of key information. In addition the Stratusphere Hub generates a magic quadrate style scatter plot showing the Machine and IO experience of the sessions. The Stratusphere hub was deployed onto the core network therefore its monitoring did not impact the servers being tested. This core network represents an existing customer environment and also includes the following services:

- Active Directory
- DNS
- DHCP
- Anti-Virus

Stratusphere UX calculates the User Experience by monitoring key metrics within the Virtual Desktop environment, the metrics and their thresholds are shown in the following screen shot:

Machine Experience Indicators						
	Weight (%)	Good		Fair		Poor
Login Delay : Time it takes to login (sec.) 2	20	0	<= 15	<= 60	<= unbounded	
Application Load Time : Avg. startup time for applications (sec.) 2	20	0	<= 10	<= 30	<= unbounded	
CPU Queue Length : Length of CPU queue at inspection time 2	20	0	<= 3	<= 6	<= unbounded	
Page Faults : Number of page faults during inspection interval 2	20	0	<= 2,000	<= 10,000	<= unbounded	
Non-Responding Applications : Number of unresponsive applications at inspection time 2	20	0	<= 2	<= 3	<= unbounded	

I/O Experience Indicators						
	Weight (%)	Good		Fair		Poor
Disk Load : Avg. disk IO per second 2	25	0	<= 25	<= 75	<= unbounded	
Disk Queue Length : Avg. length of disk queue(s) 2	25	0	<= 1	<= 3	<= unbounded	
Network Latency : Avg. network roundtrip time (ms) 2	25	0	<= 150	<= 300	<= unbounded	
Failed Connections : Number of outgoing connection attempts that failed 2	25	0	<= 5	<= 15	<= unbounded	

7.1.3 EqualLogic SAN HQ

EqualLogic SANHQ was used for monitoring the Dell EqualLogic storage units in each bundle. SAN HQ has been used to provide IOPS data at the SAN level; this has allowed the team to understand the IOPS required by each layer of the solution. This report details the following IOPS information;

- Citrix Provisioning Server IOPS (Read IOPS for the vDisks)
- Desktop Write IOPS (WriteCache temporary session data)
- File Server IOPS for User Profiles and Home Directories
- SQL Server IOPS required to run the solution databases
- Infrastructure VM IOPS (the IOPS required to run all the infrastructure Virtual servers)

7.1.4 VMware vCenter


VMware vCenter was used for VMware vSphere-based solutions to gather key data (CPU, Memory and Network usage) from each of the desktop hosts during each test run. This data was exported to .csv files for each host and then consolidated to show data from all hosts. While the report does not include specific performance metrics for the Management host servers, these servers were monitored during testing and were seen to be performing at an expected performance level.

7.1.5 Microsoft Perfmon

Microsoft Perfmon was utilized to collect performance data for tests performed on the Hyper-V platform.

7.2 Performance Analysis Methodology

In order to ensure the optimal combination of end-user experience (EUE) and cost-per-user, performance analysis and characterization (PAAC) on Dell Wyse Datacenter solutions is carried out using a carefully designed, holistic methodology that monitors both hardware resource utilization parameters and EUE during load-testing. This methodology is based on the three pillars shown below. Login VSI is currently the load-testing tool used during PAAC of Dell Wyse Datacenter solutions; Login VSI is the de-facto industry standard for VDI and server-based computing (SBC) environments and is discussed in more detail below.



Each of the pillars shown in the above diagram is discussed in more detail below.

7.2.1 Resource Utilization

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

Resource Utilization Thresholds	
Parameter	Pass / Fail Threshold
Physical Host CPU Utilization	85%
Physical Host Memory Utilization	85%
Network Throughput	85%
Storage IO Latency	20ms

7.2.2 EUE (Tools Info)

Good EUE is one of the primary factors in determining the success of a VDI implementation. As a result of this fact, a number of vendors have developed toolsets that monitor the environmental parameters that are relevant to EUE. PAAC on Dell Wyse Datacenter solutions uses the Liquidware Labs Stratusphere UX tool to ensure that good EUE is delivered for the density numbers defined in our RAs. More specifically, our PAAC analysis uses a scatter plot provided by Stratusphere UX which presents end-user experience for all load-tested users. Stratusphere UX does this by algorithmically combining relevant parameters in relation to virtual desktop machine experience (e.g. login duration) and virtual desktop IO experience (e.g. disk queue length) to provide a plot that shows end-user experience as good, fair or poor using a magic-quadrant type approach.

7.2.3 EUE (Real User Info)

In order to complement the tools based end-user experience information gathered using Stratusphere UX (as described above) and to provide further certainty around the performance of Dell Wyse Datacenter solutions, PAAC on our solutions also involves a user logging into one of the solutions when they are fully loaded (based on the density specified in the relevant RA) and executing user activities that are representative of the user type being tested (e.g. task, knowledge or power user). An example is a knowledge worker executing a number of appropriate activities in Excel. The purpose of this activity is to verify that the end-user experience is as good as the user would expect on a physical laptop or desktop.

7.2.4 Dell Wyse Datacenter Workloads and Profiles

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

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

7.2.5 Dell Wyse Datacenter Profiles

The table shown below presents the profiles used during PAAC on Dell Wyse Datacenter solutions. These profiles have been carefully selected to provide the optimal level of resources for common use cases.

Profile Name	Number of vCPUs per Virtual Desktop	Nominal RAM (GB) per Virtual Desktop	Use Case
Standard	1	2	Task Worker
Enhanced	2	3	Knowledge Worker
Professional	2	4	Power User
Pass-through Graphics	4 + Pass-through GPU	32	Workstation type user e.g. producing complex 3D models.

7.2.6 Dell Wyse Datacenter Workloads

Load-testing on each of the profiles described in the above table is carried out using an appropriate workload that is representative of the relevant use case. In the case of the non-graphics workloads, these workloads are Login VSI workloads and in the case of graphics workloads, these are specially designed workloads that stress the VDI environment to a level that is appropriate for the relevant use case. This information is summarized in the table below:

Profile Name	Workload	OS Image
Standard	Login VSI Light	Shared
Enhanced	Login VSI Medium	Shared
Professional	Login VSI Heavy	Shared + Profile Virtualization
Pass-through Graphics	eFigures / AutoCAD - SPEC Viewperf	Persistent

As noted in the table above, further information for each of the workloads is given below. It is noted that for Login VSI testing, the following login and boot paradigm is used:


- For single-server / single-host testing (typically carried out to determine the virtual desktop capacity of a specific physical server), users are logged in every 30 seconds.

- For multi-host / full solution testing, users are logged in over a period of 1-hour, to replicate the normal login storm in an enterprise environment.
- All desktops are pre-booted in advance of logins commencing.

For all testing, all virtual desktops run an industry-standard anti-virus solution (currently McAfee VirusScan Enterprise) in order to fully replicate a customer environment.

7.2.6.1 Login VSI 3 versus Login VSI 4

PAAC on Dell Wyse Datacenter solutions is currently carried out using Login VSI version 4. However, some previous Dell Wyse Datacenter solutions used Login VSI version 3 for this PAAC work. Login VSI version 3 used a slightly different set of workloads to those used by Login VSI 4 and in order to allow comparison of results obtained using these different Login VSI versions, it is useful to be aware of the information presented in the figure below. This information (for Login VSI medium workload) illustrates higher CPU utilization and lower disk IO for Login VSI 4 (green / upper set of graphs) when compared to Login VSI3 (red / lower set of graphs). The exact variation between these Login VSI versions will vary between environments.



7.2.6.2 Login VSI Light Workload

Compared to the Login VSI medium workload described below, the light workload runs fewer applications (mainly Excel and Internet Explorer with some minimal Word activity) and starts/stops the applications less frequently. This results in lower CPU, memory and disk IO usage.

7.2.6.3 Login VSI Medium Workload

The Login VSI medium workload is designed to run on 2vCPU's per desktop VM. This workload emulates a medium knowledge worker using Office, IE, PDF and Java/FreeMind. The Login VSI medium workload has the following characteristics

- Once a session has been started the workload will repeat (loop) every 48 minutes.
- The loop is divided in four segments; each consecutive Login VSI user logon will start a different segment. This ensures that all elements in the workload are equally used throughout the test.
- The medium workload opens up to five applications simultaneously.
- The keyboard type rate is 160 ms for each character.
- Approximately two minutes of idle time is included to simulate real-world users.

Each loop opens and uses:

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

7.2.6.4 Login VSI Heavy Workload

The heavy workload is based on the medium workload except that the heavy workload:

- Begins by opening four instances of Internet Explorer. These instances stay open throughout the workload loop.
- Begins by opening two instances of Adobe Reader. These instances stay open throughout the workload loop.
- There are more PDF printer actions in the workload.
- Instead of 480p videos a 720p and a 1080p video are watched.
- Increased the time the workload plays a flash game.
- The idle time is reduced to two minutes.

7.2.7 Workloads Running on Shared Graphics Profile

Graphics hardware vendors (e.g. Nvidia) typically market a number of graphics cards that are targeted at different markets. Consequently, it is necessary to provide two shared graphics workloads – one for mid-range cards and the other for high-end cards. These workloads are described in more detail below. It is noted that technologies such as the Citrix / Nvidia vGPU technology (where the Nvidia drivers reside on the virtual desktop, giving shared-level density with near pass-through functionality) mean that in some cases, the higher-end workloads, traditionally used for pass-through GPU PAAC, may be more appropriate for shared GPU PAAC. Such scenarios will explicitly state the workload used.

7.2.8 Workloads Running on Pass-through Graphics Profile

Similarly for pass-through graphics, two workloads have been defined in order to align with graphics cards of differing capabilities.

Mid-Range Pass-through Graphics Workload

The mid-range pass-through graphics workload consists of one desktop running Heaven Benchmark and n-1 desktops running eFigures Advanced Animation activity where n = per-host virtual desktop density being tested at any specific time.

High-End Pass-through Graphics Workload

One desktop running Viewperf benchmark; n-1 desktops running AutoCAD auto-rotate activity where n = per host virtual desktop density being tested at any specific time.

7.3 Testing and Validation

7.3.1 Testing Process

The purpose of the single server testing is to validate the architectural assumptions made around the server stack. Each user load is tested against four runs. A pilot run is conducted to validate that the infrastructure is functioning and valid data is captured. Subsequently three more runs are conducted allowing for correlation of data. Summary of the test results is listed out in the below mentioned tabular format.

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

Login VSI has two modes for launching user's sessions:

- Parallel

Sessions are launched from multiple launcher hosts in a round robin fashion; this mode is recommended by Login Consultants when running tests against multiple host servers. In parallel mode the VSI console is configured to launch a number of sessions over a specified time period (specified in seconds)

- Sequential

Sessions are launched from each launcher host in sequence, sessions are only started from a second host once all sessions have been launched on the first host and this is repeated for each launcher host. Sequential launching is recommended by Login Consultants when testing a single desktop host server. The VSI console is configured to launch a specified number of sessions at a specified interval specified in seconds

All test runs which involved the six desktop hosts were conducted using the Login VSI "Parallel Launch" mode and all sessions were launched over an hour to try and represent the typical 9am logon storm. Once the last user session has connected, the sessions are left to run for 15 minutes prior to the sessions being instructed to logout at the end of the current task sequence, this allows every user to complete a minimum of two task sequences within the run before logging out. The single server test runs were configured to launch user sessions every 60 seconds, as with the full bundle test runs sessions were left to run for 15 minutes after the last user connected prior to the sessions being instructed to log out.

7.4 XenDesktop Test Results

This validation was performed for XenDesktop 7.1 delivering VDI desktops on two x R720 compute hosts, one running vSphere ESXi 5.5 and the other Windows 2012 R2 Hyper-V. Each compute host was outfitted with 256 GB RAM (1600 MHz), dual 3.0 GHz E5-2690v2 processors, and 10 300GB 15K RPM SAS drives in a RAID 10 configuration. Each host comes with Broadcom 5720 4-port NICs running 1Gb Ethernet. Two NICs were connected on each compute host and teamed. All VDI desktops were hosted on the local disk storage on the compute hosts.

The management host was a dedicated R720 running Windows Server 2012 R2 with the same hardware configuration as the compute hosts, with the exception of an added 10GbE dual-port SFP NIC. The management host connects to an EqualLogic PS6010 storage array via a PowerConnect 8024F 10GbE switch to store all of the management server VMs. For each hypervisor scenario, the compute host had its own separate stack of management VMs. The management VMs for each compute host were deployed on the management host using the following configurations:

- 1 VM, 4 vCPU and 8 GB RAM for XenDesktop 7.1 DDC
- 1 VM, 1 vCPU and 2 GB RAM for StoreFront
- 1 VM, 1 vCPU and 2 GB RAM for Licensing Server
- 1 VM, 4 vCPU and 24 GB RAM for PVS 7.1 (powered off during MCS testing)
- 1 VM, 2 vCPU and 8 GB RAM each for VMware vCenter 5.5 or Microsoft System Center 2012 R2
- 1 VM, 2 vCPU and 8 GB RAM for Microsoft SQL Server 2012 SP1

For testing with MCS provisioning, each compute host used a total of 5 management VMs (DDC, Storefront, Licensing, SQL, and vCenter/System Center). For testing with PVS provisioning, a PVS Server was enabled for each compute host.

The Desktop images were built using Windows 8.1 Enterprise and Microsoft Office 2010. For Standard and Enhanced workloads the 32-bit versions were used, but for Professional workloads the 64-bit versions were installed.

The PVS desktops were each configured to boot from vDisk hosted on the PVS Server VM. Each PVS desktop had a local write cache disk attached that was sized according to the formula $1.5 * \text{Memory} + 1 \text{ GB}$. The write cache disks were created to use thin provisioning in all of the test cases below, and were formatted with NTFS. In PVS the vDisks were each configured to use only the target device's local disk for the write cache ("Cache on device hard drive").

The following table illustrates the CPU and memory configuration of the user profiles as tested for both ESXi and Hyper-V.

User Profile	vCPUs	Hyper-V Start up Memory	Hyper-V Minimum Memory	Hyper-V Max Memory	ESXi Memory Reservation	ESXi memory configured
Standard	1	1GB	1GB	2GB	1GB	2GB
Enhanced	2	1.5GB	1GB	3GB	1.5GB	3GB
Professional	2	1.5GB	1GB	4GB	1.5GB	4GB

Windows 8 desktops were configured with some optimizations to enable the Login VSI workload to run and in order to prevent long delays in the login process. Previous experience with Windows 8 has shown that the login delays are somewhat longer than experienced with Windows 7. These were alleviated by performing the following customizations:

- Bypass Windows Metro screen to go straight to the Windows Desktop. This is performed by a scheduled task provided by Login Consultants at logon time
- Disable the **"Hi, while we're getting things ready..."** first time login animation. In random assigned Desktop groups each login is seen as a first time login. This registry setting can prevent the animation and therefore the overhead associated with it.
[HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\System]
"EnableFirstLogonAnimation"=dword:00000000
- McAfee antivirus is configured to treat the Login VSI process VSI32.exe as a low risk process and to not scan that process. Long delays during login of up to 1 minute were detected as VSI32.exe was scanned
- Before closing or sealing the Golden image (either for PVS or MCS), perform a number of logins using domain accounts. This was observed to significantly speed up the logon process for VMs deployed from the Golden image. It is assumed that Windows 8 has a learning process when logging on to a domain for the first time.

Validation was performed using Dell Wyse Solutions Engineering standard testing methodology using LoginVSI 4 load generation tool for VDI benchmarking that simulates production user workloads. Stratusphere UX was not used in this validation since it is not yet compatible with the Windows 8.1 Enterprise.

Each test run was configured to allow for 30 minutes of steady state operations after the initial session launch and login period. This allowed for all sessions to be logged in and active for a significant period of time to allow statistical measuring during steady state operations.

The following table summarizes the steady state test results for the various workloads and configurations. The charts in the sections below show more details about the max values for CPU, memory, IOPS, and Network during the test run.

Hyper-	Provis-	Workload	Tested	Avg CPU %	Avg Mem	Avg	Avg Net
--------	---------	----------	--------	-----------	---------	-----	---------

visor	ioning		Density		Usage GB	IOPS/User	KBps/User
ESX	PVS	Standard	170	95% *	205	5.0	109
		Enhanced	110	86% *	236	8.3	275
		Professional	90	75% *	227	9.1	250
ESX	MCS	Standard	170	92% *	220	4.2	100
		Enhanced	110	83% *	220	7.3	124
		Professional	90	74% *	232	7.0	113
Hyper-V	PVS	Standard	170	92%	229	5.2	121
		Enhanced	110	84%	189	9.3	313
		Professional	90	80%	198	12.3	306
Hyper-V	MCS	Standard	170	88%	228	3.4	98
		Enhanced	110	82%	187	8.5	123
		Professional	90	80%	194	9.5	113

* CPU % for ESX Hosts was adjusted to account for the fact that on Intel E5-2600v2 series processors the ESX host CPU metrics will exceed the rated 100% for the host if Turbo Boost is enabled (by default). E5-2690v2 CPUs rated at 100% utilization in vSphere is 60000 MHz, while actual usage with Turbo Boost has been seen to reach 67000 MHz in some cases. The Adjusted CPU % Usage is based on 100% = 66000 MHz usage and is used in all charts for ESX to account for Turbo Boost. Windows Hyper-V metrics by comparison do not report usage in MHz, so only the reported CPU % usage is used here.

In terms of CPU usage, ESX seems to show less usage at the higher-intensity workloads (Enhanced and Professional desktops use 2 vCPUs) but higher usage at the Standard Workload (1 vCPU). Hyper-V seems to do somewhat better on the Standard Workload with 170 users but worse on the heavier workloads.

Memory usage is a different story however, as many of the ESX test runs approached very close to the memory limit for the host even while CPU was safely below limits. Hyper-V on the other hand shows some room for increased density at the heavier workloads since the Memory usage was significantly lower than the host limit.

Network utilization was very consistent across all tests, showing as expected that PVS required more Network IO per desktop since it was offloading the disk IOPS to the network. Also in the PVS charts below you can clearly see that the network utilization graph resembles the IOPS graph, while in MCS the network graph is clearly shaped differently from IOPS. The only thing of note is that Hyper-V produced 10-15% more network traffic per user in PVS than the same workloads on ESXi.

7.4.1 PVS Write Cache Comparisons (for XenDesktop)

The IOPS results above were surprising in that via the PVS IO offload mechanism, the local storage should be somewhat relieved from a significant number of IOPS. In the chart above however it shows that PVS test cases consistently produced more IOPS per user than equivalent MCS test cases when caching to local disk. To further investigate and optimize we ran additional tests on PVS with some different cache options. We tested caching to local device/ disk and caching to RAM + HD overflow with two different cache sizes. This provides the performance benefits of write cache

in RAM along with the security of overflowing those write caches to disk should the RAM buffers fill. Intermediate buffering was enabled for all ESXi runs leveraging cache in RAM. The results of these tests are summarized in the following table:

Hyper-visor	PVS Cache Type	Workload	Density	Avg CPU %	Avg Mem Usage GB	Avg IOPS/User	Avg Net KBps/User
ESXi	Device HD only	Standard	170	95%	1.2	5	109
ESXi	256MB RAM + Overflow	Standard	170	76%	1.5	0.4	113
ESXi	512MB RAM + Overflow	Standard	170	77%	1.5	0.3	124
ESXi	Device HD only	Enhanced	110	86%	2.1	8	275
ESXi	256MB RAM + Overflow	Enhanced	110	72%	2.2	1.2	284
ESXi	512MB RAM + Overflow	Enhanced	110	73%	2.2	0.2	286
ESXi	HD only, thin provisioned	Professional	90	75%	2.5	9.1	250
ESXi	HD only thick provisioned	Professional	90	79%	2.6	11.7	272
ESXi	256 MB RAM + overflow	Professional	90	61%	2.6	1.9	255
ESXi	512 MB RAM + overflow	Professional	90	64%	2.7	0.3	272

Intermediate buffering was disabled for all Hyper-V runs. The results of these tests are summarized in the following table:

Hyper-visor	PVS Cache Type	Workload	Density	Avg CPU %	Avg Mem Usage GB	Avg IOPS/User	Avg Net KBps/User
Hyper-V	Device HD only	Standard	170	92%	1.3	5.2	121
Hyper-V	256MB RAM + Overflow	Standard	170	78%	1.5	0.3	104
Hyper-V	512MB RAM + Overflow	Standard	170	78%	1.5	0.2	110
Hyper-V	Device HD only	Enhanced	110	85%	1.7	9.3	323
Hyper-V	256MB RAM + Overflow	Enhanced	110	80%	2	0.8	275
Hyper-V	512MB RAM + Overflow	Enhanced	110	81%	2.1	0.4	273
Hyper-V	HD only, thin provisioned	Professional	90	80%	2.2	12.3	306
Hyper-V	HD only thick provisioned	Professional	90	80%	2.2	10.5	308
Hyper-V	256MB RAM + overflow	Professional	90	80%	2.5	2.0	294
Hyper-V	512MB RAM + overflow	Professional	90	79%	2.7	1.4	294

Here you can see that CPU utilization remained almost unchanged while IOPS per user was drastically different. At its best, utilizing PVS cache to RAM reduced per user IOPS load to local disk by nearly **100%** on both ESXi and Hyper-V with an acceptable increase in RAM utilization of 20GB. This clearly shows the benefits of PVS cache to RAM by reducing the required number of IOPS each user would otherwise require.

7.4.2 Recommended Densities

Based on the test results at the selected densities, we propose that it is reasonable to achieve higher densities in some cases. For instance on Hyper-V using MCS for an Enhanced workload with 110 users, CPU usage was 82% in steady state; and only 187 GB (73% of 256 GB) of memory was used. This implies that, all things being equal, a higher density of users at that workload could also be successful. Scaling the CPU usage to 90% yields a potential density of 120 users, and scaling the RAM usage to 90% yields a potential density of 135 users. Taking the lesser of these two densities implies that 120 Enhanced users utilize the host resources without exceeding any limits, so the new recommended user density for that workload is 120 users. In other cases, the recommendation may be lower than the tested density. Applying the similar reasoning to other workloads, while adjusting to maintain a similar pattern of densities

Hyper-visor	Provisioning	Profile	Tested Density	Avg CPU %	Avg Mem Usage GB	Density at 90% CPU	Density at 90% RAM	Recommended Density
ESX	PVS	Standard	170	95%	205	161	191	165 *
		Enhanced	110	86%	236	115	107	110 *
		Professional	90	75%	227	108	91	90
ESX	MCS	Standard	170	92%	220	166	178	165
		Enhanced	110	83%	220	119	115	110 *
		Professional	90	74%	232	109	89	90 *
Hyper-V	PVS	Standard	170	92%	229	166	171	170 *
		Enhanced	110	84%	189	117	134	120 *
		Professional	90	80%	198	101	104	100
Hyper-V	MCS	Standard	170	88%	228	173	171	170
		Enhanced	110	82%	187	120	135	120
		Professional	90	80%	194	101	106	100


*** In these instances we increased the number slightly up or down to maintain the consistency with other similar workloads.**

7.4.3 VMware vSphere 5.5 Test Charts

These sections details the test results charts on an ESXi 5.5 host.


7.4.3.1 PVS Standard User Workload (170 Users)

These graphs show CPU, memory, local disk IOPS, and network results.




7.4.3.2 PVS Enhanced User Workload (110 users)

These graphs show CPU, memory, local disk IOPS, and network results.




7.4.3.3 PVS Professional User Workload (90 users)

These graphs show CPU, memory, local disk IOPS, and network results.




7.4.3.4 MCS Standard User Workload (170 users)

These graphs show CPU, memory, local disk IOPS, and network results.




7.4.3.5 MCS Enhanced User Workload (110 users)

These graphs show CPU, memory, local disk IOPS, and network results.



7.4.3.6 MCS Professional User Workload (90 users)


These graphs show CPU, memory, local disk IOPS, and network results.



7.4.4 Microsoft Windows 2012 R2 Hyper-V Test Charts


7.4.4.1 PVS Standard User Workload (170 Users)

These graphs show CPU, memory, local disk IOPS, and network results.




7.4.4.2 PVS Enhanced User Workload (110 users)

These graphs show CPU, memory, local disk IOPS, and network results.




7.4.4.3 PVS Professional User Workload (90 users)

These graphs show CPU, memory, local disk IOPS, and network results.




7.4.4.4 MCS Standard User Workload (170 users)

These graphs show CPU, memory, local disk IOPS, and network results.




7.4.4.5 MCS Enhanced User Workload (110 users)

These graphs show CPU, memory, local disk IOPS, and network results.



7.4.4.6 MCS Professional User Workload (90 users)

These graphs show CPU, memory, local disk IOPS, and network results.



7.4.5 Local Tier 1 for Blades Test Results

7.4.5.1 Configuration

The purpose of this validation was to investigate the feasibility of providing a XenDesktop MCS based VDI solution on Dell M620 blade servers in Dell M1000E chassis that does not require access to an external storage array. All storage requirements are provided by internal solid state disks installed in the M620 blades. This includes both management and compute storage. External access is only required by the solution for client access to the XenDesktop VMs and for access to other infrastructure such as Active Directory within the environment

This validation was performed for XenDesktop 7.1 delivering VDI desktops on a single M620 Compute server running vSphere ESXi 5.5. The Compute Host was outfitted with 256 GB RAM (1600 MHz), dual 3.0 GHz Ivy Bridge E5-2690v2 processors, and 2 800 MLC SSD drives in a RAID 1 configuration. The compute host connected to the SSD drives through a Dell H310 PERC controller

In the solution validated the management host was a separate M620 with 2 * 15K 300GB disks in a RAID 1 configuration connected by a PERC H310. However, for a full solution it is envisaged that there would not be a dedicated management host in the blade chassis. Instead, the management components will run on any of the Compute nodes. Each node is identically configured in the same way as the compute node above. Based on the latencies observed during the Login VSI tests on the compute nodes the solution should scale well to accommodate up to 16 blades. In addition the Virtual Center appliance was deployed with sufficient capacity for up to 20 hosts and 2000 VMs.

The M1000E blade chassis was fitted with a Dell MIO aggregator. This supported the external communication from the chassis for VDI and Active directory access, and provided support for all the internal communication between the various VMs hosted on the blades in the chassis. A simple management configuration was implemented using 3 management VMs only, a vSphere Virtual Center appliance, Citrix XenDesktop DDC, and Citrix Storefront server and license server combined. SQL Server 2008 Express was installed on the DDC to support XenDesktop

- 1 VM, 2 vCPU and 4 GB RAM for XenDesktop 7.1 DDC plus 30 GB system disk.
- 1 VM, 1 vCPU and 4 GB RAM for Storefront and Licensing server
- 1 VM, 2 vCPU and 8 GB RAM for Virtual Center Appliance, plus 30GB system disk plus 100GB disk for vCenter database

So for testing with MCS provisioning, each Compute host used a total of 3 management VMs (DDC, Storefront and Licensing, and vCenter). For testing with PVS provisioning, a PVS Server was enabled for each Compute Host.

The Desktop images were built using Windows 8.1 Enterprise and Microsoft Office 2010. For Standard and Enhanced workloads the 32-bit versions were used, but for Professional workloads the 64-bit versions were installed.

7.4.5.2 Test Results

Validation was performed using Dell Wyse Solutions Engineering standard PAAC testing methodology using LoginVSI 4 load generation tool for VDI benchmarking that simulates production user workloads. Stratusphere UX was not used in this validation since it is not compatible with the desktop OS, Windows 8.1 Enterprise.

Each test run was configured to allow for 15 minutes of steady state operations after the initial session launch and login period, so that all sessions is logged in and active for a significant period of time so that statistics could be measured during steady state to simulate daily operations.

The following table summarizes the steady state test results for the various workloads and configurations. The charts in the sections below show more details about the max values for CPU, memory, IOPS, and Network during the test run.

Hyper-visor	Provis- ioning	Workload	Density	Avg CPU %	Avg Mem Usage GB	Max IOPs/User during login	Avg IOPS/User steady state	Avg Net KBps/User
ESX	MCS	Standard	170	85% *	250	10.5	7	300
		Enhanced	110	83% *	230	18	7.4	325
		Professional	90	91% *	240	31	8.3	489
		Professional 0% preboot	90	91	240	31	8.3	490

*** CPU % for ESX Hosts was adjusted to account for the fact that on Intel E5-2600v2 series processors the ESX host CPU metrics will exceed the rated 100% for the host if Turbo Boost is enabled (by default). With E5-2690v2 CPUs the rated 100% in vSphere is 60000 MHz usage, while actual usage with TB has been seen to reach 67000 MHz in some cases. The Adjusted CPU % Usage is based on 100% = 66000 MHz usage and is used in all charts for ESX to account for Turbo Boost.**





Many of the ESX test runs approached very close to the memory limit for the host even while CPU was safely below limits. ESXi reports however show Active memory usage much lower than consumed with very little memory ballooning in evidence

Network utilization was very consistent across all tests. User profiles were stored on a file server external to the test environment. This accounts for some extra network utilization. However, even for a fully populated bladed with 16 * 170 Standard users the network requirement is 800Mbps. The installed Dell IO aggregators provide 10Gbps connectivity to the external environment so provide plenty of capacity. It is also possible to support the profile storage within the chassis.





Disk latency on the SSD disks remained very low at 0 milliseconds on all the tests. Even during the boot storm when 170 VMs were powered on at approximately the same time the disk latency did not rise above 15 milliseconds.

An addition test was run against the Professional image with none of the desktops prebooted. So, for each session launched the VM had to be powered on before the session started and the test could begin. It is interesting to note that the results are almost identical to the professional test with all desktops powered on. Additionally, the SSD disk latency stayed at 0 ms for the duration of the test.

The SSD volumes provided enough space including the ESXi swap file for all the VMs in for 170 VMs. With the VMs powered off the space usage was as follows. The SSDs also hosted the Windows 32 bit and 64 images required to create the MCS catalogs

General		Capacity	
Location:	ds:///vmfs/volumes/52e64d7a-0eacf0...		Refresh
Type:	VMFS	Capacity:	744.50 GB
Number of Hosts Connected:	1	Provisioned Space:	652.86 GB
Virtual Machines and Templates:	173	Free Space:	579.18 GB
		Last updated on:	05/02/2014 11:54:17
Commands		Storage Capabilities	
	Refresh		Refresh
	Enter SDRS Maintenance Mode	System Storage Capability:	N/A
	Browse Datastore...	User-defined Storage Capability:	N/A
	Assign User-Defined Storage Capability		

When all the VMs are powered on the swap files reduce the free space to 224GB as follows:

General		Capacity	
Location:	ds:///vmfs/volumes/52e64d7a-0eacf0...		Refresh
Type:	VMFS	Capacity:	744.50 GB
Number of Hosts Connected:	1	Provisioned Space:	5.87 TB
Virtual Machines and Templates:	172	Free Space:	224.23 GB
		Last updated on:	05/02/2014 12:05:57
Commands		Storage Capabilities	
	Refresh		Refresh
	Enter SDRS Maintenance Mode	System Storage Capability:	N/A
	Browse Datastore...	User-defined Storage Capability:	N/A
	Assign User-Defined Storage Capability		

The SQL Express data file usage was 112MB. Adding a further 15 hosts should not increase this to more than 1.75GB. With the allocation used in the tests there was 13.7GB free on the DDC so there is room for expansion of the database as further hosts and VMs are added. Space usage on the Management Datastore was as follows. This also includes space used by Windows 8.1 32 bit and 64 bit templates to provision the XenDesktop VMs. It is intended however that the Management node should also be provisioned on 800GB SSDs

Getting Started		Summary	Virtual Machines	Hosts	Performance	Configuration	Tasks & Events	Alarms	Permissions
<div> <div> General </div> <div> Location: ds:///vmfs/volumes/52d3d7fd-c036c6... Type: VMFS Number of Hosts Connected: 1 Virtual Machines and Templates: 7 </div> </div>									
<div> <div> Capacity </div> <div> <div>Refresh</div> Capacity: 278.75 GB Provisioned Space: 381.07 GB Free Space: 10.23 GB Last updated on: 05/02/2014 12:05:38 </div> </div>									
<div> <div> Commands </div> <div> <div>Refresh</div> <div>Enter SDRS Maintenance Mode</div> <div>Browse Datastore...</div> <div>Assign User-Defined Storage Capability</div> </div> </div>									
<div> <div> Storage Capabilities </div> <div> <div>Refresh</div> System Storage Capability: N/A User-defined Storage Capability: N/A </div> </div>									

Smoke further tests were also performed to investigate the capability of the local SSD volumes within the VDI solutions:-

- Generate maximum IOPs throughput using IOMETER on a 20 GB disk attached to a Windows 8.1 VM
- Add extra IOPS (22) within each VSI session for a standard Login VSI test to produce a total of 8000 IOPs on the host during the VSI test. Measure the effect on the latency of the SSD volume
- Perform a logon storm with 170 Standard VSI sessions logging on within a 10 minute interval instead of sessions logging on at 30 second intervals. Measure the effect of the increased logon rate on the SSD volumes


The results of each of these tests are document in sections 7.4.4.5, 7.4.4.6, and 7.4.4.7 respectively.

7.4.5.3 VMware vSphere 5.5 Test Charts

These sections details the test results charts on an ESXi 5.5 host.

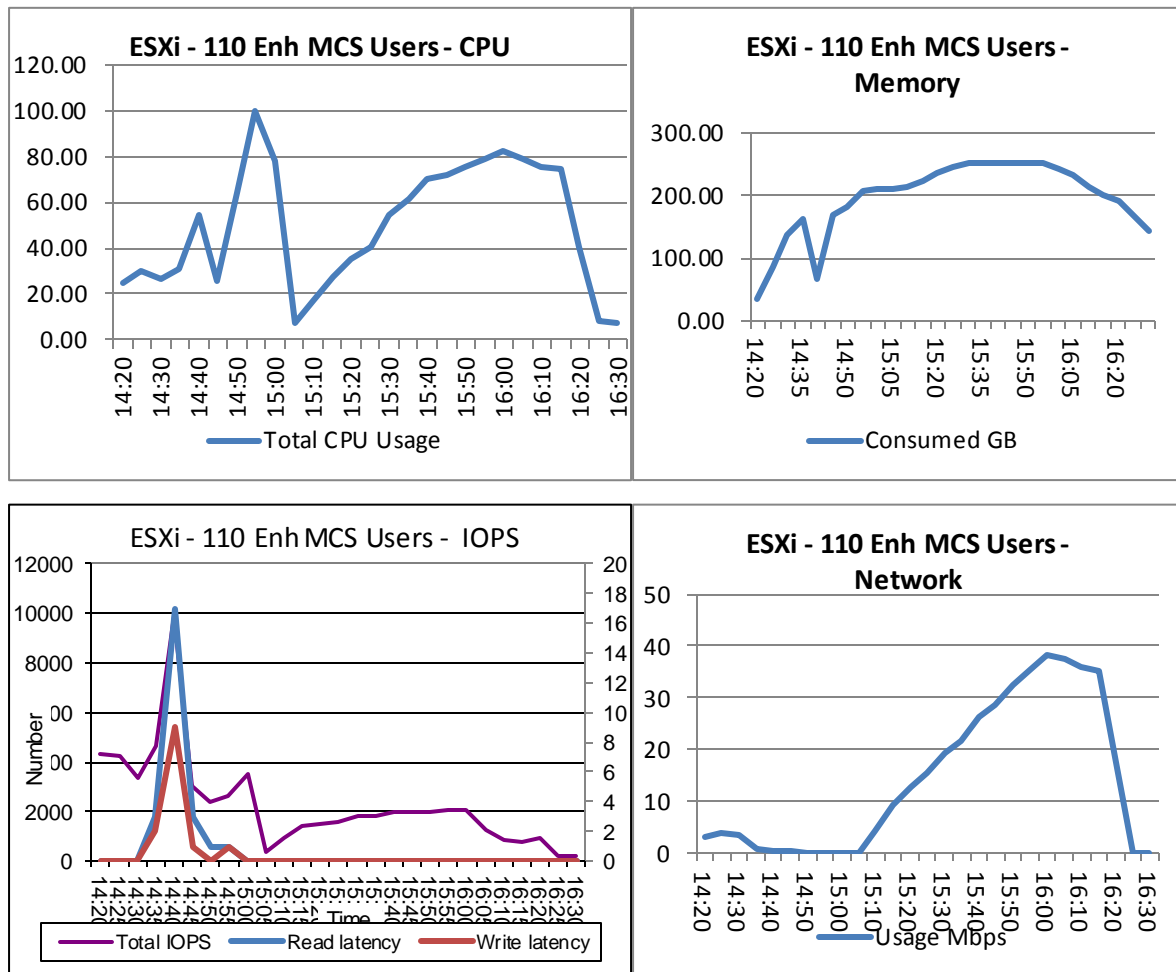
7.4.5.4 MCS Standard User Workload (170 Users)

These graphs show CPU, memory, local disk IOPS, and network results.



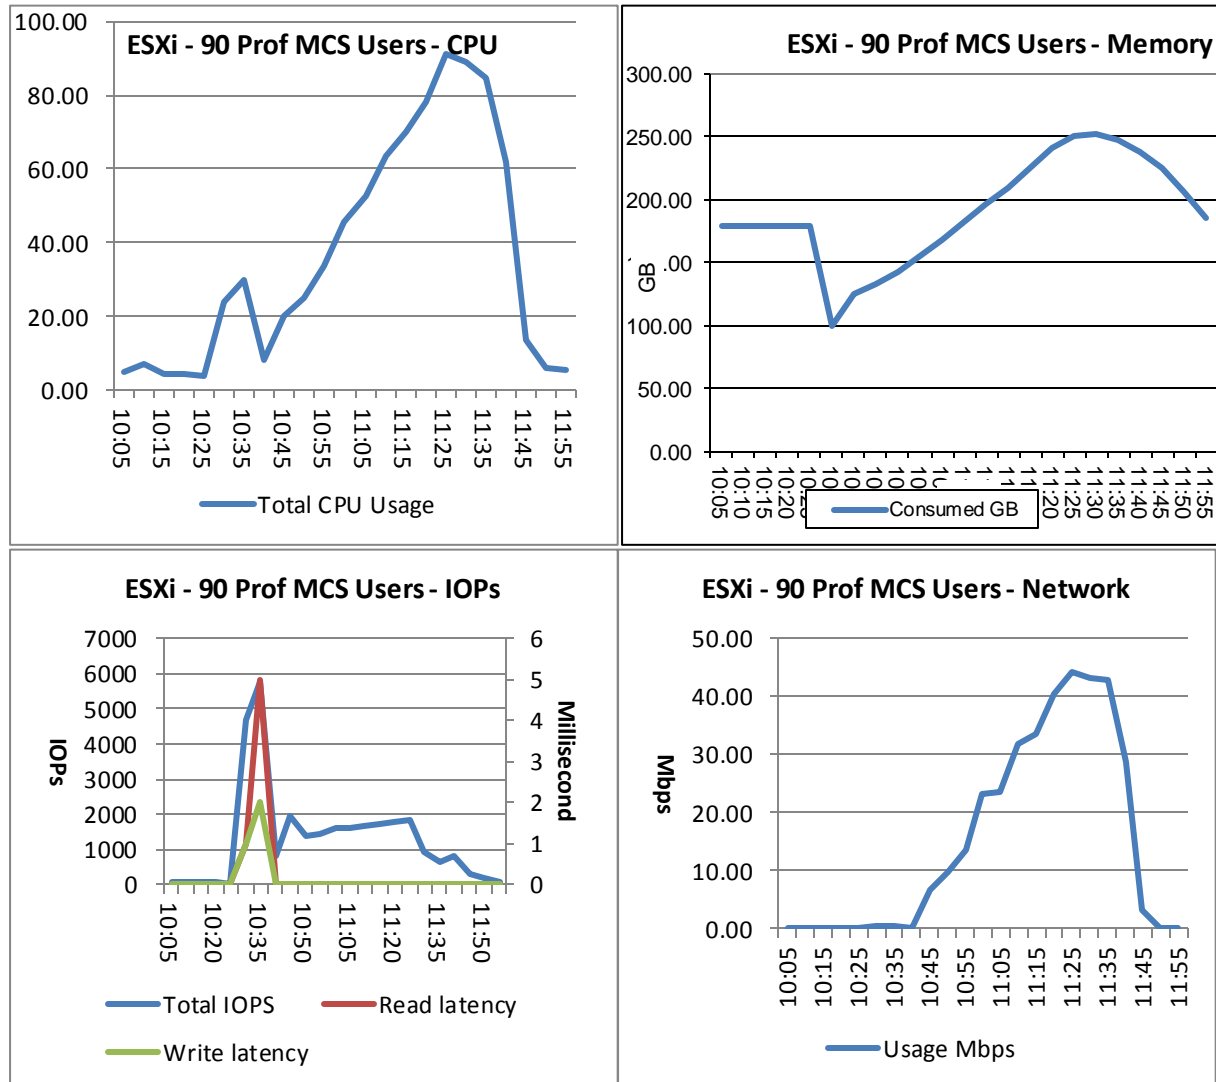
7.4.5.5 MCS Enhanced User Workload (110 users)

These graphs show CPU, memory, local disk IOPS, and network results.




7.4.5.6 MCS Professional User Workload (90 users)

These graphs show CPU, memory, local disk IOPS, and network results.



7.4.5.7 MCS Professional User Workload (90 users) with 0% preboot

These graphs show CPU, memory, local disk IOPS, and network results.



7.4.5.8 IOMeter test against SSD disks

These graphs show IOMETER results against the SSD disks using the following IOMETER profile. This was performed on a single VM on the compute host using a 20GB disk configured on the SSD volume. The test was run to gather information that would indicate the overall capability of the SSD volume when hosting VDI VMs. The access specifications were configured as proposed by Jim Moyle of Atlantis at the following [location](#).

Edit Access Specification

Name: Default Assignment:

Size	% Access	% Read	% Random	Delay	Burst	Alignment	Reply
0MB 4KB 0B	100	20	80	0	1	0MB 128KB 0B	none

Insert Before
Insert After
Delete

Transfer Request Size: Megabytes Kilobytes Bytes

Percent of Access Specification: Percent

Percent Read/Write Distribution: Write Read

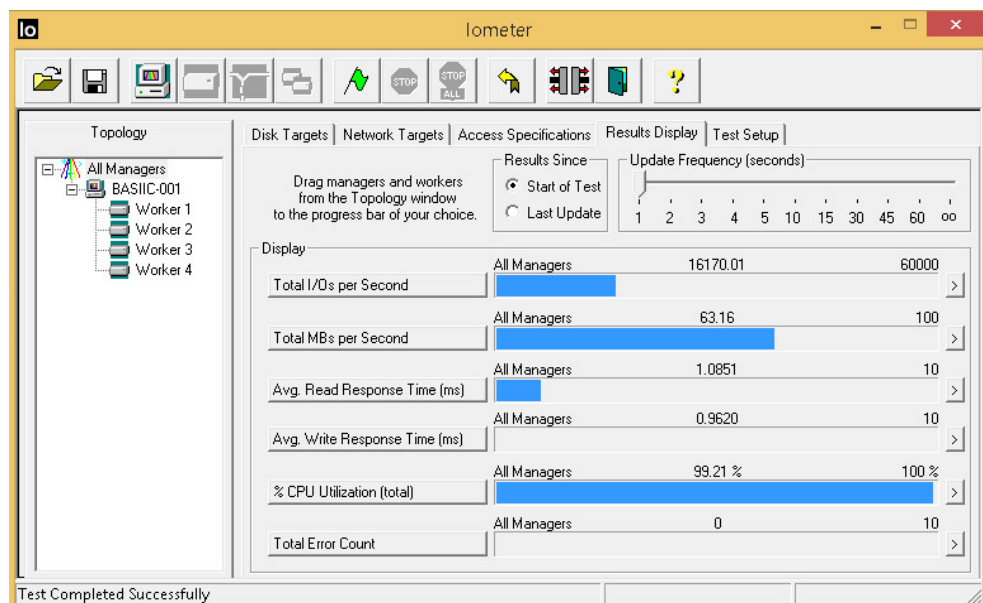
Percent Random/Sequential Distribution: Sequential Random

Burstiness: Transfer Delay: ms Burst Length: I/Os

Align I/Os on: ☐ Sector Boundaries ☒ Megabytes Kilobytes Bytes

Reply Size: ☒ No Reply ☐ Megabytes Kilobytes Bytes


OK Cancel



It can be seen that at 16000 IOPs with 80% write the latency remains at approx. 1 millisecond. While IOMETER test scenarios are very subjective and it is arguable whether the test specification used above truly represents a VDI workload, it does represent a reasonable approximation and the result of 16000 IOPs available on the SSD volume does indicate that it is possible for the solution to support much more IO intensive workloads as indicated in the following graph


Number of Desktops Per Server	Approximate Max IOPS Per Desktop
170 Standard	95 IOPS
110 Enhanced	145 IOPS
90 Premium	180

7.4.5.9 MCS Standard User Workload (170 Users) + extra IOPs generated by IOMETER



The login script launched an automated Iometer workload in each VDI session to generate extra IOPs within the session. This generated sustained IOPS of approximately 8000. As can be seen from the graphs above the latency of the local tier 1 storage on the SSD volumes did not rise above 2 milliseconds which is well within the bounds that are considered appropriate for VDI performance (i.e. < 20ms)

7.4.5.10 MCS Standard User Workload (170 Users) Logon storm



During the logon period of 10 minutes the Read and Write latency on the SSD volume rose to 17 milliseconds but quickly dropped to 0 once the sessions reached steady state. A latency of 17 seconds is still below a value that could affect VDI performance. The shortened logon interval produced a peak CPU of 100% usage but this dropped off once the steady state was reached.

7.4.6 XenDesktop with Personal vDisk Enabled

The Personal vDisk technology provides a savings in storage by leveraging a smaller private vDisk versus using a full dedicated desktop image. The impact on storage was found to be minimal. While evaluating the desktop density per server, server CPU was monitored closely. Based on the use of a low level filter driver to determine and then write user/application data, a CPU hit was expected and validated. The overall impact on storage both Tier 1 and Tier 2 is approximately 2 IOPS per user. It was therefore determined that on the same R720 server configuration where 145 basic users are supported, 116 basic user sessions are supported with PvDisk enabled. Also it was determined that on an M620 blade server configuration where 135 basic users are supported, 108 basic user sessions are supported with PvDisk enabled.

7.4.6.1 XenDesktop Scaling with PvDisk enabled on the R720 (Win7- Sandy Bridge)

Basic Users	Basic Users w/PvDisk	Standard Users	Standard Users w/PvDisk	Premium Users	Premium Users w/PvDisk	Physical Mgmt Servers	Physical Host Servers
145	116	116	93	95	76	1	1
1000	800	800	640	655	524	1	7
2000	1600	1600	1280	1310	1048	2	14
3000	2400	2400	1920	1966	1572	3	21
4000	3200	3200	2560	2621	2097	3	28
5000	4000	4000	3200	3276	2621	4	35
6000	4800	4800	3840	3931	3145	5	42
7000	5600	5600	4480	4586	3669	5	49
8000	6400	6400	5120	5241	4193	6	56
9000	7200	7200	5760	5897	4717	7	63
10000	8000	8000	6400	6552	5241	8	69

7.4.7 EqualLogic PS6210XS Hybrid Array Testing Summary


The following tests were conducted using a VMware View solution stack but demonstrate the capability of the array given the following common VDI events:

1. **Boot storm:** Boot storms represent the worst-case scenario where many virtual desktops are powered on at the same time and they all contend for the system resources simultaneously. This test was used to evaluate if the storage array hosting the desktops was capable of handling huge spikes in storage I/O without causing significant impact on other services.
2. **Login storm:** Login storms also represent a heavy I/O scenario where many users are logging into their virtual desktops at the beginning of a work day or a shift at the same time. In this test, all the desktops were pre-booted and left in an idle state for more than 20 minutes to let their I/O settle before running the Login VSI Medium workload to simulate users logging in to their virtual desktops.
3. **Steady state workload for standard users:** Once the login storm for the previous test is completed, the Login VSI Medium workload is allowed to run for at least one hour to simulate the real world scenario of users performing their daily tasks. Stratusphere UX is used to capture user and machine performance during testing.

The following sections provide results from the boot storm, login storm, and steady state testing for a standard LoginVSI user on a single EqualLogic PS6210XS hybrid storage array.

Boot storm I/O


To simulate a boot storm, the virtual desktops were reset simultaneously from the VMware vSphere client. The PS6210XS array delivered 17,514 IOPS (8.7 IOPS per VM) with less than 3 ms average latency under the peak load during this test and all 2,000 desktops were available in about 25 minutes.



SAN HQ data showing EqualLogic PS6210XS array IOPS and latency during boot storm

The above screenshot represents the boot storm for 2,000 VMware Linked Clone VMs with a read/write ratio of 65% read and 35% write. The Replica volumes contributed to the majority of the I/O. Latency was extremely low at a weighted average of 2.78ms.

Due to the number of VMs being powered on, each replica generated maximum IOPS at different times in the boot storm. Figure 3 below shows two Replica volumes generating the majority of I/O operations at peak of the boot storm. The IOPS recorded are predominantly read operations.



Read and write IOPS on the replica volumes

The remaining I/O operations were handled by the linked clone volumes with an average ratio of 40% read and 60% write operations. The table below shows that the majority of the I/O operations were handled by the SSDs during the boot storm.

Table 7 Disk usage on EqualLogic PS6210XS during boot storm

Member	Disk	Description	Average IOPS	Read I/O Rate	Write I/O rate	Status
IXT-6210XS	0	SSD 800GB SAS	3,025.16	31.81 MB/sec	25.79 MB/sec	online
IXT-6210XS	1	SSD 800GB SAS	3,030.44	31.87 MB/sec	25.80 MB/sec	online
IXT-6210XS	2	SSD 800GB SAS	3,003.07	31.78 MB/sec	25.85 MB/sec	online
IXT-6210XS	3	SSD 800GB SAS	3,045.40	31.77 MB/sec	25.84 MB/sec	online
IXT-6210XS	4	SSD 800GB SAS	3,010.06	31.65 MB/sec	25.85 MB/sec	online
IXT-6210XS	5	SSD 800GB SAS	3,020.25	31.86 MB/sec	25.76 MB/sec	online
IXT-6210XS	6	SSD 800GB SAS	3,030.88	31.70 MB/sec	25.75 MB/sec	online
IXT-6210XS	7	10K 1TB SAS	79.31	2.20 MB/sec	541.45 KB/sec	online
IXT-6210XS	8	10K 1TB SAS	49.67	1.96 MB/sec	516.03 KB/sec	online
IXT-6210XS	9	10K 1TB SAS	47.79	2.08 MB/sec	542.55 KB/sec	online
IXT-6210XS	10	10K 1TB SAS	47.06	1.91 MB/sec	501.67 KB/sec	online
IXT-6210XS	11	10K 1TB SAS	46.34	1.99 MB/sec	540.34 KB/sec	online
IXT-6210XS	12	10K 1TB SAS	45.01	1.91 MB/sec	501.67 KB/sec	online
IXT-6210XS	13	10K 1TB SAS	43.85	1.88 MB/sec	512.72 KB/sec	online
IXT-6210XS	14	10K 1TB SAS	42.38	1.87 MB/sec	510.51 KB/sec	online
IXT-6210XS	15	10K 1TB SAS	42	1.85 MB/sec	509.40 KB/sec	online
IXT-6210XS	16	10K 1TB SAS	42.13	1.85 MB/sec	509.40 KB/sec	online
IXT-6210XS	17	10K 1TB SAS	43.32	1.85 MB/sec	509.40 KB/sec	online
IXT-6210XS	18	10K 1TB SAS	45.2	1.87 MB/sec	511.61 KB/sec	online
IXT-6210XS	19	10K 1TB SAS	43.53	1.87 MB/sec	507.19 KB/sec	online
IXT-6210XS	20	10K 1TB SAS	44.08	1.87 MB/sec	509.40 KB/sec	online
IXT-6210XS	21	10K 1TB SAS	44.3	1.91 MB/sec	499.46 KB/sec	online
IXT-6210XS	22	10K 1TB SAS	45.26	1.91 MB/sec	499.46 KB/sec	online
IXT-6210XS	23	10K 1TB SAS	0	0 KB/sec	0 KB/sec	spare

Storage network utilization was well within the available bandwidth. The peak network utilization during the boot storm reached approximately 6% of the total storage network bandwidth and then gradually declined once all the VMs were booted up. There were also no retransmissions on the iSCSI SAN.

The above results show that EqualLogic PS6210XS hybrid array can handle a heavy I/O load like a boot storms in a VDI environment with no issues.

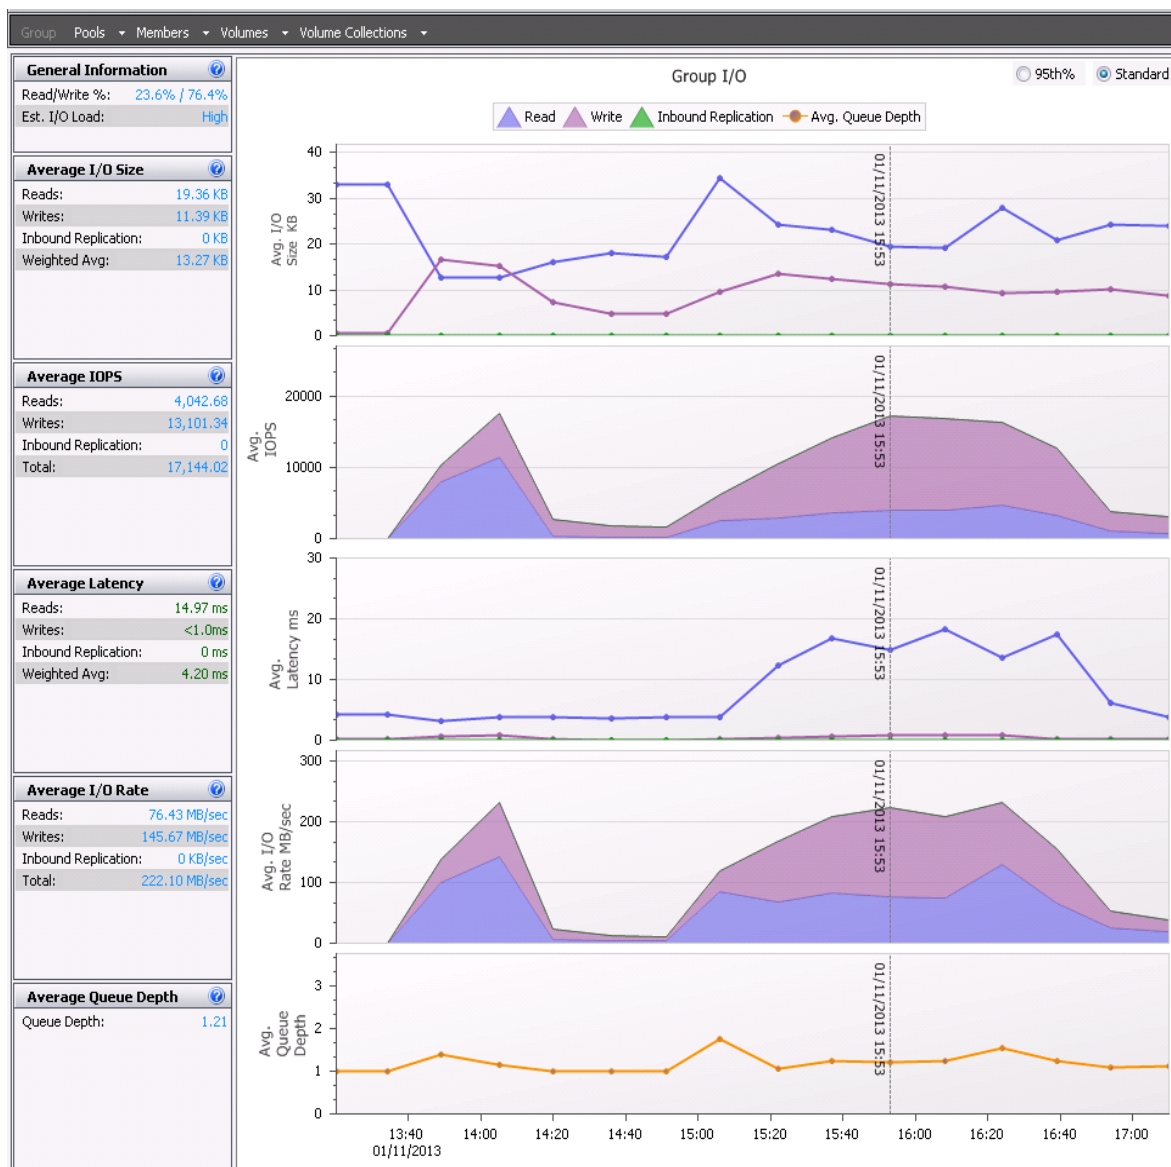
Login storm I/O

Login VSI was configured to launch 2,000 virtual desktops over a period of about 30 minutes after pre-booting the virtual desktops. The peak IOPS observed during the login storm was about 17,144 IOPS (8.5 IOPS per VM).

Login storms generate significantly more write IOPS than a boot storm due to multiple factors including:

- User profile activity
- Starting operating system services on the desktop
- First launch of applications

Once a virtual desktop has achieved a steady state after user login, the Windows 7 OS has cached applications in memory and does not need to access storage each time the application is launched. This leads to lower IOPS during the steady state. Figure below shows the IOPS and latency observed during the login storm.



SAN HQ data showing login storm I/O

EqualLogic PS6210XS easily handled logging in 2,000 sessions in a short time delivering the required 17,144 IOPS with 4.2 ms of average latency at the peak of the login storm. The table below shows the overall disk usage in the array during the login storm.

Disk usage on EqualLogic PS6210XS during login storm

Member	Disk	Description	Average IOPS	Read I/O Rate	Write I/O Rate	Status
IXT-6210XS	0	SSD 800GB SAS	4,945.26	52.19 MB/sec	52.61 MB/sec	online
IXT-6210XS	1	SSD 800GB SAS	4,948.01	52.14 MB/sec	52.68 MB/sec	online
IXT-6210XS	2	SSD 800GB SAS	4,961.39	52.29 MB/sec	52.78 MB/sec	online
IXT-6210XS	3	SSD 800GB SAS	4,937.53	52.33 MB/sec	52.79 MB/sec	online
IXT-6210XS	4	SSD 800GB SAS	4,940.79	52.25 MB/sec	52.80 MB/sec	online
IXT-6210XS	5	SSD 800GB SAS	4,934.91	52.24 MB/sec	52.72 MB/sec	online
IXT-6210XS	6	SSD 800GB SAS	4,932.05	52.09 MB/sec	52.55 MB/sec	online
IXT-6210XS	7	10K 1TB SAS	65.15	2.50 MB/sec	752.52 KB/sec	online
IXT-6210XS	8	10K 1TB SAS	67.55	2.36 MB/sec	745.92 KB/sec	online
IXT-6210XS	9	10K 1TB SAS	63.98	2.48 MB/sec	748.12 KB/sec	online
IXT-6210XS	10	10K 1TB SAS	65.68	2.25 MB/sec	734.92 KB/sec	online
IXT-6210XS	11	10K 1TB SAS	62.22	2.35 MB/sec	740.42 KB/sec	online
IXT-6210XS	12	10K 1TB SAS	60.11	2.23 MB/sec	736.02 KB/sec	online
IXT-6210XS	13	10K 1TB SAS	59.46	2.19 MB/sec	750.32 KB/sec	online
IXT-6210XS	14	10K 1TB SAS	58.56	2.18 MB/sec	745.92 KB/sec	online
IXT-6210XS	15	10K 1TB SAS	57.25	2.15 MB/sec	737.12 KB/sec	online
IXT-6210XS	16	10K 1TB SAS	57.18	2.14 MB/sec	732.72 KB/sec	online
IXT-6210XS	17	10K 1TB SAS	56.71	2.13 MB/sec	739.32 KB/sec	online

IXT-6210XS	18	10K 1TB SAS	56.9	2.13 MB/sec	736.02 KB/sec	online
IXT-6210XS	19	10K 1TB SAS	58.48	2.14 MB/sec	742.62 KB/sec	online
IXT-6210XS	20	10K 1TB SAS	58.27	2.13 MB/sec	739.32 KB/sec	online
IXT-6210XS	21	10K 1TB SAS	57.64	2.16 MB/sec	717.31 KB/sec	online
IXT-6210XS	22	10K 1TB SAS	57.75	2.16 MB/sec	714.01 KB/sec	online
IXT-6210XS	23	10K 1TB SAS	0	0 KB/sec	0 KB/sec	spare

The table above shows that most of the login storm I/O is handled by SSDs and therefore the array is able to provide the best possible performance. Each SSD handled approximately 4,950 IOPS at the peak of login storm; the average latency was very low during the entire login storm time period, and the array clearly demonstrated its ability to handle the workload.

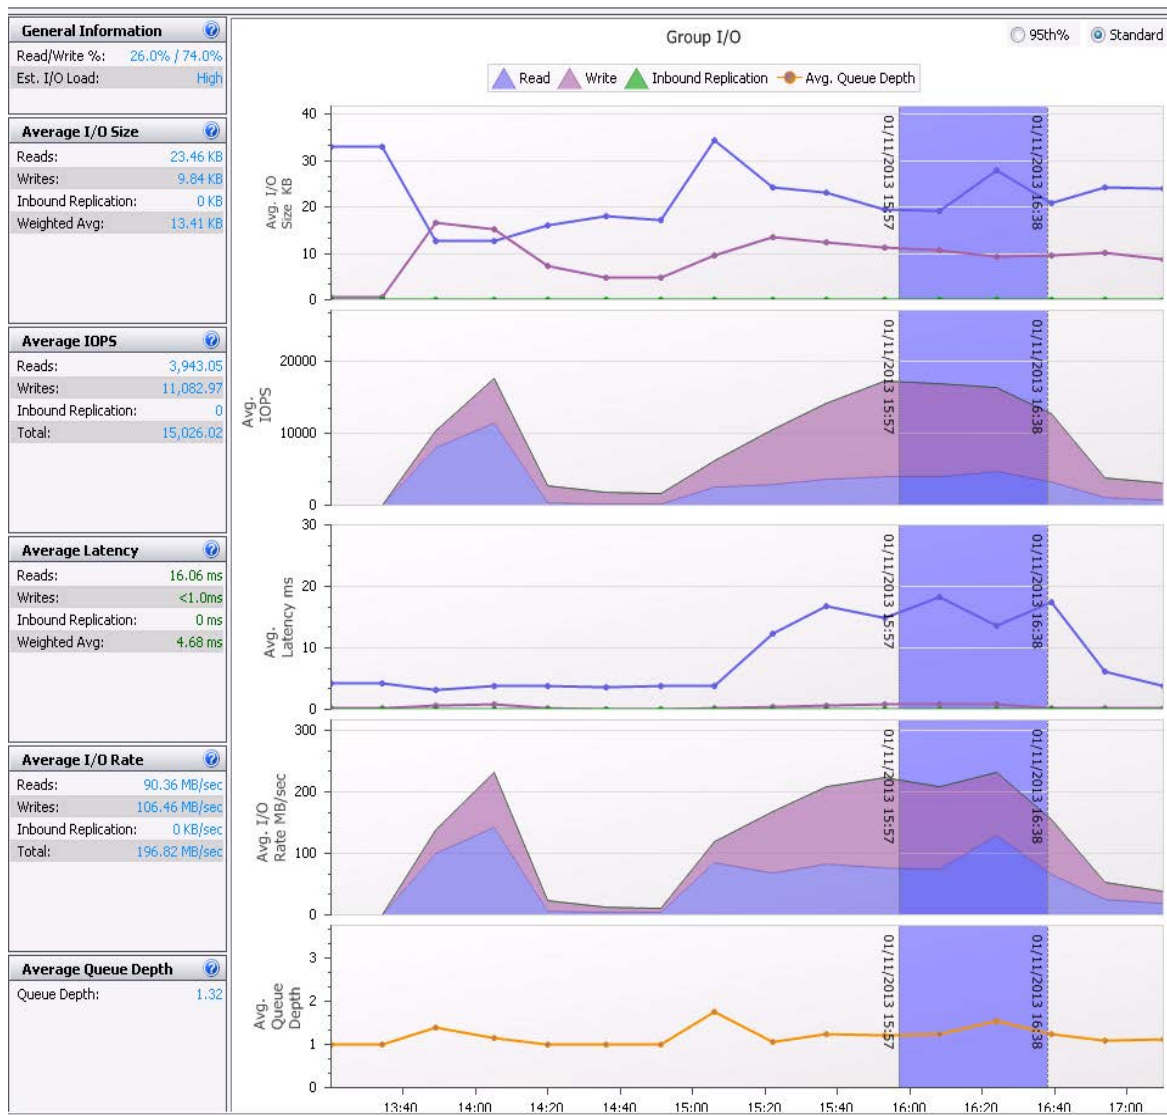
Steady state I/O

Following the completion of the login storm, the I/O profile changed to approximately 24% Read and 76% Write I/O operations at steady state. The total IOPS required during the peak load at steady state with all the users logged in was around 16,773 (8.3 IOPS per VM). EqualLogic PS6210XS delivered these IOPS with 5 ms average latency, which is well below the 20 ms threshold. The average load during the entire steady state test period was approximately 15,000 IOPS, which EqualLogic PS6210XS delivered with 4.7 ms average latency as shown in Figure 5.

All changes that occur on the virtual desktop (including temporary OS writes such as memory paging) are being written to disk. The I/O pattern is mostly writes due to this activity. Once the desktops are booted and in a steady state, the read I/O becomes minimal due to Horizon View Storage Accelerator enabling content based read caching (CBRC) on the ESXi hosts.


During steady state there is minimal activity on the replica volume and most of the activity is seen on the VDI-Images volumes that host the virtual desktops.

The figure below shows the performance of the array during the steady state test.



Based on the criteria in Section 4.3.3, Figure 9 shows Stratusphere UX scatter plot for user experience. Virtually all the users are in the Good category. This shows that the EqualLogic PS6210XS array is capable of easily providing adequate user experience for all the users in the VDI environment.

Note that Stratusphere UX was able to register 1,860 of the total 2,000 desktops for the scatter plot in the figure below, but the Good performances of the remaining ones were verified in other reports.



7.4.8 Dell Compellent Storage Testing Summary

The Dell Compellent MSV Solutions team tested and validated the fibre channel-based Dell Compellent storage with the following configuration and criteria:

- 2000 XenDesktop VMs
- 0% pre-booted desktop images ([all 2000 desktops booted on demand](#))
- VSI Medium workload to generate maximum IOPS demand.

Compellent array configuration as tested:

Storage Role	Type	Quantity	Description
Controllers	SC8000	2	SCOS 6.3
Enclosures	SC220	4	24 bay – 2.5" disk drive enclosure
Ports	FC – 8 Gb	16	Front end host connectivity
	SAS – 6 Gb	4	Back end drive connectivity
Drives	300 GB 15K RPM	96	92 Active with 4 hot spares

The results show the performance of a Dell Compellent SC8000 under load. The screen captures come directly from the Dell Compellent Charting viewer, which can display real-time load from the storage perspective. A subset of all data is presented here for clarity, with each of the screen captures showing the load from a different layer of the storage architecture.

Performance charts

The performance charts are captured from real-time data, averaged on 10 second intervals. This granularity will provide accurate results based on the volume of data and the short data collection interval.

Front end performance


Front-end performance is measured as time to complete an I/O operation, from the time an I/O is initiated from a host, until the response is sent that the I/O is complete.

Boot storm

The boot storm shows the ramp up in load from the time the first machine starts the boot process, until the last machine displays the Windows login screen. Additional I/O will continue to be generated, but at this point all machines are available for login.

In a customer environment this would show how long a cold state boot would take from a storage perspective. This does not include host start-up times, or the time required for the desktop broker to initiate the connections to the vSphere hosts and start the power-on process.

Figure 1 – Boot storm:



The peak IOPS load was about 13,000 IOPS. The write ratio is about 95% as the majority of the boot data is read from the PVS streaming service. Even under peak load conditions, the average write latency was approximately 5ms, with a corresponding average read latency of approximately 2ms.

Login storm

The Login storm displays the load ramp-up from the time the first workstation starts the login process, until the last workstation has completed login and starts the testing workload. Over approximately a 50 minute period, all virtual desktops are logged in and have started processing the assigned tasks.

This simulates the time required for all users to login, such as the start of user workday. It would also simulate performance if a complete VDI environment required a restart during a workday.

Figure 2 – Login Storm:




During the login window, the storage load peaks at approximately 22,000 IOPS. Through this peak, the average write latency measured 5 ms, with the corresponding read latency averaging at 7 ms. This results in a very smooth login process with very good overall response times inside of the virtual desktops. Since the industry considers any storage latency under 20 ms to offer acceptable performance, this low latency configuration results in an excellent user experience.

Steady state VDI load

The steady state load shows a continuous 2,000 user workload simulating all the users performing their daily duties. During the test run the steady state load averaged about 16,250 IOPS (about 8 IOPS/user), with an average write latency of 6 ms, and an average read latency of 7 ms.

Figure 3 – Steady State:



Storage Scaling

Due to the nature of spinning drives, sizing a solution are done making linear scaling assumptions. For example, if 92 drives can support 2000 users at a given I/O load, sizing for 1000 or 3000 users would require 46 or 138 active drives respectively. Along those lines, once a ratio of spindles per user has been established, expanding the user base as it grows can save an upfront purchase of the total needed drives. For example, if 500 users are deployed in each phase of the rollout, the linear math suggests that if the users being added share the same I/O profile, that adding 23 drives per phase is required to handle the desktop's performance.

Summary

In conclusion, a 2000 desktop Citrix XenDesktop based VDI solution are deployed on a Dell Compellent Storage Center platform comprised of two SC8000 and four SC220 enclosures with 96

– 300 GB 15K drives (92 active and 4 spares) for the tested workload. This mimics a typical knowledge worker as represented by the Login VSI Medium workload.

It is highly recommended to run a pilot to determine the key disk metrics typical of a particular set of users to be virtualized in the environment. Knowing the number of IOPS per user, read and write ratios, average transfer sizes, and how much data each user consumes, will help prevent under or oversizing a storage solution for a virtual desktop deployment.

Keep in mind that the performance based solution shown within this document is architected for non-persistent desktops, and thus, due to the transient nature of the data, array based snapshots on the operating system volumes were not utilized. In any final solution based on this architecture, additional capacity for user profile data and user files will need to be considered for the file servers within the infrastructure. Also if desired, additional storage may also be added to the array to allow for a snapshot schedule to store any desktop data deemed important.

7.4.8.1 Dell Compellent Storage Test Result Summary Tables

User Workload	Citrix XenDesktop Sessions
Basic User	2000
Standard User	1625
Premium User	1350

7.5 XenApp Test Results

7.5.1 Configuration

This validation was performed for XenDesktop 7.1 delivering hosted shared desktops on 2 R720 Compute servers where one was running vSphere ESXi 5.5 and the other Windows 2012 R2 Hyper-V. Each Compute Host was outfitted with 96 - 256 GB RAM (1600 MHz), dual 3.0 GHz Ivy Bridge E5-2690v2 processors, and 10 300GB 15K RPM SAS drives in a RAID 10 configuration. Each host comes with Broadcom 5720 4-port NICs running 1Gb Ethernet. 2 NICs were connected on each Compute Host and were teamed. All of the hosted shared desktops were hosted on the local disk storage on the Compute Hosts.

The Management Host was a separate host, an M620 running vSphere ESXi 5.5 with the same hardware configuration as the Compute Hosts. The Management Host used an EqualLogic PS4100E storage array connected via iSCSI on a Power Connect 8024F 1GbE network switch to store all of the management server VMs. A common Storefront, Desktop Delivery controller, and License server was used. The Management node also hosted the SCVMM virtual machine and VMware vCenter appliance. The Management VMs for the solution were deployed on the Management Host using the following configurations:

- 1 VM, 4 vCPU and 8 GB RAM for XenDesktop 7.1 DDC
- 1 VM, 2 vCPU and 4 GB RAM for Storefront and Licensing Server
- 1 VM, 4vCPU and 16 GB RAM for PVS 7.1 (powered off during MCS testing)

- 1 VM, 2 vCPU and 8 GB RAM each for VMware vCenter 5.5 appliance or Microsoft System Center 2012 R2
- 1 VM, 4 vCPU and 8 GB RAM for Microsoft SQL Server 2012 SP1

The XenApp images were built using Windows Server 2012 R2 Standard edition and Microsoft Office 2010 plus the support tools required for Login VSI testing.

The PVS desktops were each configured to boot from a vDisk hosted on the PVS Server VM. Each PVS VM had a local write cache disk attached that was sized at 32GB. The write cache disks were created to use thin provisioning in all of the test cases below, and were formatted with NTFS. In PVS the vDisks were each configured to use the target device's local disk for the write cache ("Cache on device hard drive"). In addition further tests were run on Hyper-V with the Write cache set to RAM plus overflow to local hard disk.

7.5.2 Test Results Summary

Validation was performed using DWD Solution Engineering standard testing methodology using LoginVSI 4 load generation tool for VDI benchmarking that simulates production user workloads. Stratusphere UX was not used in this validation since it is not compatible with the server OS in use here (Windows 2012 R2 Standard).

The following scenarios were tested to identify the optimal configurations:

- Virtual XenApp server optimization on both vSphere and Hyper-V to discover the right mix of vCPUs, oversubscription, RAM and total number of VMs per host
- Physical Windows 2012 R2 host running XenApp
- The performance impact and benefit of NUMA enabled
- The performance impact of various provisioning mechanisms: MCS, PVS write cache to disk, PVS write cache to RAM
- The performance impact of an increased user idle time to simulate a less than 80+% concurrency of user activity on any given host.

These scenarios were tested using the Login VSI light and medium workloads. Medium and heavy user workloads are likely better suited to full VDI desktops. The results presented in this document for virtual machine configurations were optimized for 6 x XenApp Server VMs with 5 vCPU and 16GB of memory for each VM. This configuration was found to give slightly improved results compared to other configurations using either 2 or 3 VMs with varying CPU and memory configurations. In all cases the R720 host was configured to allow NUMA operation. With 20 physical CPU cores we followed the 1.5x oversubscription model and assigned vCPUs according to 30 logical cores.

Each test run was configured to allow for 16 minutes of steady state operations after the initial session launch and login period, so that all sessions would be logged in and active for a significant period of time so that statistics could be measured during steady state to simulate daily operations.

In the CPU charts below it can be seen that the max CPU usage on the Hyper-V solution was reached before all the sessions logged in but did not rise further as more sessions logged in. This pattern is repeated in other Hyper-V tests for Standard and Enhanced workloads on MCS and PVS.

Testing of the physical server showed a percentage of users that did not complete the logoff process (approximately 15%) which was also tested using only the Login VSI medium workload (Enhanced). The light workload was not successful. Based on the results we find using virtual XenApp servers to be more scalable and higher performing than XenApp running on physical Windows servers, generally speaking. The following table summarizes the steady state test results for the various workloads and configurations. The charts in the sections below show more details about the max values for CPU, memory, IOPS, and Network during the test run.

Hyper-visor	Provision-ing	Workload	Density	Avg CPU %	Avg Memory Usage GB	Avg IOPS/User	Avg Net KBps/User
ESX	MCS	Standard	250	90% *	67	2.1	113
		Enhanced	170	88% *	72	3	138
		Dell Light **	280	85% *	91	1.6	58
ESXi	PVS Cache on disk	Standard	275	86%	101	4.9	116
		Enhanced	200	89%	118	7	138
		Enhanced without NUMA	200	90%	120	7.6	144
Hyper-V	PVS Cache on local disk	Standard	300	73%	107	5	133
		Enhanced	225	76%	107	8	140
		Dell Light **	300	54%	107	4	83
Hyper-V	PVS Cache in RAM	Standard	300	70%	203	0†	126
		Enhanced	225	74%	204	0†	155
Hyper-V	MCS	Standard	300	67%	92	2.3	130
		Enhanced	225	73%	107	4.5	151
Physical		Standard	200	85%	70	2	NA

* CPU % for ESX Hosts was adjusted to account for the fact that Intel E5-2600v2 series processors with the Turbo Boost feature enabled will exceed the ESXi host CPU metrics of 100% utilization. With E5-2690v2 CPUs the rated 100% in ESXi is 60000 MHz of usage, while actual usage with Turbo has been seen to reach 67000 MHz in some cases. The Adjusted CPU % Usage is based on 100% = 66000 MHz usage and is used in all charts for ESXi to account for Turbo Boost. Windows Hyper-V metrics by comparison do not report usage in MHz, so only the reported CPU % usage is used in those cases.

** The "Dell Light" workload is a modified VSI workload to represent a significantly lighter type of user. In this case the workload was modified to produce about 50% idle time.

†Avg IOPS is 0 because it is offloaded to RAM.

Summary of configuration recommendations:

- Enable Hyper-Threading and Turbo for oversubscribed performance gains.

- NUMA did not show to have a tremendous impact enabled or disabled.
- Virtual XenApp servers outperform dedicated physical hosts with no hypervisor so we recommend virtualized XenApp instances.
- Using 10-Core Ivy Bridge CPUs, we recommend running 6 x XenApp VMs per host, each VM assigned 5 x vCPUs and 16GB RAM.
- 1.5x CPU oversubscription per host produced excellent results. 20 physical cores x 1.5 oversubscription netting 30 logical vCPUs assigned to VMs.
- PVS cache in RAM (with HD overflow) will reduce the user IO generated to disk almost nothing but will require greater RAM densities on the compute hosts. 256GB is a safe high water mark using PVS cache in RAM based on a 21GB cache per XenApp VM.

CPU utilization was higher on the ESXi hosts compared to what was observed on the equivalent Hyper-V solution.

Memory utilization remains low for all solutions except PVS solution with write cache in RAM. This provides a significant saving over VDI solutions that typically require 256GB to provide lower user densities. Utilizing PVS write cache in RAM provides significant performance improvements and practically eliminates the need for tier 1 storage on local disk.


The IOPS results above were surprising in that PVS typically would offload IO to the network, relieving the local storage from a significant number of IOPS. In the chart above however it shows that PVS test cases consistently produced more IOPS per user than equivalent MCS test cases. As expected PVS with write cache in RAM showed practically no tier 1 IOPs.

Network Utilization was consistent between the solutions. It was expected that the PVS solution would show higher IOPs than MCS or physical solutions. However, this was not depicted in the results gathered. This is probably due to caching. Network utilization was higher for the equivalent Hyper-V solution than the ESXi solution for MCS. As expected the PVS solution produced significantly higher network utilization overall.


7.5.2.1 User Experience

The results show significant differences in user density achievable with Hyper-V compared to ESXi. Unfortunately Stratusphere UX was not available to gather User experience data but a subjective test of user experience was gathered in each test run. In all cases the user experiences was good. The subjective user test is performed when approximately 95% of users were logged in. Each test consists of browsing the internet, performing some tasks on Word, Excel, PowerPoint and playing some multimedia clips. The following graphs from VSI Analyzer indicate that the VSI tool does indeed indicate a difference in capability between the Hyper-V and ESXi solutions. The following graphs were taken for the VSI light workload on ESXi and Hyper-V deployed using MCS technology.

Hyper-V Standard workload for 300 Users



ESXi Standard workload for 250 users



The same image was used for both ESXi and Hyper-V testing except that the appropriate support tools were applied for each Hypervisor. The R720 hosts were configured identically in terms of processor, memory, BIOS, local tier1 disk and network. The performance data gathered does not indicate stress in any other area, e.g. local disk latency remains low in both solutions (≤ 5 ms).

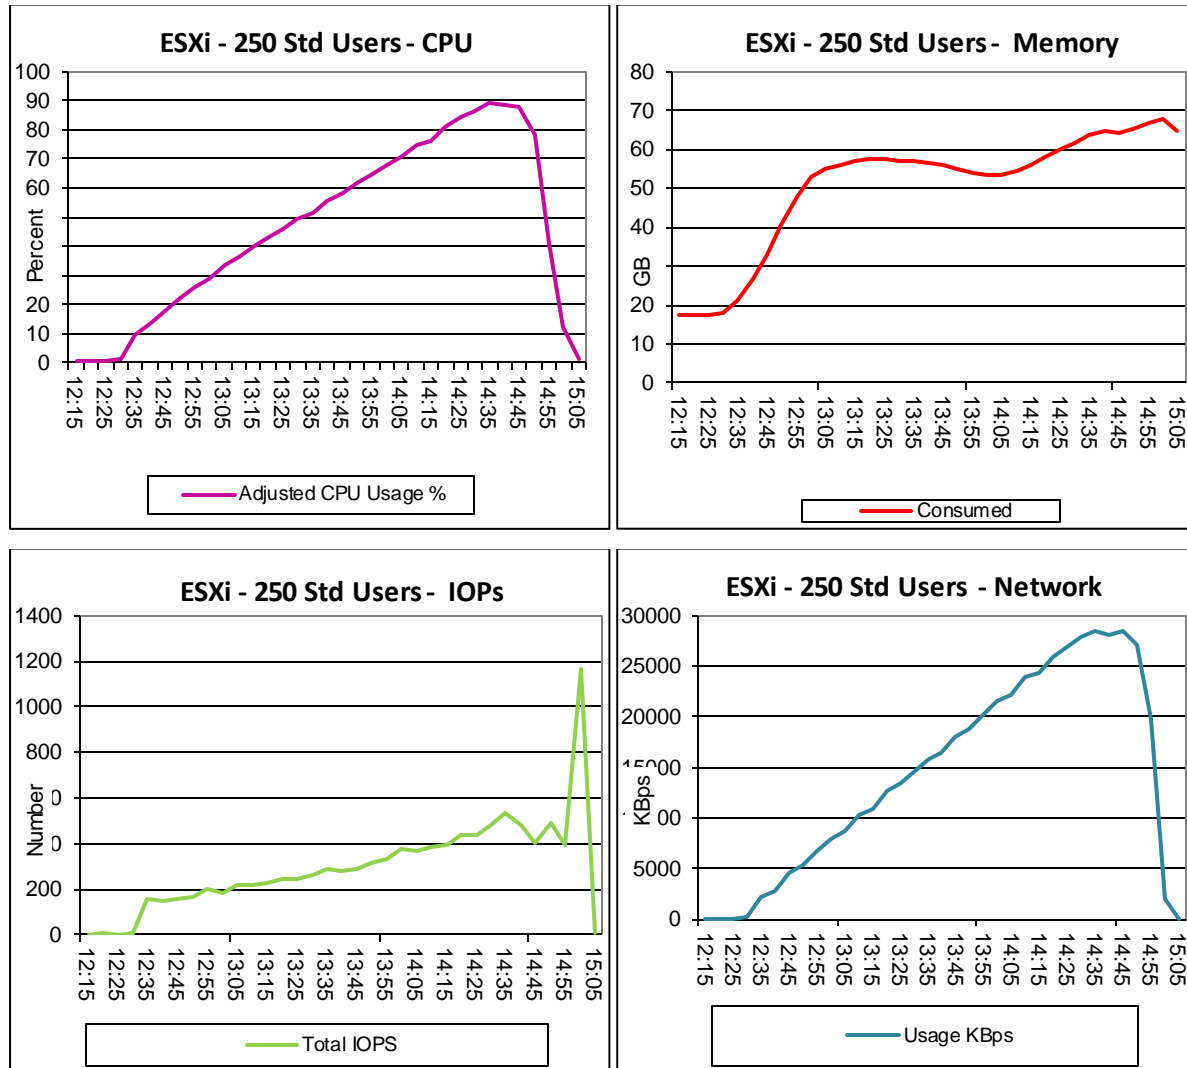
In both cases above the VSIMAX chart indicates that the solution is approaching its maximum capability. Overall, the PVS solution with write cache in RAM indicates a very viable solution provided the volume of changes in each session doesn't exceed the available memory. In any event this can be configured with an overflow to local hard disk.

7.5.3 VMware vSphere 5.5 Test Charts

These sections details the test results charts on an ESXi 5.5 host.


7.5.3.1 MCS Standard User Workload (250 Users)

These graphs show CPU, memory, local disk IOPS, and network results.




7.5.3.2 MCS Enhanced User Workload (170 users)

These graphs show CPU, memory, local disk IOPS, and network results.




7.5.3.3 MCS Dell Extra Idle time User Workload (280 users)

These graphs show CPU, memory, local disk IOPS, and network results.




7.5.3.4 PVS Enhanced User Workload (200 Users)

These graphs show CPU, memory, local disk IOPS, and network results



7.5.3.5 PVS standard User Workload (275 Users)

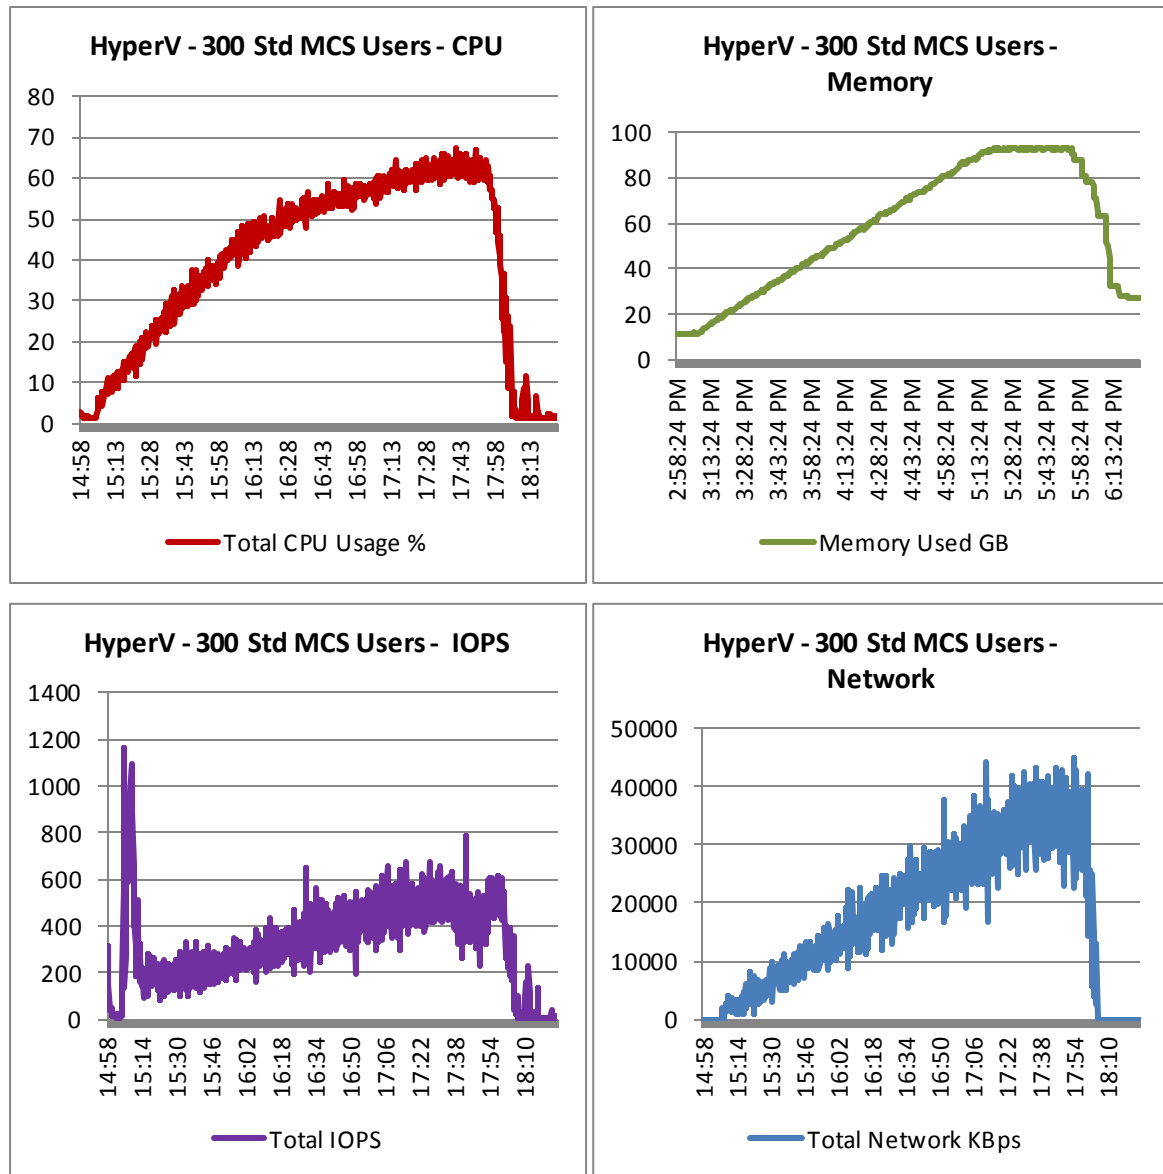
These graphs show CPU, memory, local disk IOPS, and network results:



7.5.4 Microsoft Windows 2012 R2 Hyper-V Test Charts

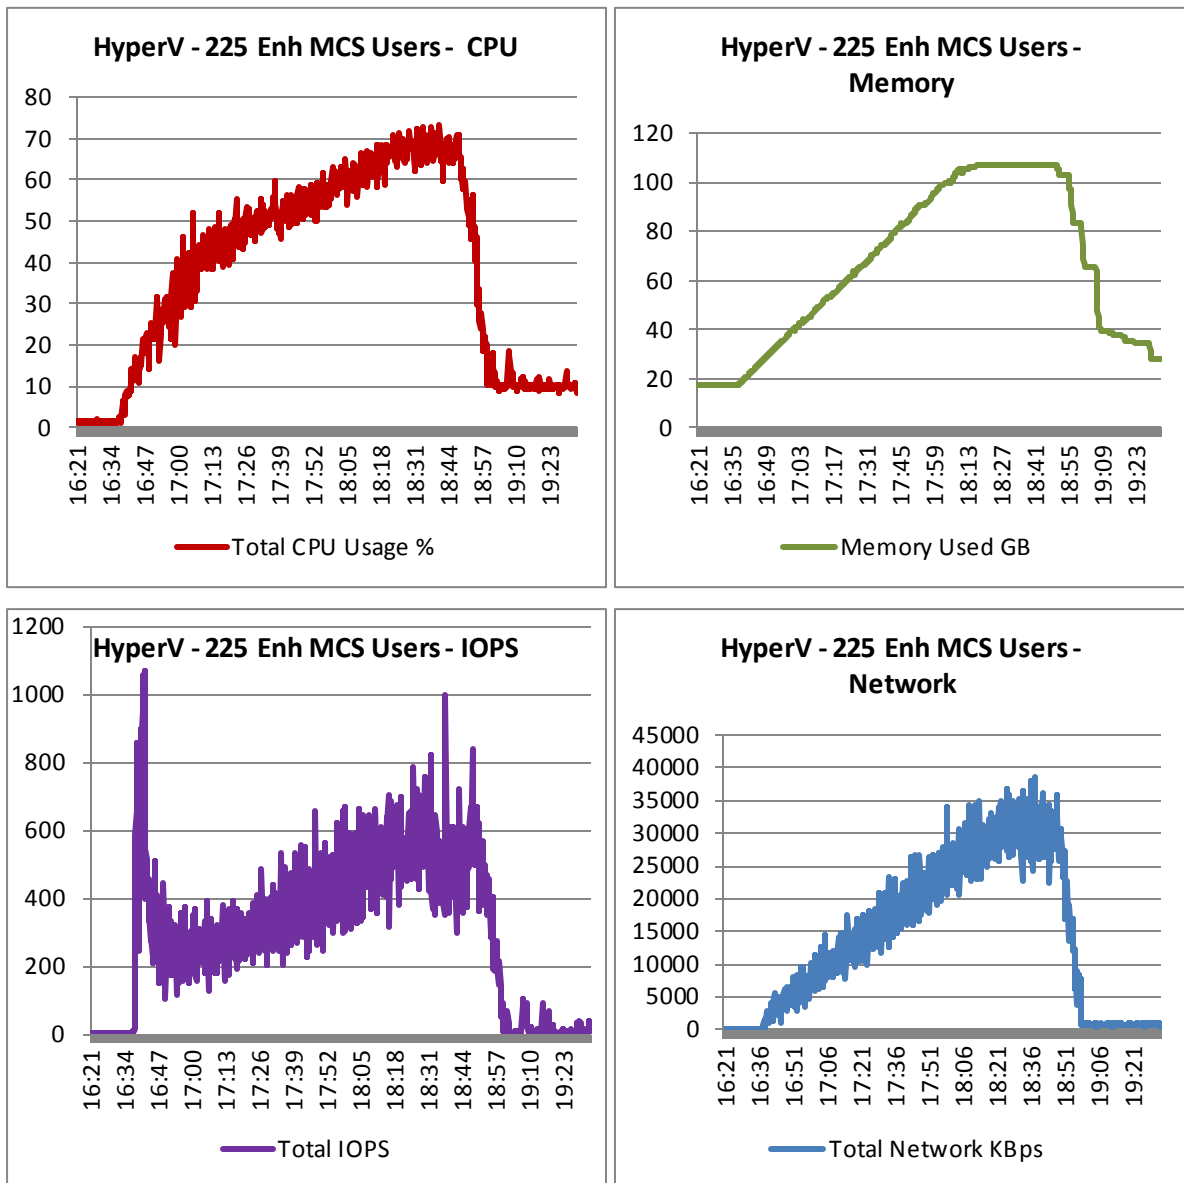
7.5.4.1 MCS Standard User Workload (300 Users)

These graphs show CPU, memory, local disk IOPS, and network results.



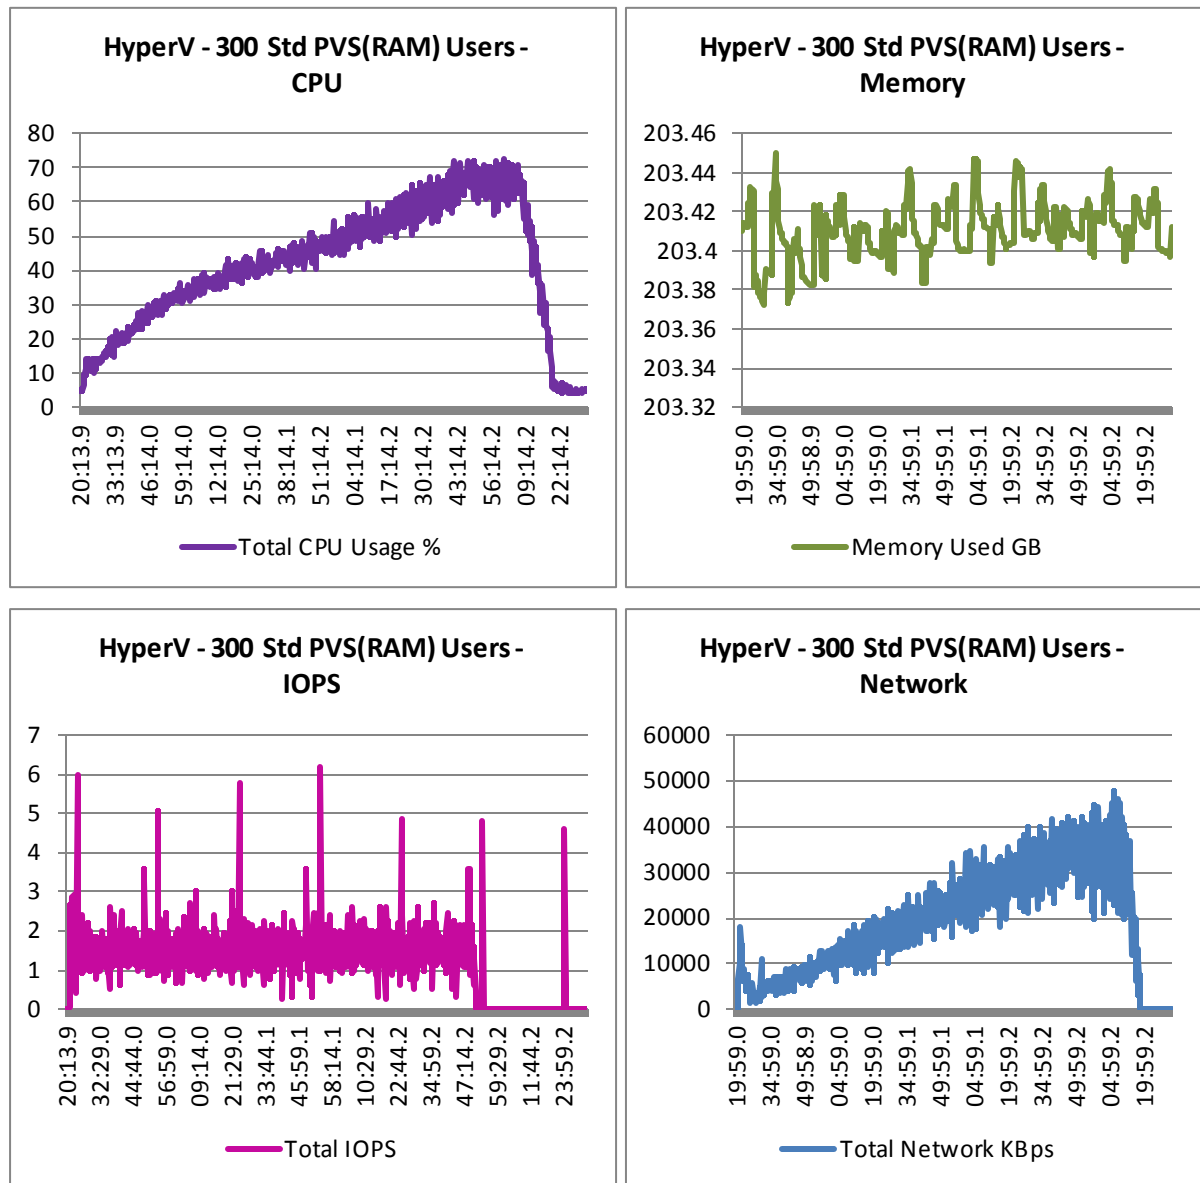
7.5.4.2 MCS Enhanced User Workload (225 users)

These graphs show CPU, memory, local disk IOPS, and network results.



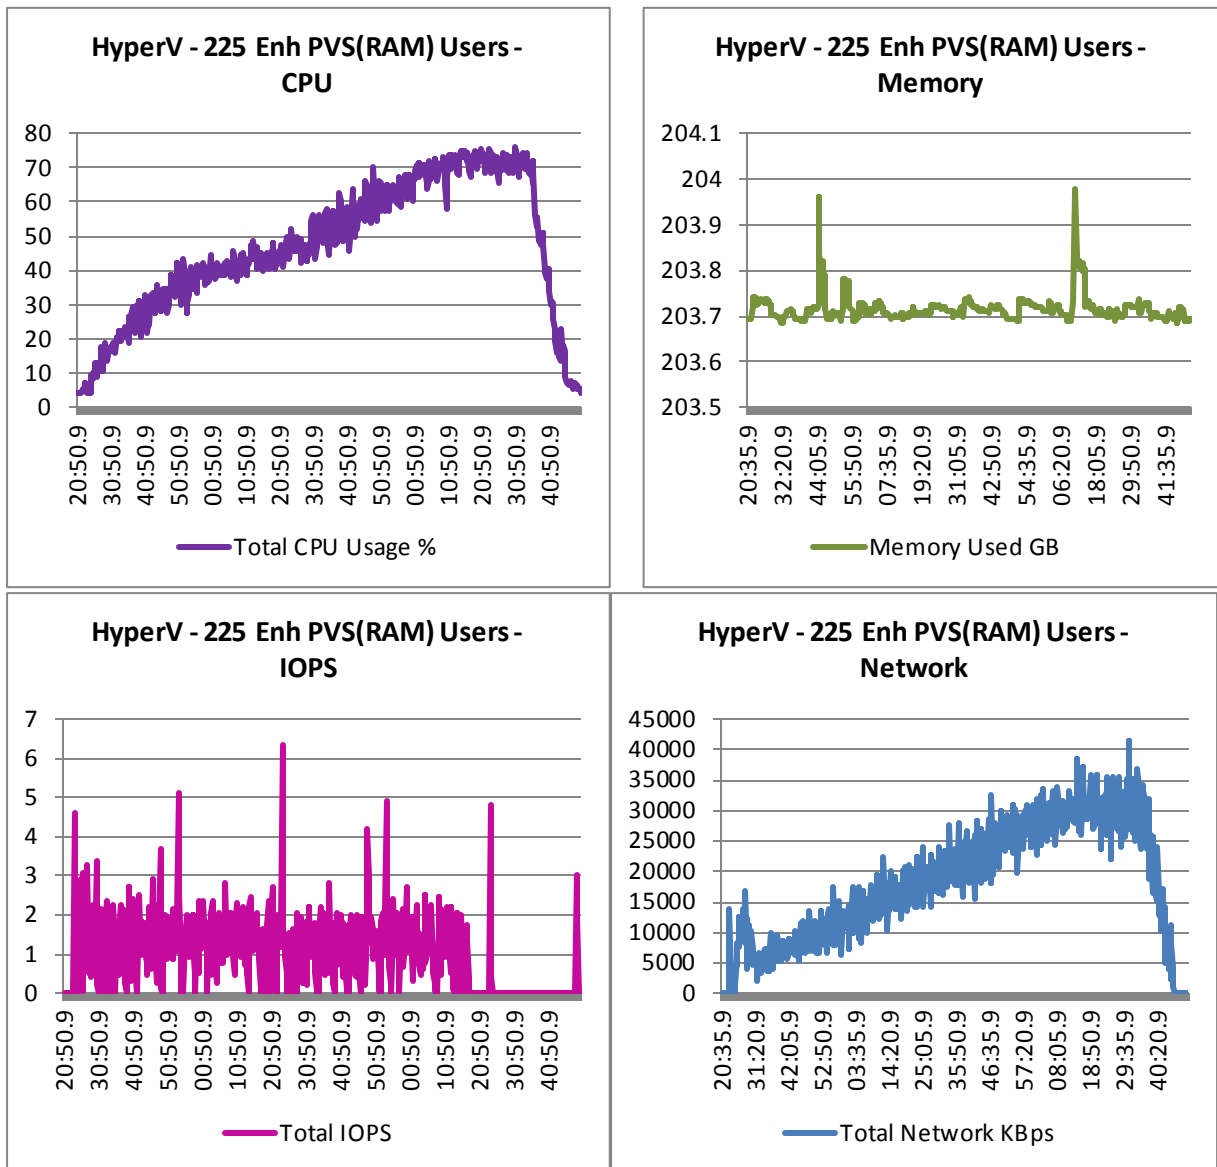
7.5.4.3 PVS Standard User Workload with Write cache in RAM (300 Users)

These graphs show CPU, memory, local disk IOPS, and network results.




7.5.4.4 PVS Enhanced User Workload with Write cache in RAM (225 users)

These graphs show CPU, memory, local disk IOPS, and network results.



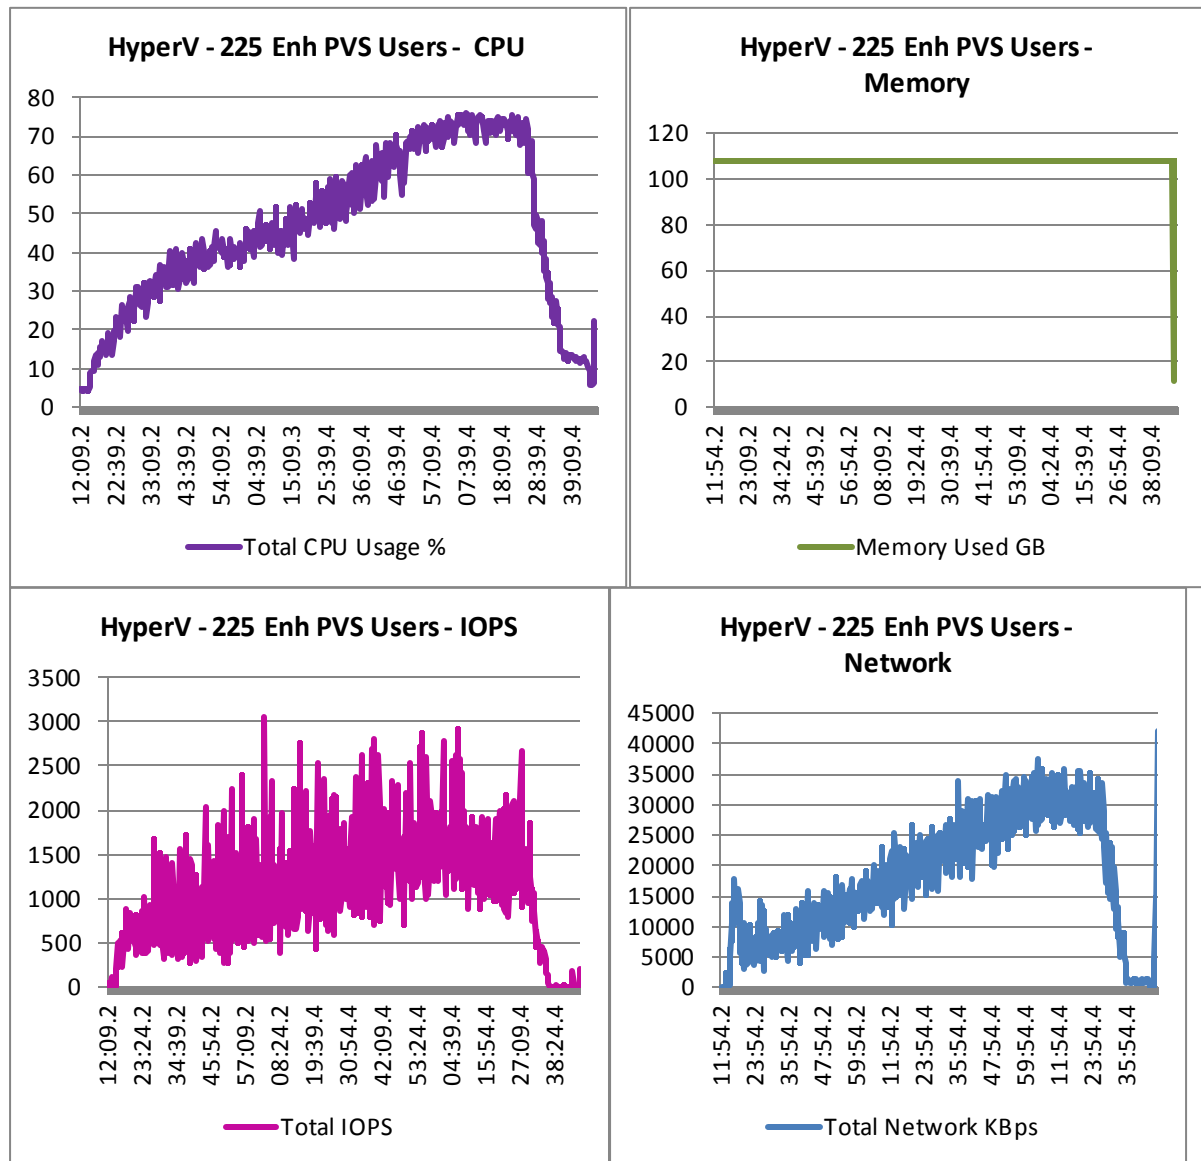
7.5.4.5 PVS Standard User Workload write cache on local disk (300 Users)

These graphs show CPU, memory, local disk IOPS, and network results.




7.5.4.6 PVS Enhanced User Workload write cache on local disk (225 Users)

These graphs show CPU, memory, local disk IOPS, and network results.



7.5.4.7 PVS Dell Modified Extra Light Workload (300 Users)


These graphs show CPU, memory, local disk IOPS, and network results.



7.5.5 Microsoft Windows 2012 R2 Physical Server

7.5.5.1 Enhanced Workload (200 Users)

These graphs show CPU, memory and local disk IOPS



7.5.6 XenApp Scaling Guidance

7.5.6.1 vSphere

Rack or Blade, Server 2012 R2/ XenApp, vSphere, PVS									
Standard User Count	Enhanced User Count	Physical Mgmt Servers	Physical Host Servers	M1000e Blade Chassis	Virtual PVS Servers	Virtual Delivery Controllers	Virtual License Servers	Virtual StoreFront Servers	Virtual vCenter Server
275	200	1	1	1	1	1	1	1	1
1000	800	1	4	1	1	1	1	1	1
2000	1600	2	8	1	2	1	1	1	1
3000	2200	3	11	1	3	1	1	1	1
4000	3000	3	15	2	4	1	1	1	1
5000	3800	4	19	2	5	1	1	1	1
6000	4400	5	22	2	5	1	1	1	1
7000	5200	5	26	2	6	1	1	1	1
8000	6000	6	30	3	7	1	1	1	1
9000	6600	7	33	3	8	1	1	1	1
10000	7400	7	37	3	9	1	1	1	1

Rack or Blade, Server 2012 R2/ XenApp, vSphere, MCS								
Standard User Count	Enhanced User Count	Physical Mgmt Servers	Physical Host Servers	M1000e Blade Chassis	Virtual Delivery Controllers	Virtual License Servers	Virtual StoreFront Servers	Virtual vCenter Server
250	170	1	1	1	1	1	1	1
1000	680	2	4	1	1	1	1	1
2000	1360	2	8	1	1	1	1	1
3000	2040	2	12	1	1	1	1	1
4000	2720	2	16	2	1	1	1	1
5000	3400	3	20	2	1	1	1	1

7.5.6.2 Hyper-V

Rack or Blade, Server 2012 R2/ XenApp, Hyper-V, PVS									
Standard User Count	Enhanced User Count	Physical Mgmt Servers	Physical Host Servers	M1000e Blade Chassis	Virtual PVS Servers	Virtual Delivery Controllers	Virtual License Servers	Virtual StoreFront Servers	Virtual SCVMM Server
300	225	1	1	1	1	1	1	1	1
1000	900	2	4	1	1	1	1	1	1
2000	1575	2	7	1	2	1	1	1	1
3000	2250	2	10	1	3	1	1	1	1
4000	3150	2	14	1	4	1	1	1	1
5000	3825	3	17	2	5	1	1	1	1
6000	4500	5	20	2	5	1	1	1	1
7000	5400	5	24	2	6	1	1	1	1
8000	6075	6	27	3	7	1	1	1	1
9000	6750	7	30	3	8	1	1	1	1
10000	7650	7	34	3	9	1	1	1	1

Rack or Blade, Server 2012 R2/ XenApp, Hyper-V, MCS								
Standard User Count	Enhanced User Count	Physical Mgmt Servers	Physical Host Servers	M1000e Blade Chassis	Virtual Delivery Controllers	Virtual License Servers	Virtual StoreFront Servers	Virtual SCVMM Server
300	225	1	1	1	1	1	1	1
1000	900	2	4	1	1	1	1	1
2000	1575	2	7	1	1	1	1	1
3000	2250	2	10	1	1	1	1	1
4000	3150	2	14	1	1	1	1	1
5000	3825	3	17	2	1	1	1	1

Acknowledgements

Thanks to the Darin Schmitz and Damon Zaylskie of the Dell Compellent MSV Solutions team for providing expertise and validation of the Dell Wyse Solutions Engineering Compellent Tier 1 array.

Thanks to the Dell Solution Centers in Austin and Limerick for their support and passionate evangelism for VDI.

Thanks to Paul Wynne and the Dell Wyse Solutions IXT team for their expertise and continued support validating VDI architectures and Tier 1 storage platforms.

Thanks to John Kelly of the Dell Wyse Solutions Engineering team for his expertise and guidance in the Dell Wyse Solutions Engineering PAAC process.

Thanks to Scott Stanford of the Dell Wyse Solutions Engineering team for his attention to detail and enterprise architecture expertise.

About the Authors

Peter Fine is the Sr. Principal Solutions Architect for 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 Citrix solutions at Dell, managing the development and delivery of Enterprise class Desktop virtualization solutions based on Dell Data center components and core virtualization platforms.

Cormac Woods is a Sr. Solution Engineer in the Dell Wyse Solutions Engineering Solutions Group at Dell building, testing, validating, and optimizing enterprise VDI stacks.

Geoff Dillon is a Sr. Solutions Engineer in the Dell Wyse Solutions Engineering Group at Dell with deep Citrix experience and validation expertise of Dell's Dell Wyse Datacenter VDI solutions.

Pranav Parekh is a Sr. Solutions Engineer in the Dell Wyse Solutions Engineering group at Dell. Pranav has extensive experience designing desktop virtualization solutions, IaaS private cloud solutions, virtualization solutions, and enterprise class blade servers. Pranav has a master's degree in Electrical & Computer Engineering from the University of Texas at Austin.